



Exam : 642-801

Title : Building Scalable Cisco Internetworks (BSCI)

Ver : 09-27-07

**QUESTION 1:**

A packet that needs to be forwarded arrives on an interface of a router. In order for a router to route data, what must that router determine? (Select the best answer)

- A. The route age of the next-hop device
- B. The subnet mask of the source network
- C. The cost metric of the path of the destination
- D. The outbound interface of the best path to the destination
- E. All of the above

Answer: D

Explanation:

In order to route the data to the correct destination the router must know the outbound interface that provides the best destination. When a receives a packet that needs to be forwarded, the router determines the destination network, looks up the best path to that destination in the routing table, and then forwards the packet out the correct interface.

Incorrect Answers:

- A. The age of the next-hop router is not consulted before forwarding information.
- B. The destination subnet mask is used by the router to determine the best path (most specific network match rule), but not the source subnet mask.
- C. The cost metric is used by routing protocols to determine the best route to a destination. This best route is then installed into the routing table. The router uses the routing table to forward packets, but does not use the metrics for that specific route before forwarding each packet.

---

**QUESTION 2:**

A new router is being installed into an existing network, and the routing table is being built for the first time on this network. Which of the following statements is true regarding the routing tables on a Cisco router?

- A. Entries are listed in the order of the route cost metric.
- B. Only the active link is shown for load-balanced routers.
- C. Privileged EXEC mode is required to view the routing table.
- D. The clear ip route \* command refreshes the entire routing table.
- E. All of the above are true.

Answer: C

Explanation:

The show ip route command is used to display IP routing table entries. It can only be run in privileged (enable) mode.

Incorrect Answers:

- A. Entries are not listed in route cost order. Entries are ordered by destination IP address.
  - B. The other links in the load balanced bundle are also shown.
  - D. The clear ip route command is used to clear delete IP routing table entries. In particular the clear ip route \* (or clear ip route all) command deletes IP routing table entries. This is not a refresh. Static routes will not be recreated.
- 

**QUESTION 3:**

Which of the following phrases is the correct term for what happens to a network when a topology change causes all the routers to synchronize their routing tables?

- A. Flooding
- B. Broadcasting
- C. Convergence
- D. Summarization
- E. None of the above

Answer: C

Explanation: A topology change forces a convergence of the routers to a new routing state. Convergence is the time that it takes for all routers to agree on the network topology after a change.

Incorrect Answers:

- A. Flooding is the process of sending out routing topology information used by link state protocols.
  - B. Broadcasting is used on LAN level communications. By default, routers do not forward broadcasts.
  - D. Summarization is the act of taking multiple routes within the routing table, and advertising them as one less specific route.
- 

**QUESTION 4:**

What is the purpose of configuring a router with the "IP Helper address" command?

- A. IP Helper is used to direct BOOTP clients to a BOOTP server.
- B. IP Helper is used to prevent the router from forwarding IP broadcasts.
- C. IP Helper is used to allow IPX clients to communicate with IP-based servers.
- D. IP Helper is used to accommodate compatibility routers using different IP routing protocols.

Answer: A

Explanation:

The ip helper-address command is used to have the Cisco IOS software forward User

Datagram Protocol (UDP) broadcasts, including BOOTP, received on an interface. DHCP protocol information is carried inside of BOOTP packets. To enable BOOTP broadcast forwarding for a set of clients, configure a helper address on the router interface closest to the client. The helper address should specify the address of the DHCP server.

Note: A DHCP server can be considered to be a BOOTP server, even though a DHCP server is more advanced.

Incorrect Answers:

B: Combined with the `ip forward-protocol` global configuration command, the `ip helper-address` command allows you to control which broadcast packets and which protocols are forwarded. However, the main purpose of the IP helper feature is not to prevent the router from forwarding IP broadcasts.

C: IP helper does not use IPX.

D: This is false.

---

### QUESTION 5:

On router CK1 the command "ip helper address" is already configured. Which of the following commands would you use if you wanted to send SNMP broadcast packets off to a specific server?

- A. `ip server udp 161`
- B. `ip helper-protocol 161`
- C. `ip forward-protocol 161`
- D. `ip directed-broadcast 161`
- E. `ip forward snmp`

Answer: C

Explanation: SNMP requests are typically sent to UDP port 161.

The `ip forward-protocol` command is used to specify which protocols and ports the router forwards when forwarding broadcast packets. If an IP helper address is defined, UDP forwarding is enabled on default ports.

Note: Syntax: `ip forward-protocol {udp [port] | nd | sdns}`

Incorrect Answers:

A: The use of `ip server udp` is not correct.

B: There is no such command.

D: The `ip directed-broadcast` command is used to enable the translation of directed broadcast to physical broadcasts. It does not apply in this scenario. Furthermore, an access-list with number 161 has to be configured.

Syntax: `ip directed-broadcast [access-list-number] | [extended access-list-number]`

---

### QUESTION 6:

When you execute the "ip helper-address" command on a router, which three UDP ports get enabled automatically by default? (Select three)

- A. 53 (DNS)
- B. 69 (TFTP)
- C. 515 (LPR)
- D. 161 (SNMP)
- E. 49 (TACACS)

Answer: A, B, E

Explanation:

To forward the BootP/DHCP request from the client to the DHCP server, the ip helper-address interface command is used. The IP helper-address can be configured to forward any UDP broadcast based on UDP port number. By default, the IP helper-address will forward the following UDP broadcasts:

1. DNS (port 53), time service (port 37)
2. Trivial File Transfer Protocol (TFTP) (port 69)
3. Terminal Access Control Access Control System (TACACS) service (port 49)
4. NetBIOS name server (port 137)
5. NetBIOS datagram server (port 138)
6. Boot Protocol (DHCP/BootP) client and server datagrams (ports 67 and 68)
7. IEN-116 name service (port 42)

Reference: Understanding and Troubleshooting DHCP in Catalyst Switch or Enterprise Networks

<http://www.cisco.com/warp/public/473/100.html>

---

### **QUESTION 7:**

Which administrative distance is given to EIGRP summary routes?

- A. 0
- B. 1
- C. 5
- D. 90
- E. 95
- F. 150

Answer: C

Explanation:

The following table displays the default AD for all routing protocols:

| Default Administrative Distances |                  |
|----------------------------------|------------------|
| Route Source                     | Default Distance |
| Connected interface              | 0                |
| Static route                     | 1                |

|                             |     |
|-----------------------------|-----|
| Enhanced IGRP summary route | 5   |
| External BGP                | 20  |
| Internal Enhanced IGRP      | 90  |
| IGRP                        | 100 |
| OSPF                        | 110 |
| IS-IS                       | 115 |
| RIP                         | 120 |
| EGP                         | 140 |
| EIGRP external route        | 170 |
| Internal BGP                | 200 |
| Unknown                     | 255 |

---

**QUESTION 8:**

If there is a route to the subnet 190.10.1.0/25 learned via RIP and a route to the subnet 190.10.1.0/24 learned via STATIC, which route would be preferred to reach the destination address 190.10.1.125?

- A. The 190.10.1.0/25 route learned via RIP will be used because its prefix has the longest match.
- B. RIP has an administrative distance of 120, and static routes have an administrative distance of 1, so the static route would be preferred.
- C. The static route to 190.10.1.0/24 will be preferred because static routes have an administrative distance of 0 and the static route looks as though it is directly connected.
- D. A show ip route to the destination will show that the destination is learned from both RIP and the static route, so the traffic to 190.10.1.125 will be load balanced between the two paths.
- E. The information given is not sufficient to determine this.

Answer: A

Explanation:

The longest match rule always takes precedence over any other routing information. The administrative distance (AD) values are only compared for destinations with subnet masks of equal length.

---

**QUESTION 9:**

What does a spoke router require to be enabled to ensure proper functionality of on-demand routing?

- A. dynamic routing

- B. static routing
- C. CDP
- D. broadcasting

Answer: C


Explanation:

On-Demand Routing (ODR) is an enhancement to Cisco Discovery Protocol (CDP), a protocol used to discover other Cisco devices on either broadcast or non-broadcast media. With the help of CDP, it is possible to find the device type, the IP address, the Cisco IOS(r) version running on the neighbor Cisco device, the capabilities of the neighbor device, and so on. In Cisco IOS software release 11.2, ODR was added to CDP to advertise the connected IP prefix of a stub router via CDP. This feature takes an extra five bytes for each network or subnet, four bytes for the IP address, and one byte to advertise the subnet mask along with the IP. ODR is able to carry Variable Length Subnet Mask (VLSM) information.

---

**QUESTION 10:**

The Certkiller WAN is displayed in the following diagram:



When performing a show run on router Certkiller 1, you notice that two static entries shown above were configured for the 10.0.0.0 network.

Which path will be selected by the router as the best path?

- A. Per-destination load sharing will be implemented.
- B. The path through Certkiller 2 will be used.
- C. The path through Certkiller 3 will be used.
- D. Per-packet load sharing will be implemented.

Answer: B

Explanation:

A static route configured using the interface is considered to be a directly connected route. Directly connected routes have an administrative distance of 0, while all other

static routes have an AD of 1. In this case, the route specifying the interface will be chosen over a static route that specifies the IP address of the next hop, because of the lower AD.

---

**QUESTION 11:**

How many update packets would a RIPv2 router send to advertise 77 routes?

- A. one update package
- B. two update package
- C. three update package
- D. four update package
- E. five update package
- F. six update package

Answer: D

Explanation:

RIP 2 Packet Format:

The RIP 2 specification (described in RFC 1723) allows more information to be included in RIP packets and provides a simple authentication mechanism that is not supported by RIP. Figure 47-2 shows the IP RIP 2 packet format.

Figure 47-2: An IP RIP 2 Packet Consists of Fields Similar to Those of an IP RIP Packet

|                             |                                       |                            |                         |                                  |  |                                    |                                 |                            |
|-----------------------------|---------------------------------------|----------------------------|-------------------------|----------------------------------|--|------------------------------------|---------------------------------|----------------------------|
| 1-octet<br>command<br>field | 1-octet<br>version<br>number<br>field | 2-octet<br>unused<br>field | 2-octet<br>AFI<br>field | 2-octet<br>route<br>tag<br>field | 4-octet<br>network<br>address<br>field | 4-octet<br>subnet<br>mask<br>field | 4-octet<br>next<br>hop<br>field | 4-octet<br>metric<br>field |
|-----------------------------|---------------------------------------|----------------------------|-------------------------|----------------------------------|--|------------------------------------|---------------------------------|----------------------------|

The following descriptions summarize the IP RIP 2 packet format fields illustrated in Figure 47-2:

1. Command-Indicates whether the packet is a request or a response. The request asks that a router send all or a part of its routing table. The response can be an unsolicited regular routing update or a reply to a request. Responses contain routing table entries. Multiple RIP packets are used to convey information from large routing tables.
2. Version-Specifies the RIP version used. In a RIP packet implementing any of the RIP 2 fields or using authentication, this value is set to 2.
3. Unused-Has a value set to zero.
4. Address-family identifier (AFI)-Specifies the address family used. RIPv2's AFI field functions identically to RFC 1058 RIP's AFI field, with one exception: If the AFI for the first entry in the message is 0xFFFF, the remainder of the entry contains authentication information. Currently, the only authentication type is simple password.
5. Route tag-Provides a method for distinguishing between internal routes (learned by RIP) and external routes (learned from other protocols).
6. IP address-Specifies the IP address for the entry.
7. Subnet mask-Contains the subnet mask for the entry. If this field is zero, no subnet mask has been specified for the entry.
8. Next hop-Indicates the IP address of the next hop to which packets for the entry should be forwarded.
9. Metric-Indicates how many internetwork hops (routers) have been traversed in the trip

to the destination. This value is between 1 and 15 for a valid route, or 16 for an unreachable route.

Note Up to 25 occurrences of the AFI, Address, and Metric fields are permitted in a single IP RIP packet. That is, up to 25 routing table entries can be listed in a single RIP packet. If the AFI specifies an authenticated message, only 24 routing table entries can be specified. Given that individual table entries aren't fragmented into multiple packets, RIP does not need a mechanism to resequence datagrams bearing routing table updates from neighboring routers.

Since a single update packet can contain at most 25 route entries, 4 packets will be needed to advertise 77 routes.

Reference: [http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito\\_doc/rip.htm#xtocid9](http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/rip.htm#xtocid9)

---

### QUESTION 12:

What are three disadvantages of a router-on-stick configuration for InterVLAN routing? (Choose three)

- A. InterVLAN routing cannot be filtered by the router.
- B. The router becomes a single point of failure for the network.
- C. Routers will not route STP BPDUs.
- D. There is a possibility of inadequate bandwidth for each VLAN.
- E. Additional overhead on the router can occur.
- F. NetFlow Switching is required for InterVLAN accounting.

Answer: B, D, E

Explanation:

When a single interface is used to route between subnets or networks, this is known as a router-on-a-stick. To assign multiple IP addresses to the same interface, secondary addresses or subinterfaces are used.

The Advantage is that it is useful when there are limited Ethernet interfaces on the router.

The Disadvantages to this design are:

Because a single link is used to connect multiple subnets, one link has to carry the traffic for multiple subnets.

Be sure this link can handle the traffic. You may wish to use a high-speed link (100 Mbps) and full-duplex.

Because all Inter-VLAN traffic now must traverse the router, the router is a single point of failure.

Generally, routers do not process the datagrams as quickly as switches and the additional overhead and CPU resources that will be consumed on the router can become an issue.

---

### QUESTION 13:

You are determining the routing protocol to use throughout your network. In doing this you compare the advantages of classless and classful protocols. Which of the

following statements are true regarding classless routing protocols?

- A. A default gateway is required.
- B. Variable-length subnet masks are not supported.
- C. Routers are automatically summarized to class boundaries.
- D. Networks with different subnet masks can exist in the same address class.
- E. All of the above

Answer: D

Explanation:

Classless routing protocols understand that different routes within a major network can have different masks. They include VLSM information in the routing updates, enabling the use of different subnet mask lengths.

Incorrect Answers:

- A. There is no requirement for a default gateway.
- B. Classless routing protocols advertise the routing mask for each route. This enables VLSM (variable length subnet masks) support.
- C. With classless routing the summarization process can manually controlled and can be invoked at any point within the network. It is not limited to class boundaries.

---

**QUESTION 14:**

Which two of the following describe advantages of implementing a classless routing protocol, when compared to a classful routing protocol?

- A. Support for VLSM.
- B. Support for FLSM.
- C. Summarization of discontinuous subnets.
- D. Auto-summarization across network boundaries.
- E. The ip classless command improves convergence time.

Answer: A, C

Explanation:

- A. Classless routing protocols support VLSM, and that, in turn, leads to more efficient allocation of subnet masks to meet different host requirements on different subnetworks, resulting in better utilization of host addresses.
- C. Because subnets routes are propagated throughout the routing domain, summarization is often required to keep the routing tables at a manageable size.

Incorrect Answers:

- B, D: Fixed Length Subnet Masks (FLSM) and auto-summarized routes across network boundaries are functions of classful routing protocols, not classless.
- E. The convergence time of a network is due to numerous factors, including the timers of

the routing protocol, as well as support for triggered updates. The "ip classless" command has no impact on the convergence time of any network.

Reference: Building Scalable Cisco Networks (Cisco Press) page 19-20.

---

**QUESTION 15:**

You wish to use a classless IP routing protocol within your network. Which of the following classless routing protocols could you use? (Select all that apply)

- A. IS-IS
- B. IGRP
- C. RIPv1
- D. OSPF
- E. EIGRP

Answer:

A, D, E

Explanation:

Intermediate System to Intermediate System (IS-IS), Open Shortest Path First (OSPF) and Enhanced IGRP are all classless routing protocols.

Note: RIPv2 and BGP are also classless routing protocols.

Incorrect Answers:

B, C: IGRP and RIPv1 are not classless. Both of these protocols do not support VLSM networks.

---

**QUESTION 16:**

The Certkiller network administrator is considering using a link state routing protocol to replace their existing distance vector protocol. Which if the following are true regarding link state routing protocols?

- A. Link-state protocols do not support summarization.
- B. Static routes must be used to accommodate link redundancy.
- C. All routers in the area know when another router joins the area.
- D. Link-state protocols utilize spanning tree to propagate routing changes.
- E. The spanning tree protocol must be enabled for link state protocols to work.

Answer: C

Explanation:

All routers within an area will have the same view of the area, meaning they will all have the same topology table. All of them will know when another router joins the area, since the routing topology table of all routers in any area must remain identical.

Incorrect Answers

A: Link state protocols support route summarization, as well as variable length subnet

masking.

B: Link redundancy does not require the use of static routes. Redundant links can be used with link state protocols, and load balancing over redundant links is also supported.

D, E: The spanning tree is not to propagate routing changes. Instead it used to find the shortest path to the destinations. The spanning tree protocol is used at layer 2 to determine bridging and switching loops, not routing loops.

---

**QUESTION 17:**

IS-IS and OSPF are two examples of link state routing protocols. Regarding the operation of link state protocols, which of the following are true?

- A. Link state protocols periodically multicast the Link State Advertisements at 90 second intervals.
- B. Link state routing protocols use Hello packets to build the link state database.
- C. Link state routing protocols use poison reverse and holddown timer to prevent routing loops.
- D. Link state routing protocols use Link State Advertisements to announce route changes.
- E. All of the above

Answer: D

Explanation:

LSA - Link State Announcement. Used by OSPF, an LSA is used to announce changes in network topology to adjacent routers

Each router periodically sends an LSA to provide information on a router's adjacencies or to inform others when a router's state changes. By comparing established adjacencies to link states, failed routers can be detected quickly and the network's topology altered appropriately. From the topological database generated from LSAs, each router calculates a shortest-path tree, with itself as root. The shortest-path tree, in turn, yields a routing table.

Incorrect Answers:

A: .

B: The Hello Protocol is responsible for establishing and maintaining neighbor relationships. It is used to build the neighbor table, not the link state database.

C: These are loop avoidance mechanisms used by distance vector routing protocols, not link state.

---

**QUESTION 18:**

One of the serial links in an OSPF network is experiencing problems and continuously changes from up to down. In a link-state environment, what does the router do when a route flaps?

- A. It enters the exstart state with its neighbors.
- B. It floods the area with new routing information.

- C. It generates a routing exchange using the hello protocol.
- D. It waits for the holdown timers to expire and then sends an update.

Answer: B

Explanation:

One of the drawbacks of using OSPF is that it can be prone to issues resulting in flapping routes. OSPF is a very strong protocol in terms of convergence time-each router is aware of the entire topology in the area. This results in fast convergence. However, if a link flaps, or changes between up and down status quickly, a flood of LSAs may be generated. This may prevent the routers in the network from converging. Administrators may use the "spf holdtime" command to force OSPF into a waiting state before computing a new route.

---

**QUESTION 19:**

You are an administrator of a network segment that uses RIP v2 as the routing protocol. How would the core router react if it were to detect a flapping link to a neighboring router?

- A. It recalculates the network topology.
- B. It purges that link from its routing table.
- C. It places a hold-down on the routes from that link.
- D. It sends a LSA to other router requesting an RIP update.

Answer: C

Explanation:

RIP is a distance vector routing protocol, which utilizes hold down timers and the rule of split horizons. The purpose of the hold-down state is to ensure the validity of any new routes for the same destination.

Incorrect Answers:

- A: RIP does not calculate a network topology.
- B: This would not serve any practical purpose, since the purged link would then reappear. The link would be placed in a held state, but not purged.
- D: OSPF utilizes link state advertisements (LSAs), but RIP routers do not.

---

**QUESTION 20:**

Which of the following routing protocols listed below use the Class D address of 224.0.0.9 to multicast its routing updates?

- A. EIGRP
- B. OSPF
- C. IGRP
- D. RIPv2

Answer: D

Explanation:

Class D addresses are multicast addresses used by many routing protocols for updating all neighbors on a link. Some of the Class D multicast addresses used by routing protocols are as follows:

1. OSPF - 224.0.0.5 and 224.0.0.6
2. Routing Information Protocol version 2 (RIPsv2) - 224.0.0.9
3. EIGRP - 224.0.0.10

Reference: Building Scalable Cisco Networks (Cisco Press) page 69

---

### **QUESTION 21:**

What are three characteristics of RIPv2? (Choose three)

- A. Supports variable-length subnet mask by default.
- B. Does not support variable-length subnet mask by default.
- C. Supports discontinuous networks by default.
- D. Multicast updates to 224.0.0.9.
- E. Broadcast updates.
- F. Does not support discontinuous networks by default

Answer: B, D, F

Explanation:

RIPV2 does not support VLSM or Discontinuous networks by default.

You have to run the no auto-summary command in order for RIPv2 to support either or both of these.

---

### **QUESTION 22:**

When comparing and contrasting the differences between classless and classful routing protocols, which of the following is true?

- A. In a classful system, a router uses a bit mask to determine the network and host portions of an address and there is no class restriction.
- B. In a classless system, a router determines the class of an address and then identifies the network and host octets based on that class.
- C. In a classful system, subnet mask information is maintained and passed along with each routing update.
- D. In a classless system, when a routing update is received about a different major network as configured on the receiving interface, the default subnet mask is applied.
- E. In a classful system, when a routing update is received about the same major network as configured on the receiving interface, the router applies the subnet mask configured on the receiving interface.

Answer: E

Explanation:

Routing protocols can be divided into various categories, such as classful or classless, and distance-vector or link-state. A classful routing protocol is one that exchanges routing information based on the classful IP boundaries. A router running a classful routing protocol such as RIPv1 or IGRP that receives an update for a network to which it is not physically connected will summarize that network at the default classful boundary. If the router received a route to network 10.1.2.0, for instance, it would summarize the route to 10.0.0.0 in its routing table, because an 8-bit (255.0.0.0) subnet mask is the default classful boundary for Class A addresses. A classless routing protocol, such as RIPv2, OSPF or EIGRP, exchanges the subnet mask with the subnet information. So if the aforementioned router was running OSPF, it would receive a route to network 10.1.2.0 with a 24-bit (255.255.255.0) subnet mask included, and would not summarize the route to 10.0.0.0.

---

**QUESTION 23:**

Which two statements are true with regard to RIPv1 and OSPF? (Choose two)

- A. RIPv1 uses the Dijkstra algorithm while OSPF uses the Bellman-Ford algorithm for calculating best path.
- B. RIPv1 uses the Bellman-Ford algorithm, OSPF uses the Dijkstra algorithm for calculating best path.
- C. RIP forwards the entire routing table incrementally, OSPF link-state advertisements are sent out when a change occurs and every thirty minutes if no change occurs.
- D. RIPv1 maintains a 15 hop count limit while OSPF maintains a 255 hop count limit.
- E. Both RIPv1 and OSPF carry subnet mask information and therefore support VLSM.

Answer: B, C

Explanation:

RIP is a distance vector routing protocol, which uses the Bellman-Ford algorithm for calculating the best path to a destination. OSPF is a link state routing protocol, which means that it uses the Dijkstra algorithm to determine the best path. RIP also periodically advertises the entire routing protocol to all neighbors, while OSPF uses LSA information to flood the area with routing information, but only after a topology change has occurred. To maintain stability, OSPF also floods out the entire OSPF table every thirty minutes if no change has occurred.

Incorrect Answers:

A: The reverse is true.


D: RIP maintains a hop limit of 15 (16 means it is unreachable) but OSPF does not have any hop limitations.

E: RIPv1 does not carry the subnet mask information in the routing updates, and so it does not support VLSM. RIPv2 does, however.

---

**QUESTION 24:**

The Certkiller IGRP network is displayed in the diagram below:



Certkiller 1 and Certkiller 2 are running IGRP. Certkiller 1 receives an update on its Fa0/0 interface from Certkiller 2 about the 199.10.50.64/26 network. Which mask will be applied to the update?

- A. 255.255.255.0
- B. 255.255.255.192
- C. 255.255.0.0
- D. 255.255.255.255

Answer: A

Explanation:

The ability to specify a different subnet mask for the same network number on different subnets is called Variable-Length Subnet Mask (VLSM). RIPv1 and IGRP are classful protocols and are incapable of carrying subnet mask information in their updates. Before RIPv1 or IGRP sends out an update, it performs a check against the subnet mask of the network that is about to be advertised and, in case of VLSM, the subnet gets dropped. In this example, first router Certkiller 2 checks to see whether 199.10.50.64/26 is part of the same major net as 10.10.0.0/16, which is the network assigned to the interface that will be sourcing the update. It is not, and so Certkiller 2 will summarize at the network boundary, which is 199.10.50.0/24 since this is a class C network.

---

**QUESTION 25:**

You're a systems administrator who's just chosen OSPF over RIP version 1. Your junior administrator Britney is confused and asks you why you didn't choose RIP. What would you tell her? (Choose all that apply)

- A. OSPF maintains smaller routing tables than RIP
- B. OSPF cost metric is based on number of hops.
- C. OSPF only sends routing updates only when necessary.
- D. OSPF supports VLSM and allows more efficient use of IP addresses.

Answer: C, D

Explanation:

C: RIP use periodic broadcast of the entire routing table, while OSPF use event-triggered announcements. OSPF uses

D: RIP Version 1 does not support VLSM, while OSPF does.

Incorrect Answers:

A: The size of the routing table depends on the number of routes. Since RIP version 1 automatically summarizes at the network boundary, the routing tables of RIP networks are generally smaller than OSPF networks.

B: RIP uses hop as cost metric. OSPF uses a metric based on the bandwidth of the links to the destination.

---

**QUESTION 26:**

OSPF is the popular choice as the standard open source routing protocol. Why is it more popular than both version 1 and version 2 of RIP?

- A. OSPF uses greater CPU overhead.
- B. OSPF has greater router memory requirements.
- C. OSPF allows for greater scalability.
- D. OSPF uses a simpler distance vector algorithm.
- E. OSPF allows for a simpler router configuration
- F. OSPF uses a simpler route selection process.

Answer: C

Explanation:

OSPF allows for greater flexibility than both RIP versions 1 and 2. OSPF allows for the use of areas, with each area connecting into the backbone area. This provides for a tiered, hierarchical design. With both RIP versions, the entire network is seen as one flat network. In addition, there is no limitation on the number of hops that a route can take with OSPF. With both RIP versions, there is a 16 hop limit. Lastly, as advantage that OSPF has over RIP version 1 is that it supports VLSM information while RIPv1 does not. However, this advantage does not apply to RIP version 2 since it also supports VLSM.

Incorrect Answers:

A, B: Although true, these are considered disadvantages to using OSPF.

D. OSPF is a link state routing protocol, while RIP is considered to be distance vector.

E. The OSPF router configuration is more complex to implement than a RIP router configuration.

F. OSPF uses the cost metrics of each link in determining routes, which is relatively simple. However, RIP uses a simple hop count as the metric, which is even simpler.

Reference: Building Scalable Cisco Networks (Cisco Press) pages 99-100.

---

**QUESTION 27:**

The Certkiller network is expanding, and there is a growing concern as to the total

number of router hops that are used. Which two of the routing protocols below use hop count to calculate their distance metric? (Select all that apply)

- A. IGRP
- B. OSPF
- C. EIGRP
- D. RIPv2
- E. RIPv1
- F. Static routes

Answer: D; E

Explanation:

Both RIP version 1 and version 2 use the hop count to measure the distance to a destination.

Incorrect Answers

A, C: IGRP and EIGRP use a composite metric that is calculated by factoring weighted mathematical values for internetwork delay, bandwidth, reliability, and load. By default, only bandwidth and delay are used in the calculation.

B: OSPF uses cost as the metric, which is found by taking the value 100,000,000 and dividing it by the bandwidth of each link.

---

**QUESTION 28:**

A network designer is considering the best routing protocol to use in a new network. What are the advantages of using OSPF over RIPv1 in a large network? (Select two)

- A. OSPF has a faster convergence time.
- B. OSPF requires less router memory.
- C. OSPF manages fewer internal tables.
- D. OSPF consumes less bandwidth through the use of incremental updates.

Answer: A, D

Explanation:

A: OSPF has a faster convergence time than RIP. OSPF supports triggered updates while RIP does not.

D: OSPF use incremental updates, while RIP broadcasts the entire routing table to all neighbors periodically.

Incorrect Answers

B: OSPF is more processor intensive and requires more memory. With faster processors this is not a big issue, but it can still be considered a disadvantage of OSPF.

C: OSPF use a more complex algorithm. These algorithms require more internal tables, such as routing, topology, and neighbor tables.

**QUESTION 29:**

The Bellman-Ford algorithm is used by the grand majority of distance vector routing protocols in calculating routes. One sophisticated routing protocol doesn't use this algorithm; instead it uses the DUAL (diffusing update algorithm). Which of the following routing protocols is it?

- A. IGRP
- B. OSPF
- C. EIGRP
- D. RIP v.2
- E. RIP v.1

Answer: C

Explanation:

DUAL is used by EIGRP. The Diffusing Update Algorithm (DUAL) is the algorithm used to obtain loop-freedom at every instant throughout a route computation. This allows all routers involved in a topology change to synchronize at the same time. Routers that are not affected by topology changes are not involved in the recomputation. The DUAL finite state machine embodies the decision process for all route computations. It tracks all routes advertised by all neighbors. The distance information, known as a metric, is used by DUAL to select efficient loop free paths. DUAL selects routes to be inserted into a routing table based on feasible successors. A successor is a neighboring router used for packet forwarding that has a least cost path to a destination that is guaranteed not to be part of a routing loop. When there are no feasible successors but there are neighbors advertising the destination, a recomputation must occur.

---


**QUESTION 30:**

**DRAG DROP**

Regarding distance vector and link state routing protocols, match each characteristic on the left to the appropriate routing protocol type on the right:

|   |                 |
|---|-----------------|
| Gains knowledge of the network based only on information learned from directly connected neighbors. | Distance Vector |
| Based on the Dijkstra algorithm.  | Place here      |
| Routing updates occur only when a change occurs.  | Place here      |
| Proprietary protocol type offered by Cisco.   | Place here      |
| Requires more memory and processing power   | Link State      |
| Based on the Bellman-Ford algorithm.  | Place here      |
| Sends complete routing tables at regular intervals  | Place here      |
|   | Place here      |

Answer:



---

**QUESTION 31:**

With regards to RIP version 1 networks, which of the following two statements are true? (Choose two)

- A. RIP v.1 route updates have a subnet mask field.
- B. RIP v.1 networks are known as classful networks.
- C. RIP v.1 networks are known as classless networks.
- D. RIP v.1 route updates do not have a subnet mask field.

Answer: B, D

Explanation:

There are two versions of RIP, namely RIPv1 and RIPv2. RIPv1 is only capable of classful routing. The routing updates do not carry subnet information, which means that a network's size is determined solely by the network class of its IP address, and there is no way to split a network into smaller subnets, each routed along a different path.

RIP cannot handle

classless routing. RIPv1 summarizes all routes it knows on classful network boundaries, so it is impossible to subnet a network properly via VLSM if you are running RIPv1, which is a Classful routing protocol.

Incorrect Answers:

A, C: RIP version 2 is a classless routing protocol that supports VLSM through the use of including subnet mask information within the routing updates. However, RIP version 1 does not.

---

**QUESTION 32:**

In an effort to increase network security, RIP routing authentication is being configured on a network. Which of the following types of authentication does RIPv2 support?

- A. MD5 only
- B. PAP and CHAP
- C. clear text and PAP
- D. clear text and MD5
- E. MD5 and CHAP

Answer: D

Explanation:

Authentication

Authentication can prevent communication with any RIP routers that are not intended to be part of the network, such as UNIX stations running routed. Only RIP updates with the authentication password are accepted. RFC 1723 defines simple plain-text authentication for RIPv2.

MD5 Authentication

In addition to plain-text passwords, the Cisco implementation provides the ability to use Message Digest 5 (MD5) authentication, which is defined in RFC 1321. Its algorithm takes as input a message of arbitrary length and produces as output a 128-bit fingerprint or message digest of the input, making it much more secure than plain-text passwords.

Incorrect Answers:

A: RIP version 2 supports plain text password authentication in addition to MD5.

B, C, E: PAP and CHAP are authentication methods used by the Point to Point Protocol (PPP), not RIP.

---

### **QUESTION 33:**

OSPF and IS-IS are examples of link state protocols. Which two features are characteristics of link-state routing protocols? (Choose two)

- A. Link-state protocols are based on the Dijkstra algorithm.
- B. Any incremental updates contain the full routing table.
- C. Link-state protocols flood the area with LSAs when a network change occurs.
- D. Routing updates are broadcast to neighbors.
- E. Link-state protocols use split-horizon updates to avoid routing loops.

Answer: A, C

Explanation:

Link-state routing protocols are based on the Dijkstra algorithm, meaning that routers exchange topology information with their nearest neighbors. The topology information is flooded throughout the AS, so that every router within the AS has a complete picture of the topology of the AS. This picture is then used to calculate end-to-end paths through the AS, normally using a variant of the Dijkstra algorithm. Therefore, in a link-state routing protocol, the next hop address to which data is forwarded is determined by choosing the best end-to-end path to the eventual destination.

Incorrect Answers:

B, D: Incremental updates and neighbor broadcasts are not utilized in link state routing protocols. Updates are triggered by a change in the network topology.  
E: Split horizons are used by distance vector routing protocols, not link state protocols.

---

**QUESTION 34:**

When an administrator uses VLSM to subnet a network address, which of the following rules must be adhered to? (Select three)

- A. A subnet is made up of all aggregated routes.
- B. A subnet can address hosts, or be further subnetted.
- C. A routing protocol must carry the subnet mask in updates.
- D. Summarized networks must have the same high-order bits.

Answer: B, C, D

Explanation:

All these statements are true regarding VLSM subnetting.

B: A subnet can address hosts, or be further subnetted.

C: The subnet mask must be included in routing updates, since the subnet mask information needs to be distributed throughout the entire network.

D: Summarized routes must have the same high order bits. All subnet masks start with a series of contiguous 1's, followed by 0's.

Incorrect Answers:

A: Route aggregation is not subnetting. Supernetting is the term being described in this choice.

---

**QUESTION 35:**

The Certkiller network is migrating to a classless routing protocol to support VLSM. Which two of the statements below correctly describe features of Variable Length Subnet Masking (VLSM)? (Select two)

- A. It supports IP version 4 and 6.
- B. It accommodates overlapping address ranges.
- C. It allows for better route summarization within routing tables.
- D. It allows subnetworks to be further divided into smaller subnets.

Answer: C, D

Explanation:

C: VLSM allows multiple levels of subnetted IP addresses within a single network, and thus allow for better route summarization within routing tables.

D: VLSMs provide the ability to include more than one subnet mask within a network, and the ability to subnet an already subnetted network address.

Incorrect Answers:

- A: VLSM was constructed for IP version 4.
  - B: Overlapping address ranges should not be used with VLSM.
- 

**QUESTION 36:**

The CTO of Certkiller INC. has called a meeting with all the system administrators to update them on the latest expansion plan of establishing multiple remote offices, and connecting those remote offices to the central office by a WAN. Management is insisting that an addressing scheme using VLSM be used, and they need six useable host addresses at each remote location. Assuming that Certkiller is using a class C IP range, which variable length subnet mask should be used to support the 6 hosts at each location?

- A. /24
- B. /28
- C. /29
- D. /30
- E. /31
- F. /32

Answer: C

Explanation:

To provide for six usable host addresses three host bits must be used. This would give exactly six host address ( $2^3-2=6$ ). This would leave 29 ( $32-3$ ) bits for the network mask. For any subnet, we must remember to subtract 2 from the usable host portion; 1 for the network IP address and one for the broadcast IP address.

Incorrect Answers:

- A: This would result in 254 hosts and only 1 network.
  - B: This would result in 14 hosts, with 14 networks.
  - D: This results in only 2 host IP addresses. This subnet mask is often used on point to point WAN links.
  - E, F: This would not allow for 6 hosts at each network.
- 

**QUESTION 37:**

Your boss has assigned you the difficult task of subnetting a Class B network, so you choose 172.29.100.0/28. Which of the following answer choices reflect a valid VLSM subnet that can be obtained from your Class B network?

- A. 172.29.100.8
- B. 172.29.100.32
- C. 172.29.100.60
- D. 172.29.100.104
- E. 172.29.100.126

Answer: B

Explanation:

We have the network 172.29.100.0/28 which is a subnetted Class B network. We want to identify a valid VLSM subnet mask of this network, which subnets this network further. We recalculate the subnetted network in binary to better study it. We must make sure that only the host bits, and not the network bits, are used in the fourth octet, so when the number is shown in binary the last 4 numbers should be 0.

First octet Second octet Third Octet Fourth Octet Comments

172 29 100 0

10101100 00011101 01100100 00000000 This is the network address

172 29 100 8

10101100 00011101 01100100 00000100 This is a host on the 172.29.100.0/28 subnet

172 29 100 32

10101100 00011101 01100100 00100000 This is a valid subnet address.

172 29 100 60

10101100 00011101 01100100 00111100 This is a host on the 172.29.100.48/28 subnet.

172 29 100 104

10101100 00011101 01100100 01101000 This is a host on the 172.29.100.96/28 subnet.

With a /28 network mask, the network boundaries will always be a multiple of 16 (172.29.100.0, 172.29.100.16, 172.29.100.32, 172.29.10.48, etc).

---

### QUESTION 38:

You are using the address space of 172.28.100.0/24. You want to divide use this address space for your WAN links by dividing the address space using a subnet mask of 172.28.100.0/30

How many subnets will this provide for WAN links?

- A. 6
- B. 14
- C. 30
- D. 62
- E. 126
- F. 254

Answer: D

Explanation:

When sub-dividing a /24 network with a subnet mask of /30 (255.255.255.252) we have 62 usable networks, with 2 usable hosts on each subnet.

First octet Second octet Third Octet Fourth Octet Comments

172 28 100 0

10101100 00011110 01100100 00000000 This is the entire address space

172 28 100 0

10101100 00011110 01100100 00000000 This is the 30-bit VLSM mask  
We can use 6 bits in the fourth octet for subnets. This gives us  $2^6-2=62$  subnets

---

**QUESTION 39:**

Certkiller is using the private IP address range of 192.168.100.0/24 for use on all of the WAN links. If this network is divided using a /30 mask, how many total WAN subnets can be supported?

- A. 14
- B. 30
- C. 62
- D. 126
- E. 254

Answer: C

Explanation:

When sub-dividing a class C network with a subnet mask of /30 (255.255.255.252) we have 62 usable networks, with 2 usable hosts on each subnet.

First octet Second octet Third Octet Fourth Octet Comments

192 168 100 0

10000000 10101000 01100100 00000000 This is the entire address space

192 168 100 0

11000000 10101000 01100100 00000000 This is the 30-bit VLSM mask

We can use 6 bits in the fourth octet for subnets. This gives us  $2^6-2=62$  subnets

Reference: Building Scalable Cisco Networks (Cisco Press) page 73

---

**QUESTION 40:**

You've secured the subnet and mask of 186.37.100.0/26 to use for some of your companies remote offices. Which VLSM mask below could provide you with six useable host addresses in each branch subnet?

- A. /24
- B. /28
- C. /29
- D. /30
- E. /31
- F. None of the above

Answer: C

Explanation:


3 bits (32-29) for the host would provide for  $2^3-2=6$  hosts as is required. In this case,

using the /29 mask on the existing /26 network would create 6 different networks with 6 hosts per network.

---

**QUESTION 41:**

Routers CK1 and CK2 are connected via a point to point serial line as shown in the diagram below:



Which command will configure a static route from the corporate network to 172.27.6.0?

- A. CK1 (config)# ip route 172.27.6.0 255.255.255.0 172.17.8.2
- B. CK2 (config)# ip route 172.27.6.0 255.255.0.0 172.17.8.2
- C. CK2 (config)# ip route 172.27.6.0 255.255.255.0 172.17.8.2
- D. CK1 (config)# ip route 172.27.6.0 255.255.0.0 172.17.8.1
- E. CK2 (config)# ip route 172.27.6.0 255.255.255.0 172.17.8.1

Answer: E

Explanation:

Only answer choice uses the correct syntax on the correct router. The static route must be placed on router CK2, and the value for the next hop should be the serial IP address of the connected router, which in this case is 172.17.8.1 on router CK1.

---

**QUESTION 42:**

The Certkiller network is using EIGRP as the network routing protocol. Which of the following statements correctly describe features and characteristics of routing using EIGRP? (Select three)

- A. It sends periodic updates every 60 seconds.
- B. EIGRP uses DUAL to achieve rapid convergence.
- C. Adjacencies exist between master routers (MRs) in each domain.
- D. It uses multicast to discover other EIGRP routers on an internetwork.
- E. EIGRP provides support for multiple network layer protocols: IPX, AppleTalk, and IP.

Answer: B, D, E

Explanation:

- B: Enhanced IGRP uses the Diffusing Update based algorithm (DUAL).
- D: EIGRP use hello multicast packets for neighbor discovery/recovery.

E: EIGRP supports IPX, Appletalk, and IP. Separate EIGRP instances can be created for each routed protocol.

Incorrect Answers:

A: EIGRP use triggered updates, not periodic updates.

C: Each EIGRP router maintains a neighbor table that lists adjacent routers. However, there is no concept of master routers in EIGRP.

Reference: "Introduction to Enhanced IGRP (EIGRP)"

<http://www.cisco.com/warp/public/459/7.html>

---

### **QUESTION 43:**

While comparing the tables associated with EIGRP and OSPF, you notice some similarities. Which EIGRP table is similar to OSPF's Link State Adjacency table?

- A. Neighbor table
- B. Routing table
- C. Topology table
- D. Successor table
- E. None of the above

Answer: A

Explanation:

Each EIGRP router maintains a neighbor table that lists adjacent routers. This table is comparable to the neighbors (adjacency) database used by OSPF.

Incorrect Answers:

B: The routing table of EIGRP would compare to the routing table of OSPF

C: Both EIGRP and OSPF contain topology tables, which would compare to each other.

D: EIGRP does not have a successor table. Successors are maintained in the routing and topology tables of EIGRP.

Reference: Building Scalable Cisco Networks (Cisco Press) page 250

---

### **QUESTION 44:**

DRAG DROP

Match the correct EIGRP term on the left to the slot in the middle that is adjacent to the correct description on the right.

| EIGRP term, use these |            | Description                        |
|-----------------------|------------|------------------------------------|
| Successor             | place here | lists adjacent routers             |
| Routing table         | place here | route entries for all destinations |
| Neighbor table        | place here | primary route to destination       |
| Feasible successor    | place here | best routes to destinations        |
| Topology table        | place here | backup route to destination        |

Answer:

Explanation:

Neighbor table - lists adjacent routers

Topology Table - route entries for all destinations.

Successor - primary route used to reach a destination

Routing table - best routes to a destination

Feasible successor - backup route to the destination.

Explanation:

\* Neighbor Table - Each EIGRP router maintains a neighbor table that lists adjacent routers. This table is comparable to the neighbors (adjacency) database used by OSPF.

\* Topology Table - An EIGRP router maintains a topology table for each network protocol configured: IP, IPX, and AppleTalk. All learned routes to a destination are maintained in the topology table.

\* Routing Table - EIGRP choose the best routes to a destination from the topology table and places these routes in the routing table. The router maintains one routing table for each network protocol.

\* Successor - This is the primary route used to reach a destination. Successors are kept in the routing table.

\* Feasible Successor - This is a neighbor that is downstream with respect to the destination, but it is not the least-cost path and thus is not used for forwarding data. In other words, this is a backup route to the destination. These routes are selected at the same time as successors, but are kept in the topology table.

Reference: Building Scalable Cisco Networks (Cisco Press) page 250

---

**QUESTION 45:**

Different types of EIGRP transmissions are sent between neighboring routers.

Which three of the following terms are known as 'reliable packets' in EIGRP?

(Select three)

- A. hello
- B. ACK
- C. reply
- D. query
- E. update

Answer: C, D, E

Explanation:

Updates are used to convey the reachability of destinations. When a new neighbor is discovered, update packets are sent so the neighbor can build up its topology table. In this case, update packets are unicast. In other cases, such as a link cost change, updates are multicast. Updates are always transmitted reliably.

Queries and replies are sent when destinations go into Active state. Queries are always multicast unless they are sent in response to a received query. In this case, it is unicast back to the successor that originated the query. Replies are always sent in response to queries to indicate to the originator that it does not need to go into Active state because it has feasible successors. Replies are unicast to the originator of the query. Both queries and replies are transmitted reliably.

EIGRP reliable packets are: Update, Query and Reply.

EIGRP unreliable packets are: Hello and Ack.

Incorrect Answers:

A, B: Hellos are multicast for neighbor discovery/recovery. They do not require acknowledgment. A hello with no data is also used as an acknowledgment (ack). Acks are always sent using a unicast address and contain a non-zero acknowledgment number.

Reference: Cisco BSCN version 1.0 study guide, pages 6-18.

---

#### **QUESTION 46:**

Which one of the following statements best describes the way EIGRP advertises subnet mask information to its destination networks?

- A. EIGRP advertises a prefix length for each destination network.
- B. EIGRP advertises a fixed length subnet mask for each destination network.
- C. EIGRP advertises only a classful subnet mask for each destination network.
- D. EIGRP, like IGRP and RIP, does not advertise a subnet mask for each destination network.

Answer: A

Explanation:

EIGRP advertises a prefix length to each destination network. It supports VLSM which allows subnet masks to be automatically summarized at the network boundary, and EIGRP can also be configured to summarize on any bit boundary at any interface.

Incorrect Answers:

B: This is incorrect because EIGRP supports Variable Length subnet masks.

C: This is incorrect because EIGRP is considered to be a classless protocol, not classful.

D: EIGRP is more sophisticated than IGRP and RIP and actually does advertise a subnet mask to each destination network.

**QUESTION 47:**

EIGRP was designed specifically not to use excessive bandwidth for routing updates and functions across WAN links. If EIGRP is configured on an interface, what is the maximum percent of its bandwidth EIGRP would use, by default?

- A. 10%
- B. 25%
- C. 50%
- D. 75%
- E. 100%

Answer: C

Explanation:

By default, EIGRP will limit itself to using no more than 50% of the available bandwidth. The maximum bandwidth that can be used is a configurable parameter. Reference: Cisco, Configuration Notes for the Enhanced Implementation of EIGRP. <http://www.cisco.com/warp/public/103/12.html>

---

**QUESTION 48:**

EIGRP supports Variable Length Subnet Masking (VLSM). Which two features of EIGRP are true regarding EIGRP support for VLSM? (Select two)

- A. It advertises a routing mask
- B. It is a classful routing protocol
- C. It is a classless routing protocol
- D. It does not advertise a routing mask.

Answer: A, C

Explanation:

EIGRP includes the routing mask when it advertises routes. Subnet mask information must be included in all routing updates for classless protocols. EIGRP is a classless protocol, meaning that it does not solely rely on the network class in order to find the subnet mask. These two features enable EIGRP to support VLSM.

Incorrect Answers

B: Classful routing protocols do not support VLSM.

D: The subnet mask must be included in the route to support VLSM.

---

**QUESTION 49:**

In an effort to decrease the size of the routing tables in the Certkiller network, summarization is being configured on all routers. Where in an EIGRP network is it most appropriate to implement route summarization?

- A. At area border routers
- B. At autonomous system boundary routers
- C. It is done automatically by the DR router in each area
- D. Manually at any interface of any EIGRP router within the network
- E. In backbone area routers.

Answer: B

Explanation:

Enhanced IGRP performs route summarization at classful network boundaries by default. Automatic route summarization occurs at major network boundaries. ASBRs (autonomous system boundary routers) are used at the major network boundaries.

Incorrect Answers

A, C, E: Area Border Router and DR are used in OSPF, but not in EIGRP. Backbone area routers are also concepts used in OSPF, but not in EIGRP.

D: Manual summarization can be done in any interface at any router within network, but this is not the preferred solution. Summarizing manually at each router can result in black hole routing. However, done properly this is an acceptable practice. If the question calls for more than one answer then the best choice would be B and D.

---

### **QUESTION 50:**

The Certkiller network is using route summarization to decrease the size of the routing tables. In an EIGRP network, which of the following describes the best method for implementing summarization?

- A. At WAN interfaces.
- B. Manually at major network boundaries.
- C. Dynamically at discontinuous interfaces.
- D. Dynamically at major network boundaries.

Answer: D

Explanation:

EIGRP supports arbitrary route summarization. Route summarization takes place automatically at major network boundaries (where network are supernetted).

Note: Route summarization is the consolidation of advertised addresses.

Incorrect Answers

A: Manually configuring summarization at WAN boundaries is not necessarily recommended in all situations. If the routers belong to the same major network boundary then summarization should not be configured.

B: Route summarization is automatic in EIGRP, and it is recommended to keep the defaults.

C: Route summarization should only be applied at contiguous interfaces. Choosing this option would most likely lead to black hole routing, making many networks unreachable.

---

**QUESTION 51:**

The Certkiller network is running EIGRP on the T1 links within their NBMA frame relay network. By default what would the hold time be for hello packets across these WAN links?

- A. 30 seconds
- B. 60 seconds
- C. 90 seconds
- D. 180 seconds
- E. 5 Seconds

Answer: D

Explanation:

EIGRP sends hello packets every 5 seconds on high bandwidth links and every 60 seconds on low bandwidth multipoint links. The hold time is typically three times the hello interval. In this scenario, on slow NBMA media, hold time will be 180 seconds.

Incorrect Answers:

A, B, C: Although the hello timers can be manually configured for these values, they are not the default time.

E: This is the default for high bandwidth links, such as Ethernet networks.

Reference: Enhanced Interior Gateway Routing Protocol

<http://www.cisco.com/warp/public/103/eigrp1.html>

---

**QUESTION 52:**

EIGRP is being used as the routing protocol within the Certkiller network. Which two of the following features allow EIGRP to support classless routing? (Select two)

- A. Djikstra's algorithm
- B. discontiguous subnets
- C. variable length subnet masks
- D. periodic update announcements
- E. unequal path-cost load balancing

Answer: B, C

Explanation:

B: Discontiguous subnets are supported by EIGRP. This is a classless routing protocol feature.

C: EIGRP support variable length subnet masks (VLSM). This is a classless routing protocol feature.

Incorrect Answers

A: EIGRP use the DUAL algorithm, not the Djikstra's algorithm. The Djikstra algorithm

is used by distance vector protocols. EIGRP is considered to be a hybrid routing protocol.

D: EIGRP use periodic hello messages, not update announcements. EIGRP use update packets to convey reachability of destinations. When a new neighbor is discovered, Update packets are sent so the neighbor can build up its. These are not periodic, however.

E: IGRP and EIGRP support unequal cost path load balancing, which is known as variance. However, this is not a classless feature of EIGRP.

Note: Classless routing protocols include the routing mask with the route advertisement.

This enables discontinuous subnets and variable length subnet masks.

Reference: Introduction to Enhanced IGRP (EIGRP)

<http://www.cisco.com/warp/public/103/1.html>

How Does Unequal Cost Path Load Balancing (Variance) Work in IGRP and EIGRP?

<http://www.cisco.com/warp/public/103/19.html>

---

### QUESTION 53:

EIGRP is being used as the routing protocol on the Certkiller network. While troubleshooting some network connectivity issues, you notice a large number of EIGRP SIA (Stuck In Active) messages. What causes these SIA routes? (Select two)

- A. The neighboring router stops receiving ACK packets from this router.
- B. The neighboring router starts receiving route updates from this router.
- C. The neighboring router is too busy to answer the query (generally caused by high CPU utilization).
- D. The neighboring router is having memory problems and cannot allocate the memory to process the query or build the reply packet.

Answer: C, D

Explanation:

SIA routes are due to the fact that reply packets are not received. This could be caused by a router which is unable to send reply packets. The router could have reached the limit of its capacity, or it could be malfunctioning.

Incorrect Answers

A: Missing replies, not missing ACKs, cause SIA.

B: Routes updates do not cause SIA.

Notes: If a router does not receive a reply to all outstanding queries within 3 minutes, the route goes to the stuck in active (SIA) state. The router then resets the neighbors that fail to reply by going active on all routes known through that neighbor, and it re-advertises all routes to that neighbor.

Reference: Enhanced Interior Gateway Routing Protocol

<http://www.cisco.com/warp/public/103/eigrp3.html>

---

### QUESTION 54:

EIGRP is being configured on the Certkiller network. After the initial configuration, the EIGRP routers form relationships with its peer routers. Which three of the

following statements correctly describe these peer relationships? (Select three)

- A. EIGRP will form neighbors if the routers are not adjacent
- B. EIGRP will not form neighbors if the metric K-values do not match.
- C. EIGRP will not form neighbors if the router AS numbers do not match.
- D. EIGRP will form neighbors over primary and secondary interface addressing.
- E. EIGRP will form neighbors even though hello and hold timers do not match on the peering interfaces.

Answer: B, C, E

Explanation:

B: To become neighbors the routers metric must be comparable.

C: EIGRP neighbors must have the same AS number.

E: It is possible for two routers to become EIGRP neighbors even though the hello and hold timers don't match.

Incorrect Answers:

A: Neighbor discovery/recovery is the process that routers use to dynamically learn of other routers on their directly attached networks.

D: EIGRP doesn't build peer relationships over secondary addresses.

Reference: Cisco White Paper, Enhanced Interior Gateway Routing Protocol

<http://www.cisco.com/warp/public/103/eigrp1.html>

---

### **QUESTION 55:**

The Certkiller network is an EIGRP network consisting of four links (Link A, Link B, Link C, & Link D). The costs associated with each link is A=1, B=3, C=2, D=4. According to the variables below, which link is going to be the feasible successor?

- A. Link A
- B. Link B
- C. Link C
- D. Link D

Answer: C

Explanation:

Link A is the best link and it therefore the successor. Link C is the next best link, and it is therefore the feasible successor.

Incorrect Answers

A: Link A is the successor, not the feasible successor.

B: Both Link A and Link C are better than Link B.

D: The feasible successor must have Advertised Destination that is less than the Feasible Distance.

Note: Feasible distance is the best metric along a path to a destination network, including the metric to the neighbor advertising that path. Reported distance, or advertised distance,

is the total metric along a path to a destination network as advertised by an upstream neighbor. A feasible successor is a path whose reported distance is less than the feasible distance.

Reference: Enhanced Interior Gateway Routing Protocol  
<http://www.cisco.com/warp/public/103/eigrp1.html>

---

**QUESTION 56:**

Part of the routing table of router CK1 is displayed below:

```
S 62.99.153.0/24 [1/0] via 209.177.64.130
172.209.12.0/32 is subnetted, 1 subnets
D EX 172.209.1
[170/2590720] via 209.179.2.114, 06:47:28, Serial0/0/0.1239
62.113.17.0/24 is variably subnetted, 2 subnets, 2 masks
D EX 99.3.215.0/24
[170/27316] via 209.180.96.45, 09:52:10, FastEthernet11/0/0
[170/27316] via 209.180.96.44, 09:52:10, FastEthernet11/0/0
25.248.17.0/24
[90/1512111] via 209.179.66.25, 10:33:13, Serial0/0/0.1400001
[90/1512111] via 209.179.66.41, 10:33:13, Serial0/0/0.1402001
62.113.1.0/24 is variably subnetted, 12 subnets, 2 masks
D 62.113.1.227/32
[90/2611727] via 209.180.96.45, 10:33:13, FastEthernet1/0/0
[90/2611727] via 209.180.96.44, 10:33:13, FastEthernet1/0/0
S* 0.0.0.0/0 [1/0] via 209.180.96.14
```

From analyzing the above command output, what is the administrative distance of the external EIGRP routes?

- A. 24
- B. 32
- C. 90
- D. 170
- E. 27316
- F. None of the above

Answer: D

Explanation:

By default an external EIGRP route has a value of 170. By examining the exhibit we see that this default value of the external EIGRP routes (see D-EX in exhibit) indeed is set to 170. The first value within the brackets display the AD, so with a value of [170/27316] the AD is 170 and the metric of the route is 27316.

Incorrect Answers:

A, B: This is the subnet mask used for some of the routes in the table.

C: This is the AD of the internal EIGRP routes, which is the default

E: This is the EIGRP metric of the external EIGRP routes.


Reference: What Is Administrative Distance?

[http://www.cisco.com/warp/public/105/admin\\_distance.html](http://www.cisco.com/warp/public/105/admin_distance.html)

---

**QUESTION 57:**

The Certkiller network is shown below, along with the relevant router configurations:



```
CK1 # show run
interface Loopback0
ip address 10.10.10.1 255.255.255.0
!
interface Ethernet0
ip address 172.29.1.1 255.255.255.0
media-type 10BaseT
!!
router eigrp 100
redistribute connected
network 172.29.0.0
auto-summary
no eigrp log-neighbor-changes
!
ip classless
no ip http server
CK2 # show run
interface Ethernet0
ip address 172.29.1.2 255.255.255.0
media-type 10BaseT
!
interface Ethernet1
ip address 172.19.2.2 255.255.255.0
media-type 10BaseT
!
router eigrp 100
network 172.19.0.0
network 172.29.0.0
!
ip classless
no ip http server
CK3 # show run
interface Ethernet1/0
ip address 172.19.2.3 255.255.255.0
```

```
!  
router eigrp 100  
network 172.19.0.0  
auto-summary  
no eigrp log-neighbor-changes  
!  
ip classless  
ip http server
```

With the topology found in the graphic, what will the CK1 loopback 0 be in the CK3 routing table?

- A. It will show up in the routing table as D 10.0.0/8.
- B. It will show up in the routing table as D EX 10.0.0.0/8.
- C. It will show up in the routing table as D 10.0.0./24.
- D. It will not show up in CK3 routing table because there is no network command on CK1 .

Answer: B

Explanation:

Because router CK1 is configured with route redistribution, it will redistribute the connected loopback network into EIGRP. Because redistributed routes will show up as external EIGRP routes in the routing table, choice B is correct. Although the loopback interface is using a /24 subnet mask, EIGRP summarizes at network boundaries by default so the network will appear as the class A network of 10.0.0.0/8 in the routing table of the other routers.

Incorrect Answers:

- A: The route will be external, since it was redistributed into EIGRP.
- C: It will be external because of redistribution, and it will also be summarized since that is the default behavior of EIGRP.
- D: Although it was not configured under the EIGRP network command, it would be redistributed because it is a connected route.

---

### **QUESTION 58:**

EIGRP has been configured as the routing protocol on the Certkiller network.  
Which statements are true regarding EIGRP? (Choose three)

- A. By default, EIGRP performs auto-summarization across classful network boundaries.
- B. EIGRP uses an area hierarchy to increase network scalability
- C. To speed convergence, EIGRP attempts to maintain a successor and feasible successor path for each destination.
- D. EIGRP uses hellos to establish neighbor relationships.
- E. By default, EIGRP uses the Dijkstra algorithm to determine the best path to destination network based on bandwidth and delay.

Answer: A, C, D

Explanation:

The default behavior of EIGRP routers is to automatically summarize routes at the network boundary.

EIGRP routers maintain information regarding the successors and the feasible successors to each network destination. This information is useful in the convergence time whenever any of the links fail.

EIGRP neighbors periodically use hellos to establish the relationship. Should any of the neighbors fail, triggered updates are sent to update the network regarding the topology change.

Incorrect Answers:

B: EIGRP networks are inherently flat, as each router is perceived as a peer to every other EIGRP router. IS-IS and OSPF utilize hierarchical network topologies.

E: The Dijkstra algorithm is used by link state routing protocols. EIGRP is considered to be a hybrid routing protocol.

---

**QUESTION 59:**

Router Certkiller 1 is the headquarters router in a hub and spoke topology supporting 24 remote offices. Point-to-multipoint Frame Relay EIGRP network is deployed between the headquarters and the remote offices. There is no bandwidth command configured under either the major serial interface or the subinterface on router Certkiller 1.

What is the bandwidth of each Frame Relay connection perceived by the EIGRP process?

- A. 64 kbps
- B. 128 kbps
- C. 1.544 Mbps
- D. 1.536 Mbps

Answer: A

Explanation:

By default in point-to-multipoint Frame Relay EIGRP network all subinterfaces share the bandwidth and each one receives  $(\text{throughput of the channel})/(\text{number of subinterfaces})$  throughput. In this case, the formula is  $1544\text{kbps}/24 = 64\text{kbps}$ .

---

**QUESTION 60:**

What administrative distance is given to EIGRP summary routes?

- A. 0
- B. 1
- C. 5

- D. 90
- E. 95
- F. 170

Answer: C

Explanation:

|   |                 |
|---|-----------------|
| Connected interface or static route that identifies the outgoing interface rather than the next hop | 0               |
| Static route  | 1               |
| EIGRP summary route   | 5               |
| External BGP  | 20              |
| EIGRP   | 90              |
| IGRP  | 100             |
| OSPF  | 110             |
| RIP   | 120             |
| External EIGRP  | 170             |
| Internal BGP  | 200             |
| An unknown network  | 255 or infinity |

---

**QUESTION 61:**

From the following choices, identify the characteristics of EIGRP feasible successors. (Select three)

- A. A feasible is selected by comparing the advertised distance of a non-successor router to the feasible distance of the best route.
- B. If the advertised distance of the non-successor route is less than the feasible of best route, then that route is identified as a feasible successor.
- C. If the successor becomes unavailable, the feasible successor can be used immediately without recalculating for a lost route.
- D. The feasible successor can be found in the routing table.
- E. Traffic will be load balanced between feasible successors with the same advertised distance.

Answer: A, B, C

Explanation:

Once a neighbor relationship has been formed, called an Adjacency, the routers exchange routing update information and each router builds its own topology table. The Updates contain all the routes known by the sender. For each route, the receiving router calculates a distance for that route based on the distance that is conveyed and the cost to that

neighbor that advertised the particular route. If the receiving router sees several routes to a particular network with different metrics, then the route with the lowest metric becomes the Feasible Distance (FD) to that network. The Feasible Distance is the metric of a network advertised by the connected neighbor plus the cost of reaching that neighbor. This path with the best metric is entered into the routing table because this is the quickest way to get to that network.

With the other possible routes to a particular network with larger metrics, the receiving router also receives the Reported Distance (RD) to this network via other routers. The Reported Distance being the total metric along a path to a destination network as advertised by an upstream neighbor. The Reported Distance for a particular route is compared with the Feasible Distance that it already has for that route. If the Reported Distance is larger than the Feasible Distance then this route is not entered into the Topology Table as a Feasible Successor. This prevents loops from occurring. If the Reported Distance is smaller than the Feasible Distance, then this path is considered to be a Feasible Successor and is entered into the Topology table. The Successor for a particular route is the neighbor/peer with the lowest metric/distance to that network. If the receiving router has a Feasible Distance to a particular network and it receives an update from a neighbor with a lower advertised distance (Reported Distance) to that network, then there is a Feasibility Condition. In this instance, the neighbor becomes a Feasible Successor for that route because it is one hop closer to the destination network. There may be a number of Feasible Successors in a meshed network environment, up to 6 of them are entered into the Topology table thereby giving a number of next hop choices for the local router should the neighbor with the lowest metric fail. What you should note here, is that the metric for a neighbor to reach a particular network (i.e. the Reported Distance) must always be less than the metric (Feasible Distance) for the local router to reach that same network. This way routing loops are avoided. This is why routes that have Reported Distances larger than the Feasible Distance are not entered into the Topology table, so that they can never be considered as successors, since the route is likely to loop back through that local router.

Incorrect Answers:

D: The feasible successors are found in the topology table, but not the active routing table.

E: With EIGRP, traffic is load balanced across equal cost links in the routing table, but not between feasible successors.

Reference: <http://www.rhyshaden.com/eigrp.htm>

---

## **QUESTION 62:**

Which of the following statements are true about EIGRP operation? (Choose three)

- A. When summarization is configured, the router will add a route to null 0
- B. The summary route remains in the route table, even if there are no more specific routes to the network.
- C. Summarization is configured on a per-interface level.
- D. The maximum metric for the specific routes is used as the metric for the summary route.

E. Automatic summarization across major network boundaries is enabled by default.

Answer: A, C, E

Explanation:

Using the `ip summary-address eigrp` summarization command will cause the creation of an EIGRP summary default route to the `null0` interface with an administrative distance of 5. Caution should be taken when using this as a default route. The low administrative distance of this default route can cause this route to displace default routes learned from other neighbors from the routing table. If the default route learned from the neighbors is displaced by the summary default route, or if the summary route is the only default route present, all traffic destined for the default route will not leave the router, instead, this traffic will be sent to the `null0` interface where it is dropped.

By default, Cisco routers will automatically summarize EIGRP routes across major network boundaries. With EIGRP, You can configure a summary aggregate address for a specified interface. If there are any more specific routes in the routing table, EIGRP will advertise the summary address out the interface with a metric equal to the minimum of all more specific routes.

Incorrect Answers:

B: When all of the specific routes used within a summary route are deleted from the routing table, the summary route will also be deleted.

D: The metric used in a summary route is the best metric from among the summarized routes, not the maximum (worst).

---

### QUESTION 63:

What action does an EIGRP router take when it cannot find a feasible successor for a network?

- A. It examines the routing and neighbor tables for the next best path.
- B. It transitions from passive to active state for that network and queries its neighbors.
- C. It examines the topology for a next best path.
- D. It transitions from active to passive state for the next network and queries its neighbors.

Answer: B

Explanation:

With EIGRP, A topology table entry for a destination can have one of two states. A route is considered in the Passive state when a router is not performing a route recomputation. The route is in Active state when a router is undergoing a route recomputation. If there are always feasible successors, a route never has to go into Active state and avoids a route recomputation.


When there are no feasible successors, a route goes into Active state and a route recomputation occurs. A route recomputation commences with a router sending a query packet to all neighbors. Neighboring routers can either reply if they have feasible

successors for the destination or optionally return a query indicating that they are performing a route recomputation. While in Active state, a router cannot change the next-hop neighbor it is using to forward packets. Once all replies are received for a given query, the destination can transition to Passive state and a new successor can be selected.

---

**QUESTION 64:**

The Certkiller EIGRP network is displayed in the following topology diagram:



You work as a network technician at Certkiller .com. Study the exhibits carefully. If the command "variance 3" was added to the EIGRP configuration of Certkiller 5, which path or paths would be chosen to route traffic from Certkiller 5 to network X?

- A. Certkiller 5- Certkiller 2- Certkiller 1
- B. Certkiller 5- Certkiller 2- Certkiller 1 and Certkiller 5- Certkiller 3- Certkiller 1.
- C. Certkiller 5- Certkiller 3- Certkiller 1 and Certkiller 5- Certkiller 4- Certkiller 1.
- D. Certkiller 5- Certkiller 2- Certkiller 1, Certkiller 5- Certkiller 3- Certkiller 1, and Certkiller 5- Certkiller 4- Certkiller 1.

Answer: B

Explanation:

In this question the variance 3 command is used. Since the lowest metric is 10.  $10 * 3 = 30$  meaning we can use all routes with a metric of 30 and under which is answer B(5-2-1 and 5-3-1), not answer C.

---

**QUESTION 65:**

What is the purpose of the "eigrp stub" configuration command?

- A. To increase scalability by limiting the EIGRP query range.
- B. To reduce the size of the routing table by blocking the D EX (External EIGRP) routes in the EIGRP stub router.
- C. To reduce the convergence time by enabling the EIGRP stub router to propagate the EIGRP queries from the EIGRP hub router.

D. To reduce the convergence time by enabling the EIGRP stub router to also perform query requests to the EIGRP hub router.

Answer: A

Explanation:

Complex, redundant EIGRP networks can cause scalability problems. The best solution to this is to provide a means within the context of the EIGRP protocol itself to control traffic flows and limit query depth. The EIGRP Stub Router functionality in Cisco IOS Software Release 12.0(7)T can achieve this solution.

The Enhanced Interior Gateway Routing Protocol (EIGRP) Stub Routing feature improves network stability, reduces resource utilization, and simplifies stub router configuration.

Stub routing is commonly used in a hub and spoke network topology. In a hub and spoke network, one or more end (stub) networks are connected to a remote router (the spoke) that is connected to one or more distribution routers (the hub). The remote router is adjacent only to one or more distribution routers. The only route for IP traffic to follow into the remote router is through a distribution router. This type of configuration is commonly used in WAN topologies where the distribution router is directly connected to a WAN. The distribution router can be connected to many more remote routers. Often, the distribution router will be connected to 100 or more remote routers. In a hub and spoke topology, the remote router must forward all nonlocal traffic to a distribution router, so it becomes unnecessary for the remote router to hold a complete routing table. Generally, the distribution router need not send anything more than a default route to the remote router.

When using the EIGRP Stub Routing feature, you need to configure the distribution and remote routers to use EIGRP, and to configure only the remote router as a stub. Only specified routes are propagated from the remote (stub) router. The router responds to queries for summaries, connected routes, redistributed static routes, external routes, and internal routes with the message "inaccessible." A router that is configured as a stub will send a special peer information packet to all neighboring routers to report its status as a stub router.

Any neighbor that receives a packet informing it of the stub status will not query the stub router for any routes, and a router that has a stub peer will not query that peer. The stub router will depend on the distribution router to send the proper updates to all peers.

Reference:

[http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products\\_feature\\_guide09186a0080087026.html](http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a0080087026.html)

---

**QUESTION 66:**

OSPF has recently been configured on router CK1 . A router running OSPF can be in one of many different states. Which state must router CK1 be in with a neighbor before it will begin to route traffic?

- A. full state
- B. active state

- C. two-way state
- D. forwarding state

Answer: A

Explanation:

Full is the normal state for an OSPF router. In this state, routers are fully adjacent with each other and the router is routing traffic. The OSPF neighbor states are described below:

**Down**

This is the first OSPF neighbor state. It means that no information (hellos) has been received from this neighbor, but hello packets can still be sent to the neighbor in this state.

During the fully adjacent neighbor state, if a router doesn't receive hello packet from a neighbor within the RouterDeadInterval time (RouterDeadInterval = 4\*HelloInterval by default) or if the manually configured neighbor is being removed from the configuration, then the neighbor state changes from Full to Down.

**Attempt**

This state is only valid for manually configured neighbors in an NBMA environment. In Attempt state, the router sends unicast hello packets every poll interval to the neighbor, from which hellos have not been received within the dead interval.

**Init**

This state specifies that the router has received a hello packet from its neighbor, but the receiving router's ID was not included in the hello packet. When a router receives a hello packet from a neighbor, it should list the sender's router ID in its hello packet as an acknowledgment that it received a valid hello packet.

**2-Way**

This state designates that bi-directional communication has been established between two routers. Bi-directional means that each router has seen the other's hello packet. This state is attained when the router receiving the hello packet sees its own Router ID within the received hello packet's neighbor field. At this state, a router decides whether to become adjacent with this neighbor. On broadcast media and non-broadcast multiaccess networks, a router becomes full only with the designated router (DR) and the backup designated router (BDR); it stays in the 2-way state with all other neighbors. On Point-to-point and Point-to-multipoint networks, a router becomes full with all connected routers.

**Exstart**

Once the DR and BDR are elected, the actual process of exchanging link state information can start between the routers and their DR and BDR.

In this state, the routers and their DR and BDR establish a master-slave relationship and choose the initial sequence number for adjacency formation. The router with the higher router ID becomes the master and starts the exchange, and as such, is the only router that can increment the sequence number. Note that one would logically conclude that the DR/BDR with the highest router ID will become the master during this process of master-slave relation. Remember that the DR/BDR election might be purely by virtue of a higher priority configured on the router instead of highest router ID. Thus, it is possible

that a DR plays the role of slave. And also note that master/slave election is on a per-neighbor basis.

#### Exchange

In the exchange state, OSPF routers exchange database descriptor (DBD) packets. Database descriptors contain link-state advertisement (LSA) headers only and describe the contents of the entire link-state database. Each DBD packet has a sequence number which can be incremented only by master which is explicitly acknowledged by slave. Routers also send link-state request packets and link-state update packets (which contain the entire LSA) in this state. The contents of the DBD received are compared to the information contained in the routers link-state database to check if new or more current link-state information is available with the neighbor.

#### Loading

In this state, the actual exchange of link state information occurs. Based on the information provided by the DBDs, routers send link-state request packets. The neighbor then provides the requested link-state information in link-state update packets. During the adjacency, if a router receives an outdated or missing LSA, it requests that LSA by sending a link-state request packet. All link-state update packets are acknowledged.

#### Full

In this state, routers are fully adjacent with each other. All the router and network LSAs are exchanged and the routers' databases are fully synchronized.

Full is the normal state for an OSPF router. If a router is stuck in another state, it's an indication that there are problems in forming adjacencies. The only exception to this is the 2-way state, which is normal in a broadcast network. Routers achieve the full state with their DR and BDR only. Neighbors always see each other as 2-way.

#### Incorrect Answers

B: An IGRP or EIGRP router can be in an active state, but an OSPF cannot be in an active state.

C: At the two-way state, a router decides whether to become adjacent with this neighbor. However, routes are not exchanged between the routers.

D: An interface (link) can be in a forwarding state, but not a router.

Note: An OSPF router can be in one of the following states: down, attempt, init, 2-way, exstart, exchange, loading, or full.

Reference: OSPF Neighbor States

<http://www.cisco.com/warp/public/104/13.html>

RFC 2328, OSPF Version 2

---

### **QUESTION 67:**

One of the WAN links on an OSPF has just gone down, causing a network topology change. What does an OSPF network do to maintain route consistency when it notices a change in the network?

- A. It enters the exstart state with its neighbors.
- B. It floods the area with new routing information.
- C. It generates a routing exchange using the hello protocol.
- D. It waits for the holddown timers to expire, then sends the update.

Answer: B

Explanation:

Link state routing protocols generate routing updates only when there is a change in the topology. When a link changes state, a link-state advertisement (LSA) concerning that link (route) is created by the device that detected the change and propagated to all neighboring devices using a special multicast address. This process is called flooding.

Incorrect Answers:

A: The exstart state is not entered when a change in the topology occurs.

Note: Once the DR and BDR have been elected, the routers are considered to be in the exstart state and are ready to discover the link-state information about the internetwork and create their link-state databases.

C: The hello protocol is not used to exchange routing information. The purpose of hellos is for identifying neighbors and verifying periodically that the neighbors exist.

D: OSPF does not use holddown timers like RIP. Topology changes are flooded immediately.

---

**QUESTION 68:**

Certain types of OSPF networks require the election of a Designated Router (DR) as well as a Backup Designated Router (BDR). In which of the following network types does this election take place? (Select two)

- A. Point-to-point
- B. Nonbroadcast multi-access
- C. Point-to-multipoint
- D. Broadcast multi-access

Answer: B, D

Explanation:

Mode Adjacency

NBMA Manual Configuration DR/BDR elected


Broadcast Automatic DR/BDR elected

Reference: Building Scalable Cisco Networks (Cisco Press) page 124

---

**QUESTION 69:**

Three Certkiller routers and connected via a frame relay network as shown below:



Examine the s0 interface on router Certkiller 1. What is the OSPF network type by default?

- A. Nonbroadcast
- B. Broadcast
- C. Point-to-multipoint
- D. Point-to-point
- E. Point-to-multipoint nonbroadcast

Answer: A

Explanation:

Based on the information above, the physical serial 0 interface is used, with no subinterfaces.

When OSPF is configured on a physical interface (for example, interface S0) of a nonbroadcast multiaccess technology such as Frame Relay, the default network type of NON\_BROADCAST is assigned. When OSPF is configured on point-to-point subinterfaces, the default interface type of POINT\_TO\_POINT is assigned. When OSPF is configured on multipoint subinterfaces, the default interface type of NON\_BROADCAST is assigned.

Reference:

[http://www.cisco.com/en/US/tech/CK365/technologies\\_configuration\\_example09186a0080094054.shtml](http://www.cisco.com/en/US/tech/CK365/technologies_configuration_example09186a0080094054.shtml)

### QUESTION 70:

Router CK1 has been configured for OSPF on all of the interfaces. A logical loopback interface has not been created in CK1. By default, what will determine the OSPF router ID of router CK1 ?

- A. It defaults to 255.
- B. The name set by the hostname command.
- C. The highest IP address configured in the router.
- D. The priority number of the router set by the priority command.
- E. The lowest IP address configured in the router.

Answer: C

Explanation:

The show ip ospf interface command verifies that interfaces have been configured in the intended areas. If no loopback address is specified, the interface with the highest address is the taken router ID. It also gives the timer intervals, including the hello interval, and shows the neighbor adjacencies.

Note: the router ID used is not pre-emptive, meaning that if a loopback address is configured or another interface is configured with a higher IP address, the router ID will not change until it is restarted.

Reference: Building Scalable Cisco Networks (Cisco Press) page 134

---

### **QUESTION 71:**

While troubleshooting some connectivity issues, you issue the "show ip ospf database" in order to examine the link state database. Which three of the statements below are true regarding the OSPF link state database? (Select three)

- A. Each router has an identical link state database.
- B. External routes are imported into a separate link state database.
- C. Synchronization of link state databases is maintained via flooding of LSAs.
- D. Information in the link state database is used to build a routing table by calculating a shortest-path tree.
- E. By default, link state databases are refreshed every 10 minutes in the absence of topology changes.

Answer: A, C, D

Explanation:

The Link state database is a collection of link state advertisement for all routers and networks. Each router in the OSPF network maintains an identical database. LSA flooding occurs whenever there is a change in the OSPF topology, ensuring that the databases are synchronized. OSPF also uses the SPF algorithm to build the database tables.

Incorrect Answers:

- B. Only one link state database is maintained, and it is used for all OSPF routes.
- E. The default refresh time is 30 minutes.

Reference: Building Scalable Cisco Networks (Cisco Press) page 178.

---

**QUESTION 72:**

When designing OSPF networks, it is not recommended to have an excessive number of routers in any single area. Which two of the following statements describe the consequences of having too many routers in an area? (Select two)

- A. Excessive LSA traffic.
- B. Frequent routing table recalculation.
- C. Frequent adjacencies table recalculation.
- D. Unacceptable number of reachability errors.

Answer: A, B

Explanation:

In OSPF, if too many routers are in the same area they spend too much time communicating with each other, because anytime a topology change in any single route occurs, LSAs are flooded to all other routers within the area. In OSPF the amount of overhead used increases exponentially with the number of routers in an area. Too many routers can cause an overwhelming amount of LSA and OSPF database calculations.

Incorrect Answers:

- C: The number of routers in an area has little or no impact with the number of neighbor adjacencies that must remain established for any given router.
- D: An excessive number of routers in any single area will not alone increase the number of errors within the network.

---

**QUESTION 73:**

Router CK1 has OSPF configured on its fast Ethernet interface. What is the OSPF cost associated with this 100Mbps Ethernet connection?

- A. 1
- B. 6
- C. 10
- D. 100

Answer: A

Explanation:

The OSPF cost of an interface is inversely proportional to the bandwidth of that interface. A higher bandwidth indicates a lower cost. There is a higher cost and time delay involved in crossing a 56k serial line than crossing a 10M ethernet line. The formula used to calculate the cost is of any single OSPF network is:

Cost= 100,000,000/bandwith in bps.

Since a 100M Ethernet link is 100,000,000 bps we have:

$100000000 / 100000000 = 1$

Reference:

<http://www.cisco.com/warp/public/104/2.html>

---

**QUESTION 74:**

OSPF is a routing protocol that supports variably subnetted IP networks. By what means is it able to support VLSM?

- A. By using route summarization.
- B. By maintaining a topological database.
- C. By carrying subnet mask information in the route updates.
- D. By allocating addresses in groups to support multiple areas.

Answer: C

Explanation:

Each route update includes subnet mask information, which is passed on to the other OSPF routers within the network.

Incorrect Answers

- A: VLSM allows route summarization, but VLSM does not use route summarization.
  - B: Although a topology table is maintained, a topological database does not contain VLSM information.
  - D: This does not describe the way OSPF supports VLSM..
- 

**QUESTION 75:**

Which of the following statements regarding OSPF and VLSM support are true?  
(Select two)

- A. A multi-area OSPF design dictates that VLSM be used throughout the AS.
- B. Summarization can be performed with different prefix lengths throughout the network.
- C. The use of VLSM enables a truly hierarchical addressing scheme to be developed.
- D. OSPF's support for VLSM does not compensate for its inability to handle discontinuous subnets.

Answer: B, C

Explanation:

B: Summarization can be performed with different prefix lengths throughout the network, as long as the ranges don't overlap.

C: VLSM enables a true hierarchical addressing scheme.

Incorrect Answers

- A: VLSM can be used locally only. VLSM is not required to be used throughout multiple areas.
- D: In order to take advantage of summarization, network numbers in areas should be assigned in a contiguous way to be able to lump these addresses into one range.

Reference: OSPF Design Guide

<http://www.cisco.com/warp/public/104/3.html>

---

**QUESTION 76:**

In order to reduce the overall size of the routing table, summarization is being configured throughout the Certkiller network. What is true about route summarization in OSPF? (Select the best answer)

- A. Type 3 and type 4 LSAs carry external summarized routes.
- B. Summarization prevents type 1 link LSAs from being propagated into the backbone area0.
- C. Route summarization can be performed at any point in the network where enough contiguous addresses are present.
- D. Route summarization reduces the amount of bandwidth, CPU, and memory resources consumed by the OSPF process.

Answer: D

Explanation:

Summarization reduces the overall number of routes within the network, which will increase performance by reducing the workload on the router.

---

**QUESTION 77:**

OSPF is considered to be an extremely scalable routing protocol. Which of the following would be reasons for this? (Select two)

- A. Neighbor adjacencies control distribution of routing protocol updates.
- B. Routing table information does not flood the network until holddown timers have expired.
- C. The hello protocol is a more efficient means of sending routing updates than table exchange used in RIPv1.
- D. Topological database is maintained with incremental updates, with full exchange occurring only every 30 minutes.

Answer: C, D

Explanation:

C: The Hello Protocol used by OSPF to establish and maintain neighbor relationship. These periodic hellos are much more scalable than the entire routing updates used by RIP.

D: Incremental updates are used to maintain the OSPF topology. By default, the entire SPF algorithm is only run every 30 minutes.

Incorrect Answers:

A: LSA flooding occurs whenever a topology change occurs in the OSPF network. The

DR controls this distribution, not the neighbor relationship.

B: LSA flooding occurs immediately, after each topology change within an area.

---

**QUESTION 78:**

Router CK1 is configured for OSPF on all of its interfaces. One of the links on this router goes down, causing a link state change. Assuming that router CK1 lies within an Ethernet segment, how will it respond to this link state change?

- A. It will broadcast the LSA on each of its interfaces.
- B. It multicasts the link-state changes to the DR and BDR.
- C. It will immediately flood the change to all outer routers on the same segment.
- D. It will update its routing table, and then floods the change to all other routers on the same segment.

Answer: B

Explanation:

The steps for OSPF convergence are as follows:

1. When a router detects a link failure, the router sends an LSA to its neighbors. If the router is on a multi-access link, it sends the update to the designated router (DR) and the backup designated router (BDR), not to all neighbors.
2. The path is removed from the originating router's tables.
3. On receipt of the LSA, all routers update the topology table and flood the LSA out its interfaces.
4. The routing protocol runs the Dijkstra algorithm to rebuild the routing table.

For OSPF, convergence is detection time, plus LSA flooding, plus 5 seconds before computing the topology table. This amounts to a few seconds.

---

**QUESTION 79:**

With which other routers on an Ethernet LAN should the DROTHER routers form OSPF full adjacency?

- A. the DR and BDR only
- B. all OSPF routers on the Ethernet LAN
- C. all OSPF routers in the area
- D. all other DROTHER routers on the Ethernet LAN
- E. the DR only

Answer: A

Explanation:

DROTHER - The router was not chosen as the DR or the BDR. If the priority on the interface had been set to zero, the state would always be DROTHER because the router

could not be elected as a DR or a BDR. Each router should form the full adjacency only with DR and BDR to receive the OSPF updates from them.

---

**QUESTION 80:**

What does an OSPF router do if it receives an LSA with a lower sequence number than what is already in its link-state database?

- A. Updates the database with the LSA information.
- B. Ignores the LSA
- C. Forwards the LSA to the DR
- D. Ignores the LSU and sends the source router the more current LSA info in its own database.

Answer: D

Explanation:

If the LSA entry was in the database, but the LSA that has just been received has an older sequence number, the router asks whether the information in the database is the same. If the information is the same and the new LSA has an older sequence number, the process discards the packet. It might be old news, but there is no inconsistency in the database.


If the information is different and the newly received LSA has an older sequence number, however, the receiving router discards the LSA update. It issues a copy of the LSA it has in its database, sending it out of the receiving interface to the source address of the out-of-date LS

A. The logic is that the sending router has bad or old information and must be updated because its topological database is obviously not synchronized with the rest of the area.

---

**QUESTION 81:**

The Certkiller OSPF network is displayed below:



| Neighbor ID | Pri | State          | Dead Time | Address     | Interface |
|-------------|-----|----------------|-----------|-------------|-----------|
| 192.168.1.2 | 1   | output omitted | 00:00:34  | 192.168.1.2 | Ethernet0 |
| 192.168.1.4 | 1   | output omitted | 00:00:32  | 192.168.1.4 | Ethernet0 |
| 192.168.1.3 | 1   | output omitted | 00:00:39  | 192.168.1.3 | Ethernet0 |

Refer to the exhibit. All routers have converged and neighbor relationships have been established. Which state is Certkiller 4 in?

- A. 2WAY/DROTHER
- B. 2WAY/BDR
- C. 2WAY/DR
- D. FULL/DROTHER
- E. FULL/BDR
- F. FULL/DR


Answer: F

Explanation:

The router with the highest OSPF priority on a segment will become the DR for that segment. The same process is repeated for the BDR. In case of a tie, the router with the highest RID will win. The default for the interface OSPF priority is one. Based on the output above, the dead timer shows that the neighbor relationships have already been established, so the state will be Full for all neighbors, and since the OSPF priorities are all set to 1, the router with the highest router ID will be the DR.

## QUESTION 82:

The Certkiller backbone OSPF network is shown below:



```
Certkiller1 (config)# router ospf 10
Certkiller1 (config-router)# router-id 1.1.1.1
```

Refer to the exhibit. All routers have been simultaneously rebooted, and neighbor

relationships have been established. Which two statements are true? (Choose two.)

- A. The Certkiller 1 RID is 1.1.1.1.
- B. The Certkiller 1 RID is 10.1.1.1.
- C. The Certkiller 1 RID is 192.168.1.1.
- D. The Certkiller 2 RID is 10.1.1.2.
- E. The Certkiller 2 RID is 192.168.1.2.
- F. The Certkiller 2 RID is 192.168.2.1.

Answer: A, D

Explanation:

The Router-ID will be taken from the loopback address first, if no loopback exists then the highest IP address on the router will be used, UNLESS manually set in the router ospf process, if this is done the router does not have to 'look' for address at all as it has one already.

As it is set manually then answer A is right not answer B.

---

**QUESTION 83:**

Which of the three areas below are valid OSPF area types? (Select three)

- A. Stub
- B. Active
- C. Remote
- D. Backbone
- E. Ordinary or standard

Answer: A, D, E

Explanation:

There is no such thing as an active area or a remote area in OSPF. The correct OSPF area types are:

- \*ordinary area (standard area)
- \*stub area
- \*totally stubby area
- \*not so stubby area (NSSA)
- \*backbone area (area 0)

---

**QUESTION 84:**

OSPF networks create Link State Advertisements (LSAs). What kind of LSA is conceived by an Area Border Router (ABR) containing route summarization of area subnets?

- A. Router link, type 1.

- B. Router link, type 8.
- C. Network summary link, type 3.
- D. AS external summary link, type 4.

Answer: C

Explanation:

The network summary link entry is a LSA type 3 that is originated by ABRs. It describes the links between the ABR and the internal routers of a local area. These entries are flooded throughout the backbone area to the other ABRs.

Incorrect Answers

- A: Router link, LSA type 1 is generated by each router for each area it belongs to. It describes the states of the router's link to the area.
- B: Router LSA is type 1.
- D. External link LSA is type 5.

---

**QUESTION 85:**

Router CK1 has an interface in one OSPF area and another interface in a different OSPF area. What kind of router is CK1 ?

- A. ABR
- B. ASBR
- C. internal router
- D. backbone router

Answer: A

Explanation:

An ABR (Area Border Router) shares an interface with at least one other OSPF area.

Incorrect Answers

- B: An ASBR (Autonomous System Border Router) have at least one interface in a non-OSPF network.
- C: An internal router is only connected to routers in the internal area.
- D: A backbone router has at least one interface in the backbone, also known as Area 0.

---

**QUESTION 86:**

When an OSPF ASBR announces external networks, what kind of LSA type does it use?

- A. Type 1 LSA
- B. Type 2 LSA
- C. Type 3 LSA
- D. Type 4 LSA
- E. Type 5 LSA

Answer: E

Explanation:

In OSPF, Type 5 LSA is used to announce external networks.

The OSPF Router LSA types are described here:

Router LSA (type 1) - information about the router and its directly connected links; flooded only w/in the area

Network LSA (type 2) - information about a LAN and the routers connected to it; originated by the DR; flooded only w/in the area

Summary LSA (type 3) - describe networks reachable outside the area; originated by the area border routers (ABR)

ASBR Summary LSA (type 4) - reachability information to the Autonomous System Border

Routers outside the area; originated by the ABRs


External LSA (type 5) - information about destinations outside the OSPF domain; originated by the ASBRs; flooded throughout the whole network

Other LSA types exist for multicast information (type 6) and other extensions, including NSSA area (type 7).

---

**QUESTION 87:**

The Certkiller OSPF network is displayed in the diagram below:



The status of external Network 22 has just changed. The router Certkiller 5 uses a type 7 LSA to update area 1 of this change. In OSPF terminology, what kind of area is area 1 classified as?

A. Backbone area

- B. Transit area
- C. Stubby area
- D. Totally stubby area
- E. Not-so-stubby area

Answer: E

Explanation:

NSSA use type 7 LSAs. The different OSPF area types are described below.

1. Stub Areas: These areas do not accept routes belonging to external autonomous systems (AS); however, these areas have inter-area and intra-area routes. In order to reach the outside networks, the routers in the stub area use a default route which is injected into the area by the Area Border Router (ABR).
2. Normal Areas: These areas can either be standard areas or transit (backbone) areas. Standard areas are defined as areas that can accept intra-area, inter-area and external routes.
3. Backbone area is the central area to which all other areas in OSPF connect.
4. Totally Stub Areas: These areas do not allow routes other than intra-area and the default routes to be propagated within the area. The ABR injects a default route into the area and all the routers belonging to this area use the default route to send any traffic outside the area.
5. NSSA: This type of area allows the flexibility of importing a few external routes into the area while still trying to retain the stub characteristic. Assume that one of the routers in the stub area is connected to an external AS running a different routing protocol, it now becomes the ASBR, and hence the area can no more be called a stub area. However, if the area is configured as a NSSA, then the ASBR generates a NSSA external link-state advertisement (LSA) (Type-7) which can be flooded throughout the NSSA area. These Type-7 LSAs are converted into Type-5 LSAs at the NSSA ABR and flooded throughout the OSPF domain

Reference:

[http://www.cisco.com/en/US/tech/ CK3 65/ CK4 80/technologies\\_tech\\_note09186a0080094a74.shtml](http://www.cisco.com/en/US/tech/ CK3 65/ CK4 80/technologies_tech_note09186a0080094a74.shtml)

---

### **QUESTION 88:**

What kind of information can you find in a type 3 OSPF network summary link LSA?

- A. Summary of routes in the AS.
- B. Summary of link states in an OSPF area.
- C. Summary of IP subnets in an OSPF area.
- D. Summary of metric cost from ABR to ASBR.
- E. None of the above.

Answer: C

Explanation:

A type 3 area LSA is Inter-area prefix LSA for ABR It advertises internal networks to routers in other areas (interarea routes). Type 3 LSAs may represent a single network or a set of networks summarized into one advertisement. Only ABRs generate summary LSAs. In OSPF for IPv6, addresses for these LSAs are expressed as prefix, prefix length instead of address, mask. The default route is expressed as a prefix with length 0.

Reference:

[http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products\\_feature\\_guide\\_chapter09186a0080145c56.html](http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide_chapter09186a0080145c56.html)

**QUESTION 89:**

**DRAG DROP**

Drag the OSPF area type icons on the left, and drop them into the centre target area that adjoins the corresponding description on the right.

| OSPF area type, use these |            | Descriptions                             |
|---------------------------|------------|--|
| Stub Area                 | place here | interconnects areas                      |
| Transit Area              | place here | uses type 7 LSAs                         |
| Totally Stubby area       | place here | does not accept external LSAs            |
| Backbone Area             | place here | interconnects discontinuous backbones    |
| Not-so-stubby Area        | place here | does not accept external or summary LSAs |

Answer:

| OSPF area type, use these |  | Descriptions                             |
|---------------------------|--|--|
| Backbone Area             |  | interconnects areas                      |
| Not-so-Stubby area        |  | uses type 7 LSAs                         |
| Totally Stubby area       |  | does not accept external or summary LSAs |
| Transit Area              |  | interconnects discontinuous backbones    |
| Stub Area                 |  | does not accept external LSAs            |

Explanation:

1. Stub Areas: These areas do not accept routes belonging to external autonomous systems (AS); however, these areas have inter-area and intra-area routes. In order to reach the outside networks, the routers in the stub area use a default route which is injected into the area by the Area Border Router (ABR).
2. Normal Areas: These areas can either be standard areas or transit (backbone) areas. Standard areas are defined as areas that can accept intra-area, inter-area and external routes.
3. The Backbone area is the central area to which all other areas in OSPF connect.
4. Totally Stub Areas: These areas do not allow routes other than intra-area and the default routes to be propagated within the area. The ABR injects a default route into the area and all the routers belonging to this area use the default route to send any traffic outside the area.

5. NSSA: This type of area allows the flexibility of importing a few external routes into the area while still trying to retain the stub characteristic. Assume that one of the routers in the stub area is connected to an external AS running a different routing protocol, it now becomes the ASBR, and hence the area can no more be called a stub area. However, if the area is configured as a NSSA, then the ASBR generates a NSSA external link-state advertisement (LSA) (Type-7) which can be flooded throughout the NSSA area. These Type-7 LSAs are converted into Type-5 LSAs at the NSSA ABR and flooded throughout the OSPF domain

Reference:

[http://www.cisco.com/en/US/tech/CK365/CK480/technologies\\_tech\\_note09186a0080094a74.shtml](http://www.cisco.com/en/US/tech/CK365/CK480/technologies_tech_note09186a0080094a74.shtml)

---

### **QUESTION 90:**

In a multi-area OSPF network inter-area summarization produces type's 3 & 4 Link State Advertisements (LSAs). At which router are these LSA types produced?

- A. At the ABR.
- B. At the ASBR.
- C. At the backbone DR.
- D. Each internal router calculates and sends network updates.
- E. None of the above.

Answer: A

Explanation: Network summary link entries, LSA type 3 or LSA type 4 are originated by ABRs. They describe the links between the ABR and the internal routers of a local area. These entries are flooded throughout the backbone area to the other ABRs.

Note: Type-3 describes routes to networks within the local area and are sent to the backbone area. Type-4 describes reachability to ASBRs. These link entries are not flooded through totally stubby areas.

Incorrect Answers:

B: Autonomous system external link entries, LSA type 5, are originated by ASBRs. They describe routes to destinations external to the autonomous system.

C: Backbone DRs does not generate these LSAs.

D: Internal routers produce router link entries, LSA type 1. They describe the states of the router's link to the area.

---

### **QUESTION 91:**

What does Cisco recommend for those who set up networks with multiple OSPF areas? (Select two)

- A. There should not be more than three areas per route.
- B. Area 0 must be larger than any subsequent OSPF area.
- C. A router can be a DR or BDR for more than one LAN.

D. You should not run more than one instance of the OSPF process on an ABR.

Answer: C, D

Explanation:

C: A router can be a DR or BDR for more than one LAN. Since the loopback address is normally used as the router ID, this ID is used for all LAN segments.

D: Running multiple OSPF processes on the same router is not recommended because it creates multiple database instances that add extra overhead.

Incorrect Answers:

A: There is no such limit. The hierarchical structure could be larger.

B: There is no requirement that Area 0 must be the largest OSPF area.

---

**QUESTION 92:**

When designing and implementing OSPF in multiple areas; what should be taken in consideration? (Select two)

A. Each area must connect to area 0.

B. Each area must have a unique AS number.

C. Remote areas must be configured as stub or NSSA areas.

D. Traffic between two areas must travel across the backbone area.

Answer: A, D

Explanation:

In OSPF, all areas must be connected to a backbone area, which is area 0. It is also required that traffic from one area to another non-backbone area must traverse area 0, since all areas must connect to area 0.

Incorrect Answers

B: OSPF routers within the same Autonomous System should be configured with the same AS number. The AS number is also called the OSPF process ID.

C: Remote areas are not required to be either stub or not so stubby areas.

---

**QUESTION 93:**

In OSPF, a Designated Router (DR) floods LSAs to all the routers on the segment that it has a neighbor relationships with. Over what kind of link does this happen and what kind of LSA type is sent?

A. Router link, type 1.

B. Network link, type 2.

C. External link, type 5.

D. Network summary link, type 3.

E. None of the above

Answer: B

Explanation:

A network Link LSA type 2 is generated by DRs in multi-access networks. They describe the set of routers attached to a particular network. Flooding occurs within the area that contains the network only.

Incorrect Answers:

A: Router link, LSA type 1 is generated by each router for each area it belongs to. It describes the states of the router's link to the area.

C: External link, LSA type 5 is originated by the ASBR. It describes routes to destinations external to the autonomous system. It is flooded throughout an OSPF autonomous system except for stub and totally stubby areas.

D: Network summary link entry, LSA type 3 is originated by ABRs. It describes the links between the ABR and the internal routers of a local area. These entries are flooded throughout the backbone area to the other ABRs.

---

**QUESTION 94:**

In an OSPF network, neighbor relationships allow the topology of the network to scale. What are two reasons as to why this can help a network topology scale? (Select two)

- A. Neighbor adjacencies control distribution of routing protocol updates.
- B. Routing table information does not flood the network until holddown timers have expired.
- C. The hello protocol is a more efficient means of sending routing updates than table exchanges used in RIPv1.
- D. Topological database is maintained with incremental updates, with full exchange occurring only every 30 minutes.

Answer: A, C

Explanation:

When a network topology changes and becomes more sophisticated, fewer routers are burdened with the responsibility of having to relearn the entire network. The hello protocol is also efficient because rather than focusing on knowing what all the routers know, the routers only focus on who is still around, therefore available for routing.

---

**QUESTION 95:**

In an OSPF network; which IP address multicasts all OSPF Designated Routers and Backup Designated Routers?

- A. 224.0.0.5
- B. 224.0.0.6
- C. 224.0.0.9
- D. 224.0.0.11

E. 224.0.0.2

Answer: B

Explanation:

224.0.0.6 is the IP address used by all OSPF DRs and BDRs.

Incorrect Answers:

A: This is the multicast address used by all OSPF routers.

C: This address is used by RIP version 2 routers.

D: This multicast address is used by mobile agents.

E: This multicast IP address is used by all routers on the subnet.

Reference: Building Scalable Cisco Networks (Cisco Press) page 114

---

**QUESTION 96:**

What kind of OSPF Link State Advertisement (LSA) is found in a type 3 network summary link?

- A. Summary of routes in the AS.
- B. Summary of link states in an OSPF area.
- C. Summary of IP subnets in an OSPF area.
- D. Summary of metric costs from ABR to ASBR.

Answer: C

Explanation:

Type 3 LSAs are INTER-AREA advertisements and they are not used within the local area.

The various LSA types used in OSPF are described below:

LS Type Link State ID \_\_\_\_\_ 1 The originating router's Router ID. 2 The IP interface address of the network's Designated Router. 3 The destination network's IP address. 4 The Router ID of the described AS boundary router. 5 The destination network's IP address.

---

**QUESTION 97:**

You are the administrator of the large Certkiller OSPF and are considering dividing it to multiple areas. What are two advantages of configuring multiple OSPF areas? (Select two)

- A. It eliminates the need for a DR or BDR in the OSPF network.
- B. It eliminates security concerns by segregating portions of the network.
- C. Type 1 and 2 LSAs are confined to a single area, reducing routing overhead.
- D. Area members have smaller topological databases than if the network was one large area, requiring less CPU to derive routes.

Answer: C, D

Explanation:

C: Link State update (LSU) traffic is reduced. Rather than send an LSU about each network within an area, you can advertise a single or fewer summarized routes between areas to reduce the overhead associated with linkstate updates when they are crossing areas.

Note: All LSA types, except the AS-external-LSAs (LS type = 5), are flooded throughout a single area only.

D: Topological database size would decrease and fewer CPU cycles would be required to calculate routes.

Incorrect Answers:

A: A DR and BDR for each area would still be required.

B: Security concerns would not be eliminated. There are no security advantages in using more than one area in an OSPF network.

---

**QUESTION 98:**

OSPF has been configured on router CK1 , which lies on a multi point network. On this multi-access network, what type of router floods type 2 Link State Advertisements (LSAs) into an OSPF area?

- A. DR
- B. ABR
- C. BDR
- D. ASBR

Answer: A

Explanation:

Type 2 LSAs are generated by Designated Routers (DRs) in multi-access networks. They describe the set of routers attached to a particular network and are flooded within the area that contains the network only.

Note: Instead of routers exchanging link-state information with every other router on the segment, each router sends the link-state information to the DR and Backup Designated Router (BDR). The DR sends each router's link-state information to all other routers in the network. This flooding process significantly reduces the router-related traffic on a segment.

Incorrect Answers:

B: Area Border Routers generate Type 3 and Type 4 LSAs.

C: BDRs are just used for backup, should the DR fail.

D: Autonomous System Boundary Routers (ASBRs) generate Type 5 LSAs.

---

**QUESTION 99:**

You want router one of the OSPF areas on your network to be able to accept type 7

external routers, but not type 5 external routes. What kind of area should be configured to accomplish this?

- A. A not-so stubby area
- B. A stubby area
- C. A totally stubby area
- D. A backbone area
- E. An on-demand area

Answer: A

Explanation:

The NSSA external LSA

-Identified as Type 7, these LSAs are created by the ASBR residing in a not so stubby area (NSSA). This LSA is similar to an autonomous system external LSA, except that this LSA is contained within the NSSA area and is not propagated into other areas, but it is converted into a Type 5 LSA by the ABR.

An NSSA is an area that is seen as a stub area but can receive external routes, which it will not propagate into the backbone area and thus the rest of the OSPF domain. Another LSA, Type 7, is created specifically for the NSSA

A. This LSA can be originated and communicated throughout the area, but it will not be propagated into other areas, including Area 0. If the information is to be propagated throughout the AS, it is translated into an LSA Type 5 at the NSSA ABR.

---

**QUESTION 100:**

DRAG DROP

Match the OSPF area type on the left to the type of OSPF routes (as would be seen in the routing table) it will support on the right hand side:

|                      |                |            |
|----------------------|----------------|------------|
| O, O* IA             | Stub           | Place here |
| O, O IA, O N2, O* N2 | Totally Stubby | Place here |
| O, O IA, O*IA        | NSSA           | Place here |

Answer:

|                |                      |
|----------------|----------------------|
| Stub           | O, O IA, O*IA        |
| Totally Stubby | O, O* IA             |
| NSSA           | O, O IA, O N2, O* N2 |

---

**QUESTION 101:**

Which of the following OSPF routes are supported by the Totally Stubby type?  
(Select two)

- A. O
- B. O IA
- C. O N2
- D. O\* N2
- E. O\* IA

Answer: A, E

Explanation:

The only routes that are seen in a totally stubby area are intra-area routes (O) and the default Intra area route (O\*IA).

Example routing table of a totally stubby router:

```
CK1 #show ip route
```

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP

D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP

i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, \* - candidate

default

Gateway of last resort is not set

203.250.15.0 255.255.255.252 is subnetted, 1 subnets

C 203.250.15.0 is directly connected, Serial0

131.108.0.0 255.255.255.240 is subnetted, 1 subnets

O 131.108.79.208 [110/74] via 203.250.15.1, 00:31:27, Serial0

O\*IA 0.0.0.0 0.0.0.0 [110/74] via 203.250.15.1, 00:00:00, Serial0

Note that the only routes that show up are the intra-area routes (O) and the default-route 0.0.0.0. The external and inter-area routes have been blocked. The cost of the default route is now 74 (64 for a T1 line + 10 advertised by RTC). No configuration is needed on RTE in this case. The area is already stub, and the no-summary command does not affect the Hello packet at all as the stub command does.

---

### **QUESTION 102:**

Which type of OSPF network requires manual configuration of OSPF neighbors?

- A. broadcast multi-access
- B. nonbroadcast multi-access
- C. point-to-point
- D. point-to-multipoint

Answer: B

Explanation:

Non-broadcast multi-access networks, as found in frame-relay and ATM, require the neighbors to be explicitly specified. This is the case where the PVCs are configured on the physical interface using frame map commands. When sub-interfaces are used, the neighbor commands do not need to be configured.

Note: On the physical interface of an NBMA network, the command "ip ospf network-type point-to-multipoint" command can be used instead of specifying the neighbors individually.

---

**QUESTION 103:**

Which of the following OSPF routes are supported by the Stub Area type? (Select three)

- A. O
- B. O IA
- C. O N2
- D. O\* N2
- E. O\* IA

Answer: A, B, E

Explanation:

When an area is configured as a stub area, it will learn about all routes, except for the external routes.

Example routing table of a stub router:

```
CK2 #show ip route
```

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP

D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP

i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, \* - candidate

default

Gateway of last resort is 203.250.15.1 to network 0.0.0.0

203.250.15.0 255.255.255.252 is subnetted, 1 subnets

C 203.250.15.0 is directly connected, Serial0

O IA 203.250.14.0 [110/74] via 203.250.15.1, 00:26:58, Serial0

128.213.0.0 255.255.255.252 is subnetted, 1 subnets

O IA 128.213.63.0 [110/84] via 203.250.15.1, 00:26:59, Serial0

131.108.0.0 255.255.255.240 is subnetted, 1 subnets

O 131.108.79.208 [110/74] via 203.250.15.1, 00:26:59, Serial0

O\*IA 0.0.0.0 0.0.0.0 [110/65] via 203.250.15.1, 00:26:59, Serial0

Note that all the routes show up except the external routes which were replaced by a default route of 0.0.0.0.

---

**QUESTION 104:**

Which of the following OSPF routes are supported by the Not-So-Stubby-Area

(NSSA) type? (Select three)

- A. O N1
- B. O IA
- C. O N2
- D. O\* N2
- E. O\* IA

Answer: B, C, D

Explanation:

The various route types used by OSPF are as follows:

O - OSPF

IA - OSPF inter area

N1 - OSPF NSSA external type 1

N2 - OSPF NSSA external type 2

E1 - OSPF external type 1


E2 - OSPF external type 2

An OSPF NSSA will receive inter-area, external type 2, and default external type 2 routes.

---

**QUESTION 105:**

Routers CK1 , CK2 , and CK3 are part of OSPF area 100 as shown below:



In the Certkiller network shown above, router CK1 was elected the DR and CK2 was elected the BDR. Assume that CK1 fails, and that CK2 takes the place of the DR while CK3 becomes the new BDR. What will happen when CK1 comes back online?

- A. CK1 will take the place of DR immediately upon establishing its adjacencies.
- B. CK1 will take the place of DR only if CK2 fails.
- C. CK1 will take the place of DR only if both CK2 and CK3 fail.
- D. A new election will take place establishing an all new DR and BDR based on configured priority levels and MAC addresses.

Answer: C

Explanation:

If a router with a higher priority value gets added to the network, it does not preempt the DR and BDR. The only time a DR and BDR changes is if one of them is out of service. If the DR is out of service, the BDR becomes the DR, and a new BDR is selected. If the BDR is out of service, a new BDR is elected. In a multi-access network, the router that is powered on first will generally become the DR, since the DR/BDR process is not pre-emptive.

Reference:: CCNP Self-Study Second Edition P.243

---

**QUESTION 106:**

Which of the following are features are characteristics of an OSPF ABR? (Choose three)

- A. maintains one link-state database for all areas
- B. routes traffic to and from other areas
- C. maintains a separate link-state database for each area
- D. originates type 3 link-state advertisements which can be summarized
- E. has at least one physical interface in area 0.
- F. has at least one interface in an external internetwork.

Answer: B, C, D


Explanation:

ABR forwards summary LSAs. It forwards both Type 3 LSAs and Type 4 LSAs. Type 3 LSAs are forwarded to the other ABRs, and Type 4 LSAs are forwarded to the ASBRs. ABR also forwards Type 3 LSAs from other areas into its own area. If the ABR has multiple links in the same area, it also forwards Type 1 and Type 2 LSAs in its capacity as an internal router.

---

**QUESTION 107:**

The Certkiller OSPF network is shown in the diagram below:



Which type of Link State Announcement will Certkiller E use to announce changes in network 22 to Area 1?

- A. Type 1 LSA
- B. Type 2 LSA
- C. Type 3 LSA
- D. Type 4 LSA
- E. Type 5 LSA

Answer: E

Explanation:

The external link LSA is originated by AS boundary routers and is flooded throughout the AS. Each external advertisement describes a router to a destination in another autonomous system. Default routes for the AS can also be described by AS external advertisements. This is identified as a Type 5 LSA.

---

**QUESTION 108:**

What is the OSPF cost for the 100 Mbps Ethernet media type?

- A. 1
- B. 2
- C. 6
- D. 10
- E. 100

Answer: A

Explanation:

By default, OSPF uses the formula of 100,000,000 (100M) divided by bandwidth, or 100,000,000/BW. For a 100M fast Ethernet connection, the formula is 100M/100M = 1.

---

**QUESTION 109:**

Which of the following OSPF routes are supported by the Not-So-Stubby-Area (NSSA) type? (Select three) (801)

- A. O
- B. O N2
- C. O\* N2
- D. O E1
- E. O E2

Answer: A, B, C

Explanation:

The various route types used by OSPF are as follows:

O - OSPF

IA - OSPF inter area

N1 - OSPF NSSA external type 1

N2 - OSPF NSSA external type 2

E1 - OSPF external type 1

E2 - OSPF external type 2

An OSPF NSSA will receive inter-area, external type 2, OSPF routes.

---

**QUESTION 110:**

Which three are advantages to creating multiple areas in OSPF? (Select three)

- A. less frequent SPF calculations
- B. fewer hello packets
- C. smaller routing tables
- D. reduced LSU overhead
- E. fewer adjacencies needed

Answer: A, C, D

Explanation:

OSPF employs the concept of an area, which is a grouping of contiguous OSPF networks and hosts. OSPF areas are logical subdivisions of OSPF autonomous systems whose internal topology is hidden to routers outside the area. Areas allow an additional level of

hierarchy different from that provided by IP network classes, and they can be used to aggregate routing information and mask the details of a network. These features make OSPF particularly scalable to large networks. When using areas, OSPF routers need to maintain information only within their own area so less frequent SPF calculations are required, smaller routing tables occur with area summarizations, and the link state update overhead is reduced.

---

**QUESTION 111:**

Using OSI terms, what type of system is a router?

- A. end system
- B. routing system
- C. inter-domain system
- D. intermediate system
- E. intra-domain system

Answer: D

Explanation:

Data communication sessions generally occur between end stations, or end hosts. Routers are considered to be intermediate systems, since they are used as a means for transporting data between end stations. An easy way to remember this is the OSI routing protocol IS-IS, which stands for "Intermediate System to Intermediate System."

---

**QUESTION 112:**

The network layer of the OSI protocol suite defines two of the routing protocols below. Which two are they? (Select two)

- A. End System-to-End System
- B. Routing Information Protocol
- C. Interior Gateway Routing Protocol
- D. Enhanced Interior Gateway Routing Protocol
- E. Intermediate System-to-Intermediate System
- F. End System-to-Intermediate System

Answer: E, F

Explanation:

Both are defined by ISO in the OSI protocol suite.

Intermediate System-to-Intermediate System (IS-IS) Protocol is an intradomain Open System Interconnection (OSI) dynamic routing protocol specified in International Organization for Standardization (ISO) 10589. The protocol is designed to operate in OSI Connectionless Network Service (CLNS). Data is carried using the protocol specified in ISO 8473.

OSI CLNS is a network layer service similar to bare IP service. A CLNS entity communicates over Connectionless Network Protocol (CLNP) with its peer CLNS entity.

Incorrect Answers:

A: This is not a routing protocol.

B: RIP is a standards based routing protocol, but it is not defined by the OSI or ISO organizations.

C, D: These answer choices refer to the Cisco proprietary routing protocols, IGRP and EIGRP.

Reference:

[http://www.cisco.com/en/US/tech/CK365/CK381/technologies\\_white\\_paper09186a00800a3e6f.shtml](http://www.cisco.com/en/US/tech/CK365/CK381/technologies_white_paper09186a00800a3e6f.shtml)

---

### **QUESTION 113:**

Which of the following routing protocols belongs in the network layer, as defined by the OSI protocol suite?

- A. End System-to End System
- B. Routing Information Protocol
- C. Interior Gateway Routing Protocol
- D. Enhanced Interior Gateway Routing Protocol
- E. Intermediate System-to-Intermediate System
- F. None of the above

Answer: E

Explanation:

Intermediate System-to-Intermediate System (IS-IS) Protocol is an intradomain Open System Interconnection (OSI) dynamic routing protocol specified in International Organization for Standardization (ISO) 10589. The protocol is designed to operate in OSI Connectionless Network Service (CLNS). Data is carried using the protocol specified in ISO 8473.

OSI CLNS is a network layer service similar to bare IP service. A CLNS entity communicates over Connectionless Network Protocol (CLNP) with its peer CLNS entity.

Incorrect Answers:

A: This is not a routing protocol.

B: RIP is a standards based routing protocol, but it is not defined by the OSI or ISO organizations.

C, D: These are the Cisco proprietary routing protocols, IGRP and EIGRP.

Reference:

[http://www.cisco.com/en/US/tech/CK365/CK381/technologies\\_white\\_paper09186a00800a3e6f.shtml](http://www.cisco.com/en/US/tech/CK365/CK381/technologies_white_paper09186a00800a3e6f.shtml)

---

### **QUESTION 114:**

In network engineering and computer science the term domain can be interpreted in multiple contexts. How is a domain defined in OSI terminology?

- A. A set of non-routing network nodes.
- B. A contiguously connected area that can reach all other areas.
- C. All devices configured for OSI protocols within an internetwork.
- D. A contiguous set of routers and hosts and the data links that connect them.
- E. Any portion of an OSI network that is under a common administrative authority.
- F. None of the above

Answer: E

Explanation: A domain is any portion of an OSI network that is under a common administrative authority. A domain is a collection of End Systems, Intermediate Systems, and subnetworks operated by a single organization or administrative authority. The components which make up the domain are assumed to interoperate with a significant degree of mutual trust among them, but interoperate with other Administrative Domains in a mutually suspicious manner

Reference: ISO TR 9575

---

**QUESTION 115:**

How is network layer addressing accomplished in the OSI protocol suite?

- A. Internet Protocol address
- B. Media Access Control address
- C. Packet Layer Protocol address
- D. Network Service Access Point address
- E. Authority and Format Identifier address
- F. None of the above

Answer: D

Explanation: OSI network layer addressing is implemented by using two types of hierarchical addresses: network service access point addresses and network entity titles.

A network service access point (NSAP) is a conceptual point on the boundary between the network and the transport layers. The NSAP is the location at which OSI network services are provided to the transport layer. Each transport layer entity is assigned a single NSAP, which is individually addressed in an OSI internetwork using NSAP addresses.

A network entity title (NET) is used to identify the network layer of a system without associating that system with a specific transport layer entity (as an NSAP address does). NETs are useful for addressing intermediate systems (ISs), such as routers, that do not interface with the transport layer. An IS can have a single NET or multiple NETs, if the IS participates in multiple areas or domains.

Reference:

[http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito\\_doc/osi\\_prot.htm#1022230](http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/osi_prot.htm#1022230)

---

**QUESTION 116:**

**DRAG DROP**

Drag the routing protocol characteristics on the left to the corresponding OSPF box in the middle or the corresponding IS-IS box on the right:

| EIGRP term                         | Drop OSPF Characteristics here | Drop IS-IS Characteristics here |
|------------------------------------|--------------------------------|---------------------------------|
| 1000 routers per area              | place here                     | place here                      |
| 50 routers per area                | place here                     | place here                      |
| routers belong to one level 2 area | place here                     | place here                      |
| flexible extension of backbone     | place here                     | place here                      |
| backup designated router           | place here                     | place here                      |
| boundaries lie inside routers      | place here                     | place here                      |
| boundaries lie on links            |                                |                                 |
| forms adjancies with all neighbors |                                |                                 |

Answer:

| EIGRP term | Drop OSPF Characteristics here | Drop IS-IS Characteristics here    |
|------------|--------------------------------|------------------------------------|
|            | 50 routers per area            | 1000 routers per area              |
|            | backup designated routers      | routers belong to one level 2 area |
|            | boundaries lie inside routers  | boundaries lie on links            |
|            | place here                     | forms adjancies with all neighbors |
|            | place here                     | flexible extension of backbone     |

**QUESTION 117:**

The similar routing protocols OSPF and IS-IS share many characteristics. Which four of the characteristics below do they share? (Select four)

- A. link-state database
- B. backbone design using area 0
- C. Shortest Path First (SPF) algorithm
- D. Update, Decision, and Flooding Process
- E. Hello protocol to establish and maintain adjacencies

Answer: A, C, D, E

Explanation:

A: IS-IS uses a link-state database, similar to OSPF. Both OSPF and IS-IS are considered to be link state routing protocols.

C: The IS-IS decision process runs shortest-path-first (SPF) algorithm on the link-state database, and creates the forwarding database.

D: The Update, Decision, and Flooding Process of IS-IS and OSPF are similar.

E: Routers running IS-IS will send hello packets out all IS-IS-enabled interfaces to discover neighbors and establish adjacencies. This is similar to OSPF.

Incorrect Answers:

B: IS-IS does not have a backbone area like the OSPF area 0. The IS-IS backbone is a contiguous collection of Level 2-capable routers, each of which can be in a different area.

Reference: Introduction to Intermediate System-to-Intermediate System Protocol

[http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys\\_wp.htm](http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys_wp.htm)

---

**QUESTION 118:**

OSPF and IS-IS are similar in many ways. In OSPF, the backbone area is area 0. How is an IS-IS backbone area numbered?

- A. Area 0
- B. Area 1
- C. Area 51
- D. Any legal area number.
- E. There is no backbone area number.

Answer: E

Explanation:

IS-IS does not have a backbone area like the OSPF area 0. The IS-IS backbone is a contiguous collection of Level 2-capable routers, each of which can be in a different area.

Reference: Cisco, Introduction to Intermediate, System-to-Intermediate System Protocol

[http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys\\_wp.htm](http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys_wp.htm)

---

**QUESTION 119:**

OSPF and IS-ISD are similar in many ways, but there are some important differences between them. Which three characteristics apply to IS-IS but not to OSPF? (Choose three)

- A. Encapsulates PDUs directly into a data-link frame.
- B. Uses a DIS and a backup DIS to present the pseudo-node on the LAN.
- C. Uses stubby areas to improve network scalability.
- D. Uses a default IOS metric of 10 on each interface.
- E. Runs PRC (Partial Calculations) to calculate IP reachability information.
- F. Uses an on-demand circuit to reduce the hello and LSA flooding across switched WAN links, such as ISDN.

Answer: A, D, E

Explanation:

IS-IS encapsulates protocol data units (PDUs) directly into a data link frame, as outlined

by the OSI standard. With IS-IS each link has a metric of 10, regardless of the actual bandwidth of the link. Finally, the PRC is the software's process of calculating routes without performing an SPF calculation.

Incorrect Answers:

B: On broadcast multi-access networks, a single router is elected as the DIS. There is no backup DIS elected. The DIS is the router that creates the pseudonode and acts on behalf of the pseudonode.

C: Stubby areas are only used by OSPF.

F: On demand circuits are only used by OSPF, to suppress the LSA hello packets from triggering DDR calls.

---

**QUESTION 120:**

IS-IS is often considered an alternative to OSPF in the IP world. Which two statements identify similarities between IS-IS and OSPF? Select two.

- A. support for designated intermediate systems and backup designated intermediate systems
- B. support for multiple areas per router
- C. support for classless routing
- D. support for address summarization between area
- E. support for both DIS adjacencies and neighbor adjacencies

Answer: C, D

Explanation:

Similarities Between Integrated IS-IS and OSPF

They are both link-state protocols and are based on the Dijkstra algorithm of Shortest Path First (SPF). In addition, they both have a two-level hierarchy. OSPF tends to be deployed mostly as an enterprise solution, whereas Integrated IS-IS is used for IP routing in some ISP networks. Both support address summarization between areas.

---

**QUESTION 121:**

Which of the following statements are correct about the differences in IS-IS and OSPF? (Choose three)

- A. IS-IS LSP contains TLV fields and OSPF LSU contains the LSAs.
- B. New additions to the protocol are easily implemented in OSPF but not with IS-IS
- C. For greater fine tuning there are more options available in IS-IS.
- D. OSPF has more area types than does IS-IS.
- E. IS-IS is more CPU-intensive than is OSPF.

Answer: A, C, D

Explanation:

TLV's are a main strength of IS-IS. TLV's provide flexibility and extend the functionality of the protocol. In IS-IS, the TLV fields are variable in length and appended to the various packet formats. This means that IS-IS has advantages over OSPF and the protocol can adapt to the changing needs and advances to technology by simply defining a new TLV.

---

**QUESTION 122:**

There are several varieties of IS-IS routers. Two types of IS-IS provide intra-area routing services, meaning that they are used to route traffic within the same area.. Which of the following two match this description? (Select two)

- A. L1 IS
- B. L1 ES
- C. L2 IS
- D. L2 ES
- E. L1/L2 IS

Answer: A, E

Explanation:

L1 IS and L1/L2 IS routers provide intra-area routing services. A Level 1 router knows the topology only of its own area and has Level 1 or Level1/Level2 neighbors in this area. It has a Level 1 link-state database with all the information for intra-area routing. It uses the closest Level 2-capable router in its own area to send packets out of the area, a scenario that may result in suboptimal routing.

Incorrect Answers:

B, D: End systems are OSI terms used for end hosts, such as PCs or servers. They do not provide for any type of routing functionality.

C: L2 routers make up the backbone of the IS-IS network, and they are used to provide for inter-area routing.

---

**QUESTION 123:**

In OSPF, the backbone area is area 0. How could you describe the backbone area of an IS-IS domain?

- A. A contiguous collection of L1 routers.
- B. A contiguous collection of L1 and L1/L2 routers.
- C. A contiguous collection of L1 and L2 routers.
- D. A contiguous collection of L2 and L1/L2 routers.
- E. None of the above.

Answer: D

Explanation:

IS-IS does not have a backbone area like the OSPF area 0. The IS-IS backbone is a contiguous collection of Level 2-capable routers, each of which can be in a different area. L2 routers are used for inter-area routing, and a L1 router can only route traffic to other areas via an L2 or and L1/L2 IS-IS router.

---

**QUESTION 124:**

Which one of the following statements correctly describes a characteristic of an IS-IS backbone?

- A. A router in all areas.
- B. All routers in one area.
- C. A chain of L2 and L1/L2 IS-IS routers.
- D. An unbroken string of L1 and LS IS-IS routers.
- E. A central area to which all other areas are attached.
- F. None of the above

Answer: C

Explanation:

IS-IS does not have a backbone area like the OSPF area 0. The IS-IS backbone is a contiguous collection of Level 2-capable routers, each of which can be in a different area. L1 routers are used for internal routing within the area and so only L2 and L1/L2 routers form the backbone in IS-IS.

---

**QUESTION 125:**

The routing protocol IS-IS uses areas. What must an administrator connect each IS-IS area to?

- A. Area 0
- B. Area 1
- C. Level-1 backbone
- D. Level-2 backbone
- E. External IS-IS areas

Answer: D

Explanation:

Small IS-IS networks are built as a single area that includes all the routers in the network. As the network grows larger, it is usually reorganized into a backbone area made up of the connected set of all Level 2 routers from all areas, which are in turn connected to local areas. Routers that are solely used as L1 routers must route all traffic that is destined outside of the area to a L1/L2 or a L2 router. The contiguous set of L2 routers make up the OSPF backbone, similar to area 0 in OSPF networks.

Incorrect Answers:

A, B: Area 0 or Area 1 has no special significance in IS-IS. In IS-IS, areas are broken down into levels and are not numbered like OSPF areas.

C: Level 2, not Level 1.

E: This is not a requirement.

---

**QUESTION 126:**

Which packet type is used to acknowledge LSPs on point-to-point links and to request missing pieces of information in the IS-IS link-state database?

- A. complete SNP (CSNP)
- B. partial SNP (PSNP)
- C. hello
- D. database query
- E. database description packet (DDP)

Answer: B

Explanation:

Partial SNP (PSNP) - Includes a subset of LSPs, used to request individual LSPs and to acknowledge receipt of these LSPs: Level1,Level2.

Reference: CCNP Self-Study CCNP BSCI Exam Certification Guide p.361

---

**QUESTION 127:**

In IS-IS routing, which level is used to route between different areas within the same domain?

- A. Level 0
- B. Level 1
- C. Level 2
- D. Level 3

Answer: C

Explanation:

Small IS-IS networks are built as a single area that includes all the routers in the network. As the network grows larger, it is usually reorganized into a backbone area made up of the connected set of all Level2 routers from all areas, which is in turn connected to local areas. Within a local area, routers know how to reach all system IDs. Between areas, routers know how to reach the backbone, and the backbone routers know how to reach other areas.

Routers establish Level1 adjacencies to perform routing within a local area (intra-area routing). Routers establish Level2 adjacencies to perform routing between different areas (inter-area routing).

**QUESTION 128:**

In your network you have two Level-1/ Level-2 IS-IS routers named CK1 and CK2 . The two routers are connected together via a WAN link, and they both belong to the same area. What kind of an adjacency can you expect them to establish?

- A. Level-1 only
- B. Level-2 only
- C. Both Level-1 and Level-2
- D. None unless statically configured
- E. WAN links do not support IS-IS adjacencies

Answer: C

Explanation:

IS-IS routers that are configured as L1/L2, which is the default, will run two separate SPF algorithms, one for each level. Neighbors on point-to-point networks always become adjacent unless they do not see themselves in their neighbors' hello PDU and match on certain parameters. On broadcast networks and nonbroadcast multiaccess (NBMA) networks, the DIS (Designated Intermediate System) will become adjacent with its neighbors.

Two routers will become neighbors if the following parameters are agreed upon:

1. Level 1-The two routers sharing a common network segment must have their interfaces configured to be in the same area if they are to have a Level 1 adjacency.
2. Level 2-The two routers sharing a common network segment must be configured as Level 2 if they are in different areas and want to become neighbors.

Within an Area Level 1 routing is used. Routing between areas is referred to as Level 2 routing.

A Level 1/Level 2 router may have neighbors in any area. It has two link-state databases: a Level 1 link-state database for intra-area routing and a Level 2 link-state database for inter-area routing. A Level 1/Level 2 router runs two SPFs and may require more memory and processing as a result.

---

**QUESTION 129:**

You have a network segment with an IS-IS level 1-2 router, and you're thinking of upgrading your network. What is the maximum number of Level-2 routing processes that this router can belong to?

- A. 1
- B. 2
- C. 30
- D. 40
- E. limited only by the router's resources
- F. Depends on the total number of IDB blocks.

Answer: A

Explanation:

The Cisco IOS software can handle simultaneous operation of up to 30 dynamic IP routing processes. The combination of routing processes on a router or access server consists of the following protocols (with the limits noted):

1. Up to 30 IGRP routing processes
2. Up to 30 OSPF routing processes
3. One RIP routing process
4. One IS-IS process
5. One BGP routing process

Up to 30 EGP routing processes

---

**QUESTION 130:**

Which IOS command would you enter to if you wanted to see the IS-IS level-2 routing database table?

- A. show isis route
- B. show clns route
- C. show isis database
- D. show clns neighbors

Answer: C

Explanation:

The show isis database (detail) command displays the contents of the IS-IS database. By default, this command will display all of the L1 and L2 database information.

---

**QUESTION 131:**

Which of the following types of routers can a Level-2 IS establish adjacencies with? (Select three)

- A. Any Level-1 IS in any area
- B. Any Level-2 IS in any area
- C. Any Level-1 IS in the same area
- D. Any Level-1/Level-2 IS in any area
- E. Any Level-1/Level-2 IS in the same area

Answer: B, D, E

Explanation:

The various adjacency types that an IS-IS router can be are shown below:

|                     |  |
|---------------------|--|
| <b>level-1</b>      | A Level 1 adjacency may be established if there is at least one area address in common between this system and its neighbors. Level 2 adjacencies will never be established over this interface.                                 |
| <b>level-1-2</b>    | A Level 1 and Level 2 adjacency is established if the neighbor is also configured as level-1-2 and there is at least one area in common. If there is no area in common, a Level 2 adjacency is established. This is the default. |
| <b>level-2-only</b> | Level 2 adjacencies are established if the other routers are L2 or L1L2 routers and their interfaces are configured for L1L2 or L2. Level 1 adjacencies will never be established over this interface.                           |

Reference:<http://www.cisco.com/univercd/cc/td/doc/product/software/ios113ed/cs/csprtn1/csisis.htm>

---

**QUESTION 132:**

Which IOS command displays the Level-1 routing table in Integrated IS-IS?

- A. show isis route
- B. show clns route
- C. show isis database
- D. show clns neighbors

Answer: A

Explanation:

The "show isisroute" command displays the IS-IS Level 1 forwarding table for IS-IS-learned routes.

---

**QUESTION 133:**

One of the characteristics of the routing protocol IS-IS is its ability to use level-1 routing. Where in an IS-IS network does this happen?

- A. Between domains.
- B. Between areas.
- C. Between intermediate systems in the same area.
- D. Between end systems and intermediate systems in the same area.

Answer: C

Explanation:

A two-level hierarchy is used to support large routing domains. A large domain may be administratively divided into areas. Each system resides in exactly one area.1 Routing within an area is referred to as Level 1 routing. Routing between areas is referred to as Level 2 routing. A Level 2 Intermediate System (IS) keeps track of the paths to destination areas. A Level 1 IS keeps track of the routing within its own area. For a packet destined for another area, a Level 1 IS sends the packet to the nearest Level 2 IS in its own area, regardless of what the destination area is. Then the packet travels via Level 2 routing to the destination area, where it may travel via Level 1 routing to the destination. It should be noted that selecting an exit from an area based on Level 1 routing to the closest Level 2 IS might result in suboptimal routing.2

Reference:

[http://www.cisco.com/en/US/tech/CK365/CK381/technologies\\_white\\_paper09186a00800a3e6f.shtml](http://www.cisco.com/en/US/tech/CK365/CK381/technologies_white_paper09186a00800a3e6f.shtml)

---

#### **QUESTION 134:**

Your newly appointed Certkiller trainee wants to know why Level-3 area routing is not a supported feature of integrated IS-IS on Cisco routers.

What will your reply be?

- A. The system ID on a Cisco router is limited to 6 bytes.
- B. The NET on a Cisco router is restricted to a maximum of 8 bytes.
- C. The lack of a Domain portion of the NSAP only accommodates for 2 levels of routing hierarchy.
- D. Cisco routers cannot route CLNS data that use the ISO/IEC 10589 standard of NSAP addressing.
- E. Since the NSAP service identifier (N-SEL) must always be set to 00, no other service types are available.

Answer: C

Explanation:

Integrated IS-IS is a version of the OSI IS-IS routing protocol that uses a single routing algorithm to support more network layer protocols than just CLNP. Integrated IS-IS sometimes is called Dual IS-IS, named after a version designed for IP and CLNP networks. Only one IS-IS process is allowed whether you run it in integrated mode, ISO CLNS only or IP only.

---

#### **QUESTION 135:**

Which of the following three statements correctly describe the characteristics of IS-IS? (Select three)

- A. L1 routers have no knowledge about routes outside their area.
- B. L1/L2 routers maintain a separate Level 1 link-state database and a Level 2 link-state database; they do not advertise L2 routes to L1 routers.
- C. To route packets to another area, L1 routers must forward the packets to the L2 router of the destination area.
- D. To route packets to another area, L1 routers must forward the packets to an L1/L2 router within their area.
- E. L2 routers form adjacencies with L1 and L1/L2 neighbors.

Answer: A, B, D

Explanation:

Intermediate System-to-Intermediate System (IS-IS) is an OSI link-state hierarchical routing protocol that floods the network with link-state information to build a complete, consistent picture of network topology. To simplify router design and operation, IS-IS distinguishes between Level 1 and Level 2 ISs. Level 1 ISs communicate with other Level 1 ISs in the same area. Level 2 ISs route between Level 1 areas and form an intradomain routing backbone. Hierarchical routing simplifies backbone design because Level 1 ISs need to know only how to get to the nearest Level 2 IS. The backbone routing protocol also can change without impacting the intra-area routing protocol. Reference:[http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito\\_doc/osi\\_rout.htm](http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/osi_rout.htm)

---

**QUESTION 136:**

On router CK1, IS-IS is configured on all Ethernet interfaces. In IS-IS, what network type identifies a LAN interface?

- A. broadcast
- B. point-to-point
- C. pseudo-node
- D. non-broadcast
- E. point-to-multipoint

Answer: A

Explanation:

The types of networks that IS-IS defines include Point-to-point networks and Broadcast networks. Broadcast networks are used on all LAN interfaces.

Reference: [http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys\\_wp.htm](http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys_wp.htm)

Incorrect Answers:

- B: Point-to-point networks, such as serial lines, connect a single pair of routers.
- C: A Designated Intermediate System (DIS) creates a pseudonode (a virtual node), and all the routers on a LAN, including the DIS, form an adjacency with the pseudonode instead of forming  $n*(n-1)$  order adjacencies with each other in a full mesh. DISs are not used by default however.

- D: Non-broadcast is not used by IS-IS.
- E: Point-to-multipoint is not used by IS-IS.

---

**QUESTION 137:**

The Certkiller network is utilizing IS-IS for the routing protocol. By default, what is the metric used by the Certkiller routers on IS-IS interfaces?

- A. Delay
- B. Error
- C. Default
- D. Expense
- E. Bandwidth
- F. Hops
- G. Cost

Answer: G

Explanation: The default metric used on Cisco IOS for IS-IS is cost. For each IS-IS interface, the range is from 0 to 63 with the default value being 10.

---

**QUESTION 138:**

In the IS-IS routing protocol, what metric is supported exclusively by the Cisco IOS?

- A. Cost
- B. Delay
- C. Expense
- D. Error
- E. None of the above

Answer: A

Explanation:

The original IS-IS specification defines four different types of metrics. Cost, being the default metric, is supported by all routers. Delay, expense, and error are optional metrics. The delay metric measures transit delay, the expense metric measures the monetary cost of link utilization, and the error metric measures the residual error probability associated with a link.

The Cisco implementation uses cost only. If the optional metrics were implemented, there would be a link-state database for each metric and SPF would be run for each link-state database.

IS-IS uses a single required default metric with a maximum path value of 1024. The metric is arbitrary and typically is assigned by a network administrator. Any single link can have a maximum value of 64, and path links are calculated by summing link values.


Maximum metric values were set at these levels to provide the granularity to support various link types while at the same time ensuring that the shortest-path algorithm used for route computation will be reasonably efficient. IS-IS also defines three optional metrics (costs): delay, expense, and error.

Reference:

[http://www.cisco.com/en/US/tech/CK3\\_65/CK3\\_81/technologies\\_white\\_paper09186a00800a3e6f.shtml](http://www.cisco.com/en/US/tech/CK3_65/CK3_81/technologies_white_paper09186a00800a3e6f.shtml)

**QUESTION 139:**

The Certkiller network is displayed in the diagram below:



In the network segment of the above exhibit, which router is going to become the level-2 DIS (Designated Intermediate System)?

- A. Certkiller 1
- B. Certkiller 2
- C. Certkiller 3
- D. Certkiller 4
- E. None

Answer: B

Explanation:

Since the Level 2 priority of Certkiller 2 is higher, it will become the DIS. The DIS election process is described below:

On a LAN, one of the routers elects itself the DIS, based on interface priority (the default

is 64). If all interface priorities are the same, the router with the highest subnetwork point of attachment (SNPA) is selected. The SNPA is the MAC address on a LAN, and the local data link connection identifier (DLCI) on a Frame Relay network. If the SNPA is a DLCI and is the same at both sides of a link, the router with the higher system ID becomes the DIS. Every IS-IS router interface is assigned both a L1 priority and a L2 priority in the range of 0 to 127.

The DIS election is preemptive (unlike OSPF). If a new router boots on the LAN with a higher interface priority, the new router becomes the DIS. It purges the old pseudonode LSP and floods a new set of LSPs.

Reference:

[http://www.cisco.com/en/US/tech/CK365/technologies\\_tech\\_note09186a0080094b42.shtml#subfirstthree](http://www.cisco.com/en/US/tech/CK365/technologies_tech_note09186a0080094b42.shtml#subfirstthree)

---

### **QUESTION 140:**

On a point-to-point IS-IS network connection, the PSNP (Partial Sequence Number Packet) is responsible for which of the following?

- A. Acknowledging LSPs.
- B. Replacing IIH packets.
- C. Establishing adjacencies.
- D. Sending link-state changes.

Answer: A

Explanation:

Partial sequence number PDUs (PSNPs) are used to request an LSP (or LSPs) and acknowledge receipt of an LSP (or LSPs).

Reference: Cisco, Introduction to Intermediate, System-to-Intermediate System Protocol

[http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys\\_wp.htm](http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys_wp.htm)

---

### **QUESTION 141:**

Which of the following metrics can be used by IS-IS when determining the best path to a destination? (Choose four)

- A. default
- B. load
- C. delay
- D. expense
- E. reliability
- F. error
- G. uptime

Answer: A, C, D, F

Explanation:

### IS-IS Metrics

IS-IS uses a single required default metric with a maximum path value of 1024. The metric is arbitrary and typically is assigned by a network administrator. Any single link can have a maximum value of 64, and path links are calculated by summing link values. Maximum metric values were set at these levels to provide the granularity to support various link types while at the same time ensuring that the shortest-path algorithm used for route computation will be reasonably efficient. IS-IS also defines three optional metrics (costs): delay, expense, and error. The delay cost metric reflects the amount of delay on the link. The expense cost metric reflects the communications cost associated with using the link.

The error cost metric reflects the error rate of the link. IS-IS maintains a mapping of these four metrics to the quality of service (QoS) option in the CLNP packet header. IS-IS uses these mappings to compute routes through the network.

---

### QUESTION 142:

Which packet type is used to acknowledge LSPs on point-to-point links and to request missing pieces of information in the IS-IS link state database?

- A. complete SNP (CSNP)
- B. partial SNP (PSNP)
- C. hello
- D. database query
- E. database description packet (DDP)

Answer: B

### Explanation:

The different IS-IS packet types are described below:

CSNP-Complete Sequence Number PDUs are used to tell other routers about all the LSPs in one router database; similar to an OSPF database descriptor packet

PSNP-Partial Sequence Number PDUs are used to request an LSP and acknowledge receipt of an LSP


PDUs (PSNP's) every 3 seconds and complete sequence number PDUs (CSNPs) every 10 seconds that describe the full database of the DIS. Another router can ask the DIS for a missing LSP using a partial sequence number packet (PSNP), or in turn give the DIS a new LSP.

---

### QUESTION 143:

The Certkiller network is shown below, along with the relevant information on router CK1 :

## 642-801



```
CK1 #show clns interface ethernet 0
Ethernet0 is up, line protocol is up
Checksums enabled, MTU 1497, Encapsulation
SAP
Routing Protocol: ISIS
Circuit Type: level-1-2
Interface number 0x0, local circuit ID 0x1
Level-1 Metric: 10, Priority: 64, Circuit ID: CK2 .02
Number of active level-1 adjacencies: 1
Level-2 Metric: 10, Priority: 64, Circuit ID: CK2 .01
Number of active level-2 adjacencies: 1
Next ISIS LAN Level-1 Hello in 5 seconds
Next ISIS LAN Level-2 Hello in 2 seconds
```

Based on the information above, which statement is correct about IS-IS?

- A. Router CK1 is the DIS for both L1 and L2
- B. Router CK1 is the backup DIS for both L1 and L2.
- C. Router CK2 is the DIS for both L1 and L2.
- D. Router CK1 is the DIS for L1 and a different router ( CK2 ) is the DIS for L2.
- E. Router CK1 is the DIS for L2 and a different router ( CK2 ) is the DIS for L1.

Answer: C


Explanation:

It is the DIS that generates the pseudonode Link State Packet (LSP) and is denoted with a non-zero LSP-ID - R2.01

---

### **QUESTION 144:**

The Certkiller IS-IS network is displayed below:



Given the above diagram, which statement correctly matches the letter to the type of IS-IS hello PDU?

- A. A=ISH; B=IIH; C:ESH
- B. A=IIH; B=ESH; C=ISH
- C. A=ISH; B=ESH; C:IIH
- D. A=ESH; B=ISH; C=IIH
- E. A=IIH; B=ISH; C=ESH

Answer: C

Explanation:

OSI CLNS is a network layer service similar to bare IP service. A CLNS entity communicates over Connectionless Network Protocol (CLNP) with its peer CLNS entity. In the OSI architecture there are "systems": Routers are ISs, and hosts are End Systems (ESs).

ESs themselves have no routing information; they discover ISs (routers) by listening to Intermediate System Hellos (ISHs) and sending traffic to any random router. ESs send End System Hellos (ESHs); they do not choose a designated router to handle all traffic, and optimal routing is accomplished via redirects.

ISs discover ESs by listening to ESHs, and ISs send ISHs to ESs.

Intermediate System-to-Intermediate System Hello (IIH)-Used by routers to detect neighbors and form adjacencies. In addition to the IIH, which is an IS-IS protocol data unit (PDU), there is an ISH and an ESH, which are End System-to-Intermediate System (ES-IS) PDUs.

#### QUESTION 145:

DRAG DROP

As a Certkiller .com network technician you are required to drag the appropriate IS-IS routing item to its correct definition.

**642-801**

As a Certkiller.com network technician you are required to drag the appropriate IS-IS routing them to its correct definition.

|   |                  |
|---|------------------|
| Between ES's and IS's on the same subnet. OSI routing begins at this level. | Drop answer here |
| Routing between separate domains. Similar to BGP.                           | Drop answer here |
| Between IS's on the same area. Also called area routing.                    | Drop answer here |
| Called inter-area routing.  | Drop answer here |

Select from these

|                   |                 |                 |
|-------------------|-----------------|-----------------|
| Level 1/2 routing | Level 2 routing | Level 3 routing |
| Level 0 routing   | Level 1 routing |                 |

Answer:

As a Certkiller.com network technician you are required to drag the appropriate IS-IS routing them to its correct definition.

|   |                 |
|---|-----------------|
| Between ES's and IS's on the same subnet. OSI routing begins at this level. | Level 0 routing |
| Routing between separate domains. Similar to GP.                            | Level 3 routing |
| Between IS's on the same area. Also called area routing.                    | Level 1 routing |
| Called inter-area routing.  | Level 2 routing |

Select from these

|                   |
|-------------------|
| Level 1/2 routing |
|-------------------|

---

**QUESTION 146:**

What does IS-IS use to establish and maintain neighbor relationships between ISs?

- A. IHH
- B. LSP
- C. CLNS
- D. CLNP
- E. ISH

Answer: A

Explanation:

IHH (IS-IS Hello) - These packets create and maintain neighbor relationships and adjacencies. There are three types of Integrated IS-IS Hello packet. The type of packet is

defined in the fixed header under the Type field and allows the packet to be handed off to the appropriate process.

Reference: CCNP BSCI Exam Certification Guide p.370

---

**QUESTION 147:**

In an IS-IS environment, what happens when the designated IS router crashes?

- A. The elected backup designated router takes the place of the DIS indefinitely without the necessity of a new election.
- B. The elected backup designated router takes the place of the DIS indefinitely without the necessity of a new election, until the original DIS comes back online.
- C. A new election process occurs immediately, establishing a new DIS until a router with a higher priority or MAC address establishes an adjacency.
- D. A new election process occurs immediately, establishing a new DIS that will remain in place indefinitely.

Answer: C

Explanation:

There is no backup designated router in IS-IS. Therefore, if the DIS meets an untimely death, a new DIS would be elected, based on priority or highest MAC address. If another router comes online with a higher priority, it will dislodge the existing DIS and rule in its place. This behavior is different from that of OSPF. Once a new DIS is elected, the link-state databases are purged and new LSPs are flooded.

Reference: CCNP Self-Study CCNP BSCI Exam certification guide p.806

---

**QUESTION 148:**

IS-IS Level 1 routers use LSPs for which purpose?

- A. to build topology database for the local area only
- B. to build topology database for the local and remote areas
- C. to build topology database for remote areas
- D. to build topology database for areas outside the AS
- E. to establish with L1 IS's
- F. to establish adjacency with ES's

Answer: A

Explanation:

IS-IS Level 1 routers shares LSAs with L1 routers in the same area.

---

**QUESTION 149:**

In IS-IS routing, when would the use of Level 1 IS routing be required?

- A. When routing between end systems and intermediate systems.
- B. When routing between Level 1 areas in different domains.
- C. When routing between Level 1 areas in the same domain.
- D. When routing between intermediate systems in the same area.
- E. When routing between intermediate systems in different areas.

Answer: D

Explanation:

Layer 1 links connect the routers in the same area, so layer 1 routing takes place between intermediate systems in the same area. Layer 2 routing takes place between intermediate systems in different areas.

---

**QUESTION 150:**

Given the NSAP 39.0100.0102.0001.0c00.1211.00, which portion is interpreted by IS-IS as the area?

- A. 39
- B. 39.0100
- C. 39.0100.0102
- D. 0001
- E. 0001.0c00
- F. 0001.0c00.1211

Answer: C

Explanation:

A NET is a NSAP where the last byte is always zero. On a router running IS-IS, a NET can be 8 to 20 B in length. The last byte is always the n-selector and must be zero. The n-selector indicates which transport entity the packet is sent to. An n-selector of zero indicates no transport entity, and means that the packet is for the routing software of the system. The six bytes directly preceding the n-selector are the system ID. The system ID length is a fixed size and cannot be changed. The system ID must be unique throughout each area (Level 1) and throughout the backbone (Level 2). All bytes preceding the system ID are the area ID. The area ID field is the portion of the NSAP preceding the system ID.

The following example shows a router with area ID 47.0004.004d.0001 and system ID 0001.0c11.1110:

47.0004.004d.0001.0001.0c11.1110.00

In our example, the six bytes preceding the N selector is 0001.0c00.1211 which is the system ID. The portion preceding this, 39.0100.0102, is the area ID.

---

**QUESTION 151:**

**642-801**

Router CK1 is configured as a Level-1 IS-IS router. Which of the following routers can router CK1 establish an adjacency with? (Select two)

- A. Any Level-1 IS in any area.
- B. Any Level-2 IS in any area.
- C. Any Level-1 IS in the same area.
- D. Any Level-2 IS in the same area.
- E. Any Level-1/Level-2 IS in the same area.

Answer: C, E

Explanation:

A Level-1 IS router can establish adjacencies with other routers Level-1 and Level-1/Level-2 IS routers within the same area.

Incorrect Answers:

- A: Level-1 adjacencies can only be established within the same area.
- B, D: Level-2 adjacencies require Level-2 IS routers.

---

**QUESTION 152:**

Router CK1 is an IS-IS Level-1/Level2 router. Which of the following type of router will CK1 NOT be able to establish an adjacency with?

- A. Any Level-1 IS in any area.
- B. Any Level-2 IS in any area.
- C. Any Level-1 IS in the same area.
- D. Any Level-1/Level-2 IS in any area.
- E. Any Level-1/Level-2 IS in the same area

Answer: A

Explanation:

Level-1 routers can only establish adjacencies with other Level-1 routers within the same area. This is because L1 routers in different areas must connect through a L2 capable router, so two L1 routers must be in the same area in order to become adjacent.

Incorrect Answers:

- B, C, D, E: IS-IS has a two-level hierarchy. Contiguous Level 2-capable routers form the backbone. Both Level 2 and Level 1 routers live in areas. Routers can be Level 1 (L1), Level 2 (L2), or both (L1/L2). Within Cisco IOS(r) Software, the default configuration is both Level 1 and Level 2 at the same time which allows an IS-IS network to run with minimal configuration in a plug-and-play fashion. L1/L2 routers act as 2 separate routers in terms of adjacency building, so it will become adjacent with other L2 routers, as well as other L1 routers that are in the same area.

---

**QUESTION 153:**

What is periodically sent by a DIS on a LAN to ensure that all adjacent neighbors' IS-IS link-state databases are synchronized?

- A. complete SNP (CSNP)
- B. partial SNP (PSNP)
- C. database query
- D. database description packet (DDP)
- E. link-state summary
- F. hello

Answer: A

Explanation:

All ISs maintain adjacencies with all other ISs on a broadcast network. DIS sends CSNP to all ISs. Periodic CSNPs ensure the databases are synchronized.

Reference: CCNP Self-Study CCNP BSCI Exam Certification Guide p.350

---

#### **QUESTION 154:**

Which of the following items must match in an OSPF hello packet in order for two routers to become OSPF neighbors? (Choose four)

- A. Neighbors
- B. Hello/dead intervals
- C. Area ID
- D. Address of DR/BDR
- E. Stub area flag
- F. Authentication password

Answer: B, C, E, F

Explanation:

The following hello parameters must match in order for two OSPF routers to become adjacent neighbors:

1. OSPF area number
  2. OSPF area type, such as stub or NSSA flag
  3. Subnet and subnet mask
  4. OSPF HELLO and Dead timer values
  5. Authentication values/passwords
- 

#### **QUESTION 155:**

Which two commands can be used to verify that RIPv2 authentication is working properly to ensure that RIP neighbors form adjacencies properly? (Choose two.)

- A. show interface

- B. show ip route
- C. show ip protocols
- D. debug ip rip
- E. debug ip rip authentication

Answer: B, D

Explanation:

Verifying Plain Text Authentication

RIP version 2 routers can use either plain text or MD5 authentication.

By configuring the routers using RIPv2 authentication, all routing update exchanges will be authenticated before being accepted. This can be verified by observing the output obtained from the debug ip rip and show ip route commands.

```
RB#debug ip rip
```

```
RIP protocol debugging is on
```

```
*Mar 3 02:11:39.207: RIP: received packet with text authentication 234
```

```
*Mar 3 02:11:39.211: RIP: received v2 update from 141.108.0.10 on Serial0
```

```
*Mar 3 02:11:39.211: RIP: 70.0.0.0/8 via 0.0.0.0 in 1 hops
```

```
RB#show ip route
```

```
R 70.0.0.0/8 [120/1] via 141.108.0.10, 00:00:25, Serial0
```

```
80.0.0.0/24 is subnetted, 1 subnets
```

```
C 80.80.80.0 is directly connected, Loopback0
```

```
141.108.0.0/30 is subnetted, 1 subnets
```

```
C 141.108.0.8 is directly connected, Serial0
```

Using plain text authentication improves the network design by preventing the addition of routing updates originated by routers not meant to take part in the local routing exchange process. However, this type of authentication is not secure. The password (234 in this example) is exchanged in plain text. It can be captured easily and thus exploited.

As mentioned before, MD5 authentication must be preferred over plain text authentication when security is an issue.

Verifying MD5 Authentication

By configuring the RA and RB routers as shown above, all routing update exchanges will be authenticated before being accepted. This can be verified by observing the output obtained from the debug ip rip and show ip route commands.

```
RB#debug ip rip
```

```
RIP protocol debugging is on
```

```
*Mar 3 20:48:37.046: RIP: received packet with MD5 authentication
```

```
*Mar 3 20:48:37.046: RIP: received v2 update from 141.108.0.10 on Serial0
```

```
*Mar 3 20:48:37.050: 70.0.0.0/8 via 0.0.0.0 in 1 hops
```

```
RB#show ip route
```

```
R 70.0.0.0/8 [120/1] via 141.108.0.10, 00:00:03, Serial0
```

```
80.0.0.0/24 is subnetted, 1 subnets
```

```
C 80.80.80.0 is directly connected, Loopback0
```

```
141.108.0.0/30 is subnetted, 1 subnets
```

```
C 141.108.0.8 is directly connected, Serial0
```

MD5 authentication uses the one-way, MD5 hash algorithm, acknowledged to be a strong

hashing algorithm. In this mode of authentication, the routing update does not carry the password for the purpose of authentication. Rather, a 128-bit message, generated by running the MD5 algorithm on the password, and the message are sent along for authentication. Thus, it is recommended to use MD5 authentication over plain text authentication since it is more secure.

Reference:

[http://www.cisco.com/en/US/tech/CK365/technologies\\_configuration\\_example09186a0080093f1c.shtml#veri](http://www.cisco.com/en/US/tech/CK365/technologies_configuration_example09186a0080093f1c.shtml#veri)

---

### **QUESTION 156:**

Which three of the statements below correctly describe the characteristics of Autonomous Systems in routed networks? (Select three)

- A. Within an AS, all routers must run either BGP or IBGP.
- B. An AS uses exterior gateway protocols (EGPs) to exchange information with other autonomous systems.
- C. An AS is a group of routers under the same technical administration.
- D. Within an AS, routes learned through BGP can be redistributed using interior gateway protocols.
- E. Within an AS, routes learned through an interior protocol cannot be redistributed using BGP to other autonomous systems.

Answer: B, C, D

Explanation:

1. Exterior Gateway Protocol (EGP) routing protocol used to connect between autonomous systems.
2. The use of the term autonomous system in connection with BGP stresses the fact that the administration of an autonomous system appears to other autonomous systems to have a single coherent interior routing plan, and presents a consistent picture of those networks that are reachable through it.
3. BGP is used between autonomous systems

Incorrect Answers:

A: Within an AS, all routers need not run either interior BGP or exterior BGP. While all IBGP peers need to be fully meshed, all routers within the network do not need to run BGP.


E: One of the major reasons for running BGP is so that the interior public IP networks within an AS can be distributed to other systems in the Internet via BGP.

Reference: Building Scalable Cisco Networks (Cisco Press) page 313

---

### **QUESTION 157:**

The Certkiller network is displayed in the following diagram:



On the assumption that there is no IGP running in AS 64700 and synchronization is OFF, which router in AS 64700 is going to advertise the route to 183.32.0.0?

- A. CK2 only
- B. CK5 only
- C. CK2 and CK5 only
- D. CK2 , CK3 , and CK4 only
- E. CK2 , CK3 , CK4 , and CK5

Answer: A

Explanation: We need to run IGP on AS 64700 to control the exit points from the AS. When router CK2 receives the EIGRP route, it will advertise it to all IBGP routers. However, IBGP routers do not advertise the IBGP learned route unless it is also in their IGP table, or synchronization is turned off.

Reference: Using the Border Gateway Protocol for Interdomain Routing  
<http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/icsbgp4.htm>

---

**QUESTION 158:**

Which of the following statements about IBGP routers are true? (Select one.)

- A. They must be fully meshed.
- B. They can be in a different AS.
- C. They must be directly connected.
- D. They do not need to be directly connected.

Answer: D

Explanation:

Since BGP runs over TCP, the IBGP speakers need to only be able to reach the other IBGP speakers via a TCP connection across the network and do not need to be directly connected. Even if two IBGP neighbors are separated by other routers, they will still become neighbors as long as they are reachable via TCP.

Incorrect Answers:

A: Within an AS, all IBGP speakers must be configured as a full mesh. The only exceptions to this rule occur when router reflection or confederations are used.

B: IBGP routers need to be configured with the same AS number. The only exception to this rule is when BGP confederations are in use.

C: The routers only need to be reachable via a TCP connection.

---

**QUESTION 159:**

Given the following choices, what kind of BGP router advertises routes to other IBGP neighbors?

- A. Client
- B. EBGP peer
- C. Route reflector
- D. cluster of clients
- E. None of the above

Answer: C

Explanation:

Ordinarily, with no route reflector, a full mesh of IBGP peers is required. Route reflectors modify the BGP split horizon rule by allowing the router configured as the route reflector to propagate routes learned by IBGP to other IBGP peers. Route reflectors reduce the number of BGP neighbor relationships in an AS.

---

**QUESTION 160:**

BGP has been configured on some of the Certkiller routers. The BGP routing process relies on two different types of tables. What are they? (Select two)

- A. An IP routing table.
- B. A BGP topology table.
- C. A BGP attribute table.
- D. A table that contains BGP information received from and sent to other routers.
- E. A combined table that contains both IP routes and BGP information received from and sent to other routers.

Answer: B, C

Explanation:

BGP uses one table for topology and another for attribute. The attribute table assigns values to various attributes (weight, local preference, multi-exit discriminator, origin, AS path, next hop, & community) and the topology table matches the values of these attributes to the various routes it can select.

---

**QUESTION 161:**

Attributes are a means of sending additional route information over BGP. Which of

the following statements are true regarding BGP attributes? (Select three)

- A. MED is an optional attribute.
- B. Origin is an optional attribute.
- C. Next-hop is an optional attribute.
- D. Local Preference is a discretionary attribute.
- E. AS-Path is a well-known mandatory attribute.
- F. Community is a well-known mandatory attribute.

Answer: A, D, E

Explanation:

A: The Multiple Exit Discriminator (MED) attribute is optional.

D: The local preference attribute is used to prefer an exit point from the local autonomous system (AS). It is a well-known discretionary attribute.

E: When a route advertisement passes through an autonomous system, the AS number is added to an ordered list of AS numbers that the route advertisement has traversed. The AS-Path attribute is mandatory.

Note:

BGP metrics are called path attributes. Optional attributes are recognized by some implementations, but are expected not to be recognized by everyone. A well-known mandatory attributes must be present in all update messages.

The attributes defined by BGP include:

Well-known mandatory attributes:

1. AS-path
2. Next-hop
3. Origin

Well-known discretionary attributes:

1. Local preference
2. Atomic aggregate

Optional transitive attributes:

1. Aggregator
2. Communities

Optional non-transitive attribute:

1. Multi-Exit-Discriminator (MED)

Incorrect Answers:

B: The origin attribute indicates how BGP learned about a particular route. It can have three values: IGP, EGP or incomplete. It is a required attribute.

C: The EBGP next-hop attribute is the IP address that is used to reach the advertising router. It is a required attribute.

F: The community attribute provides a way of grouping destinations, called communities, to which routing decisions can be applied. It is not mandatory.

Reference: Border Gateway Protocol

[http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito\\_doc/bgp.htm](http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm)

**QUESTION 162:**

The BGP community attribute tags route for the sake of consistency in route filtering. Which of the answer choices below correctly describe the community attribute?

- A. Optional and transitive.
- B. Optional and non-transitive.
- C. Well-known and mandatory.
- D. Well-known and discretionary.

Answer: A

Explanation:

The community attribute is an optional transitive attribute that can be in the range 0 to 4,294,967,200. Each network can be a member of more than one community.

The attributes defined by BGP include:

Well-known mandatory attributes:

1. AS-path
2. Next-hop
3. Origin

Well-known discretionary attributes:

1. Local preference
2. Atomic aggregate

Optional transitive attributes:

1. Aggregator
2. Communities

Optional non-transitive attribute:

1. Multi-Exit-Discriminator (MED)

---

**QUESTION 163:**

You are a guest lecturer at the Certkiller Academy teaching some CCNP hopefuls about BGP. One of the students, Bob asks you what a BGP community is. How would you respond?

- A. Communities are tagged by default in outgoing updates.
- B. Communities can only be used within one autonomous system.
- C. Communities are a means of tagging routes to ensure consistent filtering.
- D. Communities perform summarization of blocks of contiguous network prefixes.

Answer: C

Explanation:

A community is a group of destinations which share some common property. No tag is used by default. Communities are a means of tagging routes to ensure consistent filtering

or route-selection policy.

Incorrect Answers

A: By default, all destinations belong to the general Internet community and are not tagged.

B: Each autonomous system administrator may define which communities a destination belongs to. Community information is passed on between different autonomous systems.

D: No summarization is performed by communities. Communications can be aggregated, however.

Reference: RFC 1997, BGP Communities Attribute

---

### **QUESTION 164:**

The Certkiller IBGP routers are configured in a fully meshed fashion. Which of the following are reasons why IBGP routers should peer with each other within a single AS? (Select two)

- A. IBGP routes are not propagated to other EBGP peers.
- B. IBGP routes that a router originates are propagated to other IBGP peers.
- C. IBGP routes are propagated to other IBGP speakers in the AS that are not peers.
- D. IBGP routes that are learned from an IBGP neighbor are propagated to only EBGP peers.

Answer: A, B

Explanation:

IBGP routes are propagated to all IBGP peers and only the IBGP peers. They are not exchanged with EBGP peers, which is why all IBGP routers must be configured in a full mesh.

Note: You can configure Border Gateway Protocol (BGP) either within an autonomous system or between different autonomous systems. When run within an autonomous system, it's called internal BGP (IBGP). When run between different autonomous systems, it's called external BGP (EBGP).

---

### **QUESTION 165:**

Routers CK1 and CK2 belong to the same BGP peer group. Which two of the statements below are true regarding this BGP peer group? (Select two)

- A. The peer group name is passed to other routers in the peer group.
- B. A peer group is a group of BGP neighbors with different update policies.
- C. The peer group name is only local to the router on which it is configured.
- D. A peer group allows options that affect outbound updates to be overridden.
- E. A peer group is a more efficient way to update BGP than configuration individual neighbors.

Answer: C, E

Explanation:

C: The peer group name is only local to the router it is configured on, it is not passed to any other router.

E: Neighbors with the same update policies can be grouped into peer groups to simplify configuration and make update calculation more efficient.

Incorrect Answers:

A: A BGP peer group does not necessarily have a name, it might have a number.

B: BGP peers use the same update policy.

D: A peer group does not have this option.

Reference: RFC 1771, A Border Gateway Protocol 4 (BGP-4)

---

### **QUESTION 166:**

Routers CK1 and CK2 are configured as BGP peers. Which of the following is true regarding this peering relationship? (Select two)

- A. Periodic keepalives are used to verify connectivity.
- B. Incremental keepalives are used to verify connectivity.
- C. It provides a reliable connection between two BGP routers.
- D. It provides a "best effort" connection between two BGP routers.

Answer: A, C

Explanation:

A: The default keepalive frequency is 60 seconds.

C: BGP peering provides a reliable connection between BGP routers. BGP utilizes the inherent reliability of TCP, since all BGP sessions are maintained over TCP port 179.

Incorrect Answers:

B: Keepalives messages are not incremental.

D: BGP peers form a reliable connection.

Reference: BGP Peer Groups

<http://www.cisco.com/warp/public/459/29.html>

---

### **QUESTION 167:**

Router CK1 is configured as a route reflector. What would happen if router CK1 would receive an update from a peer in a different autonomous system?

- A. It discards the update.
- B. It sends the update to all IBGP peers.
- C. It sends the update only to non clients.
- D. It sends the update only to route reflector clients.
- E. It sends the update to all routers in the autonomous system.

Answer: B

Explanation:

When any BGP router receives an update from an EBGP peer, it will forward this routing information to all IBGP peers. This is true for all EBGP routers, including route reflectors.

Reference:

[http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products\\_configuration\\_guide\\_chapter09186a00800c](http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products_configuration_guide_chapter09186a00800c)

---

**QUESTION 168:**

Which of the following statements is true about BGP peer communications?

- A. Communication between BGP peers runs over RIP.
- B. Communication between BGP peers runs over TCP.
- C. Communication between BGP peers runs over UDP.
- D. Communication between BGP peers runs over ICMP.

Answer: B

Explanation:

BGP communicate through a TCP connection. TCP port 179 is reserved for BGP sessions.

Reference: Border Gateway Protocol

[http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito\\_doc/bgp.htm](http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm)

---

**QUESTION 169:**

The Certkiller network is considered to be a single autonomous system (AS). Which three of the following statements are characteristics of an autonomous system? (Select three)

- A. It uses only Interior Gateway Protocols (IGPs).
- B. EGPs are used to connect different autonomous systems.
- C. It is a set of routers under a single technical administration.
- D. It uses EGPs to route packets to other autonomous systems and IGPs to route packets within the autonomous system.
- E. It uses IGPs to route packets to other autonomous systems and EGPs to route packets within the autonomous system.

Answer: B, C, D

Explanation:

A BGP autonomous system is a group of routers all administered as a single team working on mutual packet switching goals. EGP's (Exterior Gateway Protocol's) handle routing to outside autonomous systems which are on the exterior of that particular

autonomous system, and IGP's (Interior Gateway Protocol's) route within, on the interior of autonomous systems.

Incorrect Answers:

A: This is incorrect because an autonomous system can use an exterior gateway protocol.

E: This is incorrect because the roles of IGP's and EGP's are reversed in this answer choice.

---

**QUESTION 170:**

Which of the following statements is true about the BGP synchronization command? (Select two)

- A. Synchronization must be enabled when implementing a multi-homed BGP connection to multiple ISPs.
- B. If it is turned ON, a prefix learned from IBGP neighbor is valid only if a non-BGP (IGP) route exists for that prefix.
- C. Synchronization is necessary when peering with an EBGP neighbor.
- D. Synchronization improves BGP routing convergence.
- E. Synchronization can be turned off if all the transit routers in an Autonomous system are running a fully meshed IBGP.

Answer: B, E

Explanation:

If your autonomou

s system will be passing traffic through it from another autonomous system to a third autonomous system, it is very important that your autonomous system be consistent about the routes that it advertises. For example, if your BGP were to advertise a route before all routers in your network had learned about the route through your IGP, your autonomous system could receive traffic that some routers cannot yet route. To prevent this from happening, BGP must wait until the IGP has propagated routing information across your autonomous system. This causes BGP to be synchronized with the IGP.

Synchronization is enabled by default.

Only if all routers in the transit path in the AS are running BGP it is safe to turn synchronization off.


Reference: Building Scalable Cisco Networks (Cisco Press) page 33

[http://www.cisco.com/en/US/products/sw/iosswrel/ps1826/products\\_configuration\\_guide\\_chapter09186a008008](http://www.cisco.com/en/US/products/sw/iosswrel/ps1826/products_configuration_guide_chapter09186a008008)

---

**QUESTION 171:**

The Certkiller BGP network is shown in the diagram below:



Router CK1 receives BGP routing advertisements from router CK2 about the network 165.50.0.0. By default, what is the value of the next hop attribute?

- A. 40.1.1.1
- B. 40.1.1.2
- C. 40.1.1.3
- D. Router CK1 does not accept the advertisement from Router CK2 because Router CK1 is not peering with Router CK3 via BGP.
- E. Router CK2 does not advertise network 165.50.0.0 to Router CK1 because the network is not directly connected to Router CK2 .

Answer: A

Explanation:

The BGP next-hop attribute is a well-known mandatory attribute that indicates the next hop IP address that is to be used to reach a destination. For EBGP, the next hop is the IP address of the neighbor specified who sent the update, Router CK2 in this scenario.

Incorrect Answers:

B: Router CK2 learned this route through IBGP with the next-hop of 40.1.1.2. This value will not be used instead of 40.1.1.1 because the next hop information is not preserved across AS 64300 when being sent to a different AS.

C: This is the local interface on Router CK1 itself.


D: The advertisement is accepted.

E: The route is advertised.

Reference: Reference: Cisco Press's BSCI 2nd Edition pg 529

## QUESTION 172:

The Certkiller network is shown in the following exhibit:



1. all routers are using BGP
2. synchronization is OFF in AS 647000

Regarding the route to 198.12.1.0; which routers will this route propagate to?

- A. Router CK6
- B. Routers CK2 and CK6
- C. Routers CK2 , CK4 , and CK6
- D. Routers CK2 , CK3 , CK4 , and CK6
- E. Routers CK1 , CK2 , CK3 , CK4 , and CK6
- F. Routers CK1 , CK2 , CK3 , CK4 , CK5 , and CK6
- G. It will not be propagated to any other router

Answer: F


Explanation:

Router CK 6 will advertise the 192.12.1.0 route to Router CK2 using EBGP. This route would propagate to Router CK1 , CK3 , and CK4 through the fully meshed IGBP sessions. Since synchronization is turned off Router CK1 will advertise the internally learned route to Router CK5 using EBGP.

---

### QUESTION 173:

The Certkiller BGP network is shown below:



1. synchronization is OFF
2. NO IGP is running in AS 65500

Which router in AS 65500 is going to end up advertising a route to network 172.16.0.0?

- A. B only
- B. E only
- C. B and E only
- D. B, C and D only
- E. B, C, D, and E

Answer: C

Explanation:

Router B learns about the 172.16.0.0 network route via the EBGP session from router A. With synchronization turned off, it will pass this route to all of its fully meshed IBGP peer routers. Although all routers will receive the route, only routers B and E will advertise this route to other routers. Router B will advertise the route to all of the fully meshed IBGP peers, and router E will advertise it to AS 64500.

Reference: Using the Border Gateway Protocol for Interdomain Routing

<http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/icsbgp4.htm>

---

#### **QUESTION 174:**

The Certkiller network administrator is planning to use a private AS number. Which of the following are choices for a private AS number?

- A. 10080
- B. 48512
- C. 64128
- D. 64524

Answer: D

Explanation:

This autonomous system designator is a 16-bit number, with a range of 1 to 65535. RFC 1930 provides guidelines for the use of AS numbers. A range of AS number, 64512 through 65535, is reserved for private use, much like the private Internet Protocol (IP) addresses defined in RFC 1918.

Reference: Building Scalable Cisco Networks (Cisco Press) page 312

---

#### **QUESTION 175:**

The router that is being used for the Internet is low on memory. Because multiple ISP's are being used, BGP is needed for Internet routing. Which of the following types of routes should be received from the Internet providers if the goal is to consume as little memory as possible?

- A. Only default routes.
- B. Only partial BGP routes.
- C. Only internal BGP routes.

D. Only redistributed routes.

Answer: A

Explanation:

Today, the Internet consists of over 100,000 routes, so accepting the full routing table from an ISP can consume a great deal of router memory. By only accepting default routes from the ISPs ensures a minimal impact on the system. There is no risk of having your autonomous system (AS) becoming a transit AS. By receiving only default routes, the entire Internet will be reachable while only using a small amount of memory for the routing table.

Note:

One recommendation of Cisco, not listed here, is to use AS\_PATH filters for accepting only routes originated from an ISP and its directly connected Autonomous Systems (AS), instead of receiving the full BGP routing table from an ISP.

Reference: Achieve Optimal Routing and Reduce BGP Memory Consumption  
<http://www.cisco.com/warp/public/459/41.shtml>

---

**QUESTION 176:**

Routers CK1 , CK2 , and CK3 are all IBGP peers within the Certkiller network. Why should these peer sessions be fully meshed within the AS?

- A. Because BGP requires redundant TCP sessions between iBGP peers.
- B. Because a full mesh allows for optimal routing within the Transit AS.
- C. Because routes learned via iBGP are never propagated to other eBGP peers.
- D. Because routes learned via iBGP are never propagated to other iBGP peers.
- E. Because routes learned via eBGP are never propagated to other iBGP peers.

Answer: D

Explanation:

Since BGP does not relay routing traffic from one Interior BGP session to another (only from an Exterior BGP session to an IBGP session or another EBGP session), BGP speakers must be fully meshed. When a BGP routing update is received from a neighboring AS, it must be relayed directly to all other BGP speakers in the AS. Do not expect to relay BGP paths from one router, through another, to a third, all within the same AS.

Incorrect Answers:

A: BGP only requires a single TCP session to each peer.

B: The full mesh is needed so that all routers receive the routing information, but it does not provide for any level of route optimization.

C, E: Routes from EBGP peers are propagated to all IBGP peers, and routes learned via IBGP are sent to EBGP peers by default.

**QUESTION 177:**

Which BGP attribute is used by BGP to prevent routing loops?

- A. AS-path
- B. Next-hop
- C. MED
- D. Weight
- E. Local preference
- F. Origin

Answer: A

Explanation:

BGP uses path information to determine routing loops throughout the network. A simplified view of AS\_PATH is that it is the list of Autonomous Systems that a route goes through to reach its destination. Loops are detected and avoided by checking for your own AS number in AS\_PATH's received from neighboring Autonomous Systems. If your AS number appears in the AS path the second time, then a loop has occurred.

---

**QUESTION 178:**

When the BGP path selection process is being performed on a Cisco router, which BGP attribute is used first when determining the best path?

- A. Local preference
- B. MED
- C. Weight
- D. Origin
- E. Next-hop
- F. AS-path

Answer: C

Explanation:

How the Best Path Algorithm Works:

BGP assigns the first valid path as the current best path. It then compares the best path with the next path in list, until it reaches the end of the list of valid paths. The following is a list of rules used to determine the best path.

1. Prefer the path with the highest WEIGHT.

Note: WEIGHT is a Cisco-specific parameter, local to the router on which it's configured.

2. Prefer the path with the highest LOCAL\_PREF. Note the following:

3. 1. Path without LOCAL\_PREF is considered as having the value set with the bgp default local-preference command, or 100 by default.

2. Prefer the path that was locally originated via a network or aggregate BGP subcommand,

or through redistribution from an IGP. Local paths sourced by network or redistribute commands are preferred over local aggregates sourced by the aggregate-address command.

3. Prefer the path with the shortest AS\_PATH. Note the following:

4. 1. The AS\_CONFED\_SEQUENCE and AS\_CONFED\_SET are not included in the AS\_PATH length.

2. Prefer the path with the lowest origin type: IGP is lower than EGP, and EGP is lower than INCOMPLETE.

3. Prefer the path with the lowest multi-exit discriminator (MED).

4. Prefer external (eBGP) over internal (iBGP) paths. If bestpath is selected, go to Step 9 (multipath).

5. Prefer the path with the lowest IGP metric to the BGP next hop. Continue, even if bestpath is already selected.

6. Check if multiple paths need to be installed in the routing table for BGP Multipath. Continue, if bestpath is not selected yet.

7. When both paths are external, prefer the path that was received first (the oldest one). This step minimizes route-flap, since a newer path will not displace an older one, even if it would be the preferred route based on the next decision criteria (Steps 11, 12, and 13).

8. Prefer the route coming from the BGP router with the lowest router ID. The router ID is the highest IP address on the router, with preference given to loopback addresses. It can also be set manually using the bgp router-id command.

9. If the originator or router ID is the same for multiple paths, prefer the path with the minimum cluster list length. This will only be present in BGP route-reflector environments. It allows clients to peer with RRs or clients in other clusters. In this scenario, the client must be aware of the RR-specific BGP attribute.

10. Prefer the path coming from the lowest neighbor address. This is the IP address used in the BGP neighbor configuration, and corresponds to the remote peer used in the TCP connection with the local router.

---

### **QUESTION 179:**

Which one of the following statements about BGP is FALSE?

- A. BGP uses TCP port 179.
- B. BGP ensures reliability of updates by using the reliable transport services of TCP.
- C. The network command with the mask option never installs a prefix into the BGP table unless there is a matching prefix exists in the IP route table.
- D. A TCP connection is required before exchanging updates.
- E. BGP uses notification and the update messages to establish and maintain the BGP neighbor relationship.

Answer: E

Explanation:

BGP relies on TCP to provide for a reliable connection between peer routers. The statement in E is false, as the notification message is used to terminate the BGP session,

not to maintain or establish the peering session. BGP systems send notification messages when an error condition is detected. After the message is sent, the BGP session and the TCP connection between the BGP systems are closed. Notification messages consist of the BGP header plus the error code and subcode, and data that describes the error.

Incorrect Answers:

A: This statement is true.

B: BGP uses TCP port 179. TCP is a reliable, connection oriented protocol.

C: This statement is true. The route must actually exist in the routing table before it will be advertised using the network BGP configuration command.

D: These are both true statements regarding BGP. BGP relies on TCP to maintain sessions.

---

**QUESTION 180:**

Which BGP state is a router in when it has found the IP address in the neighbor statement and has created and sent out a BGP packet but the router has not received a response back?

- A. Update
- B. Connect
- C. Active
- D. Established

Answer: B

Explanation:

The following table describes the various states that a BGP router can be in:

BGPStateMachine

BGP goes through six states to establish an adjacency.

Idle - incoming connections are refused, and the system gets ready to start speaking BGP.

After this is done (by way of a Start event), move to Connect.

Connect - a connection is made to the peer. Send a BGP OPEN message, and go to OpenSent.

Active - a connection comes in from a peer. Send a BGP OPEN message, and go to OpenSent.

OpenSent - Wait for an OPEN message from the peer. When received, send a KEEPALIVE and go to OpenConfirm.

OpenConfirm - Wait for the KEEPALIVE from the peer, then move to Established.

Established - Bidirectional communication is established. Start sending UPDATE and KEEPALIVE messages as required

---

**QUESTION 181:**

Which BGP attribute will not be advertised in routing updates to its neighboring routers?

- A. weight
- B. local preference
- C. origin
- D. AS\_path
- E. next hop

Answer: A

Explanation:

Weight is proprietary to Cisco and is used in route selection. It is local to the router, and because it is not propagated to other routers, there is no problem with compatibility. When there are multiple paths, it selects a path to a destination with different next hops to the same destination. Note that the weight attribute has no code. Because it is a local attribute and is not propagated to other routers, no code is needed.

Reference: Cisco Press 642-801 Study Guide p.527

---

**QUESTION 182:**

DRAG DROP

Place the BGP attributes in the correct order used for determining a route.

|                  |     |
|------------------|-----|
| originate route  | 1st |
| AS Path          | 2nd |
| Weight           | 3rd |
| local preference | 4th |
| MED              | 5th |

Answer:

|                  |
|------------------|
| weight           |
| local preference |
| originate route  |
| AS Path          |
| MED              |

---

**QUESTION 183:**

Which attribute must exist in the BGP update packet?

- A. LOCAL\_PREF
- B. AGGREGATOR

- C. AS\_Path
- D. Weight

Answer: C

Explanation:

The AS PATH is a well known mandatory attribute. Mandatory means that it is required by all routers. These attributes are required and are therefore recognized by all BGP implementations.

---

**QUESTION 184:**

Which three conditions can cause BGP neighbor establishment to fail? (Select three)

- A. There is an access list blocking all TCP traffic between the two BGP neighbors.
- B. The EBGP neighbor is not directly connected, and the ebgp-multi-hop option is set to the default value.
- C. The IBGP neighbor is not directly connected.
- D. BGP synchronization is enabled in a Transit AS with fully-meshed IBGP neighbors.
- E. The BGP update interval is different between the two BGP neighbors.
- F. The BGP neighbor is referencing an incorrect AS number in its neighbor statement.

Answer: A, B, F

Explanation:

BGP uses TCP port 179 to establish and maintain neighbor relationships, so any access lists or firewalls must permit this port for BGP to function.

By default, EBGP multi-hop is not enabled, so the EBGP peer must be directly connected in order for the local BGP router to know how to reach the EBGP peer.

BGP uses neighbor statements that specify the AS number of the BGP peer. If it is the same number as the local BGP router process, then the BGP router knows that IBGP is used. If the AS number for the specified peer is different, then EBGP is used. Either way, the specified neighbor must be configured correctly, or the BGP peers will not become neighbors.

---

**QUESTION 185:**

Which statement is true about EBGP?

- A. An internal routing protocol can be used to reach an EBGP neighbor?
- B. The next hop does not change when BGP updates are exchanged between EBGP neighbors.
- C. A static route can be used to form an adjacency between neighbors.
- D. EBGP requires a full mesh.

Answer: C

Explanation:

Only choice C is correct. To reach a EBGP peer, a static route can be used with the EBGP multi-hop command to establish a neighbor adjacency. Unlike interior routing protocols, EBGP neighbors do not need to physically be connected to each other. All that is needed is a path to the peer IP address, and a valid TCP port 179 connection between the BGP routers.

---

**QUESTION 186:**

BGP contains two paths to a destination. Assuming both routes were originated locally and have an equal weight, what will be the next determining factor in choosing the best path?

- A. lowest MED
- B. highest local preference
- C. lowest neighbor IP address
- D. lowest origin code
- E. shortest AS-path

Answer: B

Explanation:

How the Best Path Algorithm Works:

BGP assigns the first valid path as the current best path. It then compares the best path with the next path in list, until it reaches the end of the list of valid paths. The following is a list of rules used to determine the best path.

1. Prefer the path with the highest WEIGHT.

Note: WEIGHT is a Cisco-specific parameter, local to the router on which it's configured.

2. Prefer the path with the highest LOCAL\_PREF.

3. Prefer the path that was locally originated via a network or aggregate BGP subcommand, or through redistribution from an IGP. Local paths sourced by network or redistribute commands are preferred over local aggregates sourced by the aggregate-address command.

4. Prefer the path with the shortest AS\_PATH. Note the following:

5. Prefer the path with the lowest origin type: IGP is lower than EGP, and EGP is lower than INCOMPLETE.

6. Prefer the path with the lowest multi-exit discriminator (MED). Note the following:

7. Prefer external (eBGP) over internal (iBGP) paths. If bestpath is selected, go to Step 9 (multipath).

8. Prefer the path with the lowest IGP metric to the BGP next hop. Continue, even if bestpath is already selected.

9. Check if multiple paths need to be installed in the routing table for BGP Multipath. Continue, if bestpath is not selected yet.

10. When both paths are external, prefer the path that was received first (the oldest one).

This step minimizes route-flap, since a newer path will not displace an older one, even if it would be the preferred route based on the next decision criteria (Steps 11, 12, and 13).

11. Prefer the route coming from the BGP router with the lowest router ID. The router ID is the highest IP address on the router, with preference given to loopback addresses. It can also be set manually using the `bgp router-id` command.

12. If the originator or router ID is the same for multiple paths, prefer the path with the minimum cluster list length. This will only be present in BGP route-reflector environments. It allows clients to peer with RRs or clients in other clusters. In this scenario, the client must be aware of the RR-specific BGP attribute.

13. Prefer the path coming from the lowest neighbor address. This is the IP address used in the BGP neighbor configuration, and corresponds to the remote peer used in the TCP connection with the local router.

Reference:

[http://www.cisco.com/en/US/tech/CK365/technologies\\_tech\\_note09186a0080094431.shtml](http://www.cisco.com/en/US/tech/CK365/technologies_tech_note09186a0080094431.shtml)

---

### **QUESTION 187:**

The Certkiller network administrator is considering the use of policy based routing. What would happen if BGP policy-based routing was implemented and the next hop router goes down, leaving no alternative path?

- A. If the next-hop router goes down and no alternative path is in place, policy routing will route to null 0.
- B. If the next-hop router goes down and no alternative path is in place, policy routing will default to another BGP path.
- C. If the next-hop router goes down and no alternative path is in place, policy routing will deny all traffic to that destination.
- D. If the next-hop router goes down and no alternative path is in place, policy routing will default to dynamic routing decisions.
- E. None of the above.

Answer: D

Explanation:

A backup path should be in place in case the defined next-hop router goes down. If there is no alternative defined, policy routing will default to dynamic routing decisions.

With any policy based routing implementation, when a packet is unable to match any of the conditions or if the next hop is not unreachable, the default routing behavior will occur.

Reference: CCNP Self-Study CCNP BSCI Exam Certification Guide by Clare Gough, page 499.

---

### **QUESTION 188:**

BGP Policy Based routing has been configured on router CK1 . Which of the following statements are true regarding policy-based routing on this router? (Select

two)

- A. Policy routing can be configured on the inbound interface.
- B. Policy routing can be used to alter the final destination of the packet.
- C. Policy routing can be used to alter the next hop in the path to the destination.
- D. Policy routing does not allow traffic to be directed based on the source address.

Answer: A, C

Explanation:

Rules of Policy-Based Routing:

BGP can implement any of the following rules associated with the hop-by-hop paradigm. This

paradigm is the capability to influence which router will be the next-hop router, potentially dictating it at every router and thus influencing the entire path of the traffic, hop by hop. The following rules associated with policy-based routing seem repetitive, but, in fact, each point raises a subtly different nuance:

? Traffic can be directed on either the source address or both the source and destination addresses, as well as inbound or outbound traffic on an interface. These are only some of the criteria that can be used.

? Policy-based routing affects only the next hop in the path to the destination.

? Policy-based routing does not affect the destination of the packet. It affects the path used to get to the destination.

? Policy-based routing does not allow traffic sent into another autonomous system to take a different path from the one that would have been chosen by that autonomous system.

? It is possible to influence only how traffic will get to a neighboring autonomous system, not how it will be routed within that autonomous system.

? As policy-based routing examines the source address, it is configured on the inbound interface.

Reference: Self-Study CCNP BSCI Exam Certification Guide Third Edition P.523

---

### **QUESTION 189:**

Which one of the following BGP prefix lists would you use if you wanted to deny the default route 0.0.0.0?

- A. ip prefix-list abc deny 0.0.0.0/0
- B. ip prefix-list abc permit 0.0.0.0/32
- C. ip prefix-list abc deny 255.255.255.255/0
- D. ip prefix-list abc permit 255.255.255.255/32

Answer: A

Explanation:

To deny the default route 0.0.0.0/0 use the following command:

```
ip prefix-list abc deny 0.0.0.0/0
```

Incorrect Answers:

B, D: We are required to deny it, not permit it.

C: This will effectively deny all routes, not just the default route.

---

**QUESTION 190:**

Which of the following prefix-list statements would you enter if you wanted to permit all prefixes that fall between /8 and /24 on the 192.0.0.0 network?

- A. ip prefix-list abc permit 192.0.0.0/8 ge 8 le 24
- B. ip prefix-list abc permit 192.0.0.0/8 ge 24 le 8
- C. ip prefix-list abc permit 192.0.0.0/24 ge 24 le 8
- D. ip prefix-list abc permit 192.0.0.0/24 ge 8 le 24

Answer: A

Explanation:


When used with prefix lists, the "ge" means "greater than or equal to" and the term "le" means "less than or equal to." Greater than (ge) 8 and less than (le) 24.

An 8 bit subnet mask should also be used, since we are looking to match all prefixes beginning on the entire 192.X.X.X range.

---

**QUESTION 191:**

The Certkiller IBGP network is displayed below:



**AS 65100**

It is desired to configure IBGP in the best, most stable manner. Which of the following options show the best configuration choice for router Certkiller 1?

- A. router bgp 65100  
neighbor 2.2.2.2 remote-as 65100  
neighbor 2.2.2.2 update-source loopback0
- B. router bgp 65100  
neighbor 10.1.1.2 remote-as 65100  
neighbor 10.2.2.2 remote-as 65100  
neighbor 10.1.1.2 update-source loopback0  
neighbor 10.2.2.2 update-source loopback0
- C. router bgp 65100  
neighbor 10.1.1.2 remote-as 65100

```
neighbor 10.2.2.2 remote-as 65100
neighbor 2.2.2.2 remote-as 65100
D. router bgp 65100
neighbor 10.1.1.2 remote-as 65100
neighbor 10.2.2.2 remote-as 65100
neighbor 10.1.1.2 update-source s0
neighbor 10.2.2.2 update-source s1
```

Answer: A

Explanation:


Option A displays the best choice, and shows the correct syntax for the use of sourcing BGP updates from the loopback interface. Using the loopback interface provides the best option, since this interface will remain up and reachable as long as one of the serial interfaces remains active.

Note: Configuring the routers in this manner can be used for both IBGP and EBGP, and can also be used for load balancing traffic across the two serial links.

---

**QUESTION 192:**

The Certkiller WAN is displayed below:



Part of the Certkiller router configuration is displayed below:

```
Certkiller1#. permitting config
router bgp 100
network 200.200.200.0 mask 255.255.255.0
neighbor 10.10.25.2 remote-as 200
neighbor 10.10.25.2 route-map BEST out
!
access-list 101 permit ip host 200.200.200.0 host 255.255.255.0
!
route-map BEST permit 10
match ip address 101
set metric 150
!
route-map BEST permit 20
```

You work as an administrator at Certkiller .com. Study the exhibits carefully.

BGP is configured on all routers, synchronization is turned off, and one of the default attributes have been changed except the MED attribute on Certkiller 5. Which path is preferred by Certkiller 2 to reach the network 200.200.200.0/24?

- A. Certkiller 2- Certkiller 1 because it is the shortest path.
- B. Certkiller 2- Certkiller 1 because it has a higher metric.
- C. Certkiller 3- Certkiller 4- Certkiller 5 because it has a lower administrative distance.
- D. Certkiller 3- Certkiller 4- Certkiller 5 because it has a lower metric.

Answer: D

Explanation:

In BGP, the lowest MED metric is preferred over a higher metric. By default, the value of the metric is 0, so for AS 200, the path via Certkiller 5 will have a metric of 0, and the path via Certkiller 1 will have a metric of 150, so the path from Certkiller 5 will be preferred.

---

**QUESTION 193:**

Which BGP path attribute is Cisco proprietary?

- A. weight
- B. MED
- C. local preference
- D. origin
- E. next-hop
- F. AS-path

Answer: A


Explanation:

The WEIGHT attribute is a Cisco proprietary feature used to measure a route's preference. The WEIGHT attribute is local to the router and does not get exchanged between routers; therefore it is only effective on inbound route maps. Use the WEIGHT attribute to influence routes from multiple service providers to a central location. Like LOCAL\_PREF, assigning a higher weight to a route makes that route more preferred. The WEIGHT attribute also has the highest precedence of any BGP attribute. Reference: CCIE Practical Studies: Configuring Route-Maps and Policy-based Routing, Karl Solie, Cisco Press.

---

**QUESTION 194:**

The Certkiller BGP network is displayed in the following diagram:



Certkiller4# show running-configuration

```

|
router bgp 200
no synchronization
neighbor 20.20.43.3 remote-as 200
neighbor 10.10.45.5 remote-as 100
neighbor 10.10.45.5 route-map LP1 in
|
access-list 10 permit 100.100.100.0 0.255.255.255
|
route-map LPT permit 10
match ip address 101
set local-preference 300
|

```

Refer to exhibit. BGP is configured on all routers, synchronization is turned off, and none of the default attributes have been changed except the local preference attribute on Certkiller 4. Which path will be preferred by Certkiller 2 to reach the network 100.100.100.0/24

- A. Certkiller 2- Certkiller 3. Certkiller 4. Certkiller 5 because it has a lower admin distance
- B. Certkiller 2- Certkiller 3- Certkiller 4- Certkiller 5 because it has a higher local preference
- C. Certkiller 2- Certkiller 1 because it has the shortest AS-path
- D. Certkiller 2- Certkiller 1 because it has a lower local preference

Answer: B

Explanation:

The BGP decision making process fundamentally breaks down to this:

1. Highest weight. Remember weight is Cisco proprietary, and not transmitted as part of the route, so it is a decision local to the router made through a route map.
2. Highest local preference. Same idea as weight, but this is carried within the AS. The default local preference value is 100.
3. Prefer local routes.
4. Shortest AS-PATH length. Even though the number of AS hops is smaller, this doesn't necessarily mean it'll be faster
5. Lowest origin type (IGP < EGP < incomplete)

6. Lowest multi-exit discriminator (MED), AKA BGP metric. This is used to tell a neighbouring AS which path to take if there are multiple ones between the two ASes. Practically, providers let you set a community that alters their local preference, making the decision earlier in the process.

7. eBGP over iBGP. AD is 20 vs 200

8. Lowest IGP metric to NEXT\_HOP

In this case, the route over Certkiller 4 to Certkiller 5 will be preferred because it has a local preference of 300, and the other link will default to a local preference of 100.

---

**QUESTION 195:**

When you're running an Interior Gateway Protocol (like OSPF or EIGRP); why would you want to redistribute these routes into BGP?

- A. So BGP can propagate this information to other IGP neighbors.
- B. So BGP can propagate this information to other iBGP neighbors.
- C. So BGP can propagate this information to other eBGP neighbors.
- D. So BGP can propagate this information to other OSPF neighbors.

Answer: C

Explanation:

When BGP is used between autonomous systems (AS), the protocol is referred to as External BGP (EBGP). In order to advertise the IP subnets used within your network to other autonomous systems, these routes must be injected into BGP.

Note: Customer networks usually employ an Interior Gateway Protocol (IGP) such as RIP or OSPF for the exchange of routing information within their networks

Reference: Border Gateway Protocol

[http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito\\_doc/bgp.htm](http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm)

---

**QUESTION 196:**

DRAG DROP

You are the administrator of an internal network configured for OSPF. Two gateway routers are learning external routes from BGP and inserting those learned routes into your network via redistribution. One of the gateway routers is called Certkiller 2 and it is inserting the subnets 128.213.64.0 through 128.213.95.0. Your task is to configure the gateway router Certkiller 1 to summarize the subnets into one range before injecting them into OSPF. The first two lines have already been configured. Select the boxes below and place the three correct boxes in the right order to complete the configuration.

Certkiller 1(config)#router ospf 100

Certkiller 1(config-router)#redistribute bgp 50 metric 1000 subnets

Certkiller1(config-router)#

Drop keyword here

Drop keyword here

Drop keyword here

---

Select from these

128.213.64.0

summary-  
address

128.213.0.0

255.255.224.0

network

auto-summary

0.0.31.255

Answer:

Explanation:

Summary-address 128.213.64.0 255.255.224.0

External route summarization, external routes that are injected into OSPF via redistribution, is done via the following router ospf subcommand:

summary-address ip-address mask

This command is effective only on ASBRs doing redistribution into OSPF as in this scenario.

---

**QUESTION 197:**

Your network is running BGP as well as EIGRP. You are considering redistributing your BGP routes into your EIGRP. What factors must you consider before redistributing BGP routes to IGP?

- A. IGP's are limited to 250 routes
- B. A full BGP routing table may contain 100,000+ routes
- C. Because of possible routing loops, Cisco router configuration does not allow BGP routes to be redistributed into an IGP.
- D. Because BGP routes are not advertised unless they are known by the IGP, Cisco automatically redistributes routes into BGP.

Answer: B

Explanation:

The size of the BGP routing table must be taken into consideration. It is generally not recommended to redistribute BGP routes into your IGP due to the large number of routes that can be obtained from BGP.

Incorrect Answers:

A: There is no such limit.

C: BGP routes can very well be distributed into an IGP.

D: BGP routers are not redistributed automatically.

---

**QUESTION 198:**

Why is it sometimes necessary to redistribute IGP protocols like OSPF into BGP?

- A. So that BGP can propagate this information to other IGP neighbors.
- B. So that BGP can propagate this information to other IBGP neighbors.
- C. So that BGP can propagate this information to other EBGP neighbors.
- D. So that BGP can propagate this information to other OSPF neighbors.

Answer: C

Explanation:

Customer networks usually employ an Interior Gateway Protocol (IGP) such as RIP or OSPF for the exchange of routing information within their networks. In order to advertise these networks to other network across the Internet, they must be redistributed into BGP.

Reference: Border Gateway Protocol

[http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito\\_doc/bgp.htm](http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm)

---

**QUESTION 199:**

According to Cisco's official recommendations, what is the preferred method for populating the BGP table on your router with IGP routes?

- A. Use the network command.
- B. Redistribute EBGP routes into BGP.
- C. Redistribute dynamic routes into BGP.
- D. Redistribute static routes into the IGP.

Answer: A

Explanation:

For BGP the network command allows BGP to advertise a network that is already in the IP table.

Note: An alternate solution, not listed here, is to redistribute the IGP routes into BGP.

Incorrect Answers:

B: We want to distribute internal routes into BGP, not external routes.

C: Redistribution from the IGP is NOT advised because there is a great reliance on the IGP table. It could cause instability.

D: We want to populate the BGP table, not the IGP table.

---

**QUESTION 200:**

DRAG DROP

You are the midst of configuring a gateway route named Router CertK . Your task is to configure the router to redistribute the various subnets it learns via BGP AS75 into OSPF, all while using the default metric. Select three of the correct command phrases below and place them in the correct order in the command sequence.

[642-801](#)

RouterCertK(config-router)#

Select from these

|                |        |         |              |
|----------------|--------|---------|--------------|
| bgp 75         | ospf   | subnets | 200          |
| default-metric | metric | area    | redistribute |

Answer:

RouterCertK(config-router)#

Select from these

|                                   |                                  |      |
|-----------------------------------|----------------------------------|------|
| <input type="text" value="ospf"/> | <input type="text" value="200"/> |      |
| default-metric                    | metric                           | area |

Explanation:

We want to redistribute BGP AS 75 into OSPF.

Syntax of the redistribution command:

```
router(config-router)#redistribute protocol [process-id] [metric [metric-type type-value] [route-map map-tag] [subnets] [tag tagvalue]
```

We specify the BGP protocol and the Autonomous System number 75. We also use the subnets keyword, which is an optional OSPF parameter that specifies that subnetted routes should also be redistributed. Only routes that are not subnetted are redistributed if the subnets keyword is not specified.

Incorrect Answers:

ospf: We want to redistribute the routes learned from BGP AS 75 into OSPF so we should specify the BGP protocol.

area: Is not used in the redistribute command.

default-metric: Default-metric is a separate command.

metric, 200: We don't have to specify a metric since we want to use a default metric.

---

**QUESTION 201:**

DRAG DROP

You are the administrator of the Certkiller network which is configured for OSPF.

Two gateway routers are learning external routes from BGP and inserting those learned routes into your network via redistribution. One of the gateway routers is called Certkiller 1 and it is inserting the subnets 142.14.64.0 through 142.14.95.0.

Your task is to configure router Certkiller 1 to summarize the subnets into one range before injecting them into OSPF. The first two lines have already been configured.

Select the boxes below and place the three correct boxes in the right order to complete the configuration.

Certkiller 1(config)#router ospf 200

642-801

Certkiller 1(config-router)#redistribute bgp 80 metric 1100 subnets

Router Certkiller1(config-router)#

Select from these

|              |                 |            |
|--------------|-----------------|------------|
| auto-summary | summary-address | network    |
| 142.14.0.0   | 255.255.224.0   | 0.0.31.255 |
|              | 142.14.64.0     |            |

Answer:

Router Certkiller1(config-router)#

Select from these

|              |            |
|--------------|------------|
| auto-summary | network    |
| 142.14.0.0   | 0.0.31.255 |

Explanation:

The summary-address router configuration command is used to create aggregate addresses for OSPF.

Simplified syntax: summary-addressaddress mask

We study the subnets that the Certkiller 1 router is injecting in binary to decide where to summarize:

Decimal 1st octet 2nd octet 3rd octet 4th octet

142.14.64.0 10001110 00001110 01000000 00000000

142.14.95.0 10001110 00001110 01011111 00000000

255.255.224.0 11111111 11111111 11100000 00000000

The subnet we can summarize on is marked with red. This is the 142.14.64.0/19 network.

It has a 255.255.224.0 subnet mask.

**QUESTION 202:**

DRAG DROP

You are the administrator of an internal network configured for OSPF. Two gateway routers ( Certkiller 1 & Certkiller 2) are redistributing external BGP routes. Certkiller 1 is redistributing subnets ranging from 163.150.64.0 - 163.160.95.0 and Certkiller 2 is doing the same with 163.150.96.0 - 163.150.127.0. Your task is to configure the gateway router Certkiller 1 to summarize these subnets into one range before redistributing them into OSPF. The first two lines have already been configured. Select the boxes below and place the three correct boxes in the right order to complete the configuration.

Exhibit:

Certkiller 1(config)#router ospf 100

Certkiller 1(config-router)#redistribute bgp 70 metric 500 subnets

642-801

Certkiller1(config-router)# Drop keyword here Drop keyword here Drop keyword here

Select from these

163.150.0.0 255.255.224.0 0.0.31.256  
163.150.64.0 network summary-address auto-summary

Answer:

Certkiller1(config-router)# summary-address 163.150.64.0 255.255.224.0

Select from these

163.150.0.0 0.0.31.256  
network auto-summary

Explanation:

External route summarization, external routes that are injected into OSPF via redistribution, is done via the following router ospf subcommand:

summary-address ip-address mask

This command is effective only on ASBRs doing redistribution into OSPF as in this scenario..

We should summarize most matching leftmost bits as indicated below.

Decimal 1st Octet 2nd Octet 3rd Octet 4th Octet

163.50.64.0 10100011 00110010 01000000 00000000

163.50.95.0 10100011 00110010 01011111 00000000

163.50.64.0 10000000 00110010 01000000 00000000 IP address

255.255.224.0 11111111 11111111 11000000 00000000 Subnet mask

Incorrect Answers

Network:

The network command is not used to summarize routes.

auto-summary:

The auto-summary command is used to turn auto summarization on.

128.213.0.0:

We should use a more specific summarization.

0.0.31.256:

For summarization we should specify a network mask, not a wildcard.

**QUESTION 203:**

It isn't always in your best interests to redistribute dynamically learned routes from

an interior gateway protocol to a border gateway protocol. Which of the following are potential issues that can occur when doing so? (Select two)

- A. Routing loops can occur.
- B. The routes are automatically summarized.
- C. External IGP learned routes might not necessarily have originated in this AS.
- D. The BGP process will ignore the external IGP learned routes.

Answer: A, C

Explanation:

If redistribution is used, care must be taken that only local routes are redistributed. For example, routes learned from other autonomous systems (that were learned by redistributing BGP into the IGP) must not be sent out again from the IGP, or routing loops could result.

Incorrect Answers:

B: With redistribution into exterior protocols such as BGP, no routes are automatically summarized.

D: All routes will be included in the redistribution, including externally learned routes.

Reference: Building Scalable Cisco Networks (Cisco Press) page 408.

---

#### **QUESTION 204:**

You are a network administrator on a Cisco router that's running EIGRP. You want to aggregate some of the routes. Which of the following commands below would you enter to configure the summary route you need?

- A. ip auto-summary as-number address mask
- B. ip summary-address as-number address mask
- C. ip auto-summary eigrp as-number address mask
- D. ip summary-route eigrp as-number address mask
- E. ip summary-address eigrp as-number address mask

Answer: E

Explanation:

The ip summary-address eigrp command is used to configure a summary aggregate address for a specified interface.

Syntax:

ip summary-address eigrp autonomous-system-number address mask

Incorrect Answers:

A, C, D: No such command exists using this syntax.

B: We must specify the appropriate protocol with the eigrp keyword.

---

#### **QUESTION 205:**

Router CK1 has a 256kbps serial interface link to another Certkiller location. The configuration file of this interface is shown below:

```
interface serial 0/0  
bandwidth 56  
ip bandwidth-percent eigrp 1 200
```

From the command output above, how many kbps of bandwidth is allocated for EIGRP traffic?

- A. 56 kbps
- B. 112 kbps
- C. 128 kbps
- D. 256 kbps

Answer: B

Explanation:

The bandwidth-percent command tells EIGRP what percentage of the configured bandwidth it may use. The default is 50 percent. Since the bandwidth command is also used to set the routing protocol metric, it may be set to a particular value for policy reasons. In this case, it is set to 56 kbps even though the actual link is 256 kbps. The bandwidth-percent command can have values greater than 100 if the bandwidth is configured artificially low due to such policy reasons. In this specific case, the bandwidth configured on the interface is 56kbps, so 200 percent of this value is 112 kbps.


Reference:

[http://www.cisco.com/en/US/tech/CK3\\_65/CK2\\_07/technologies\\_tech\\_note09186a0080094063.shtml](http://www.cisco.com/en/US/tech/CK3_65/CK2_07/technologies_tech_note09186a0080094063.shtml)

---

**QUESTION 206:**

The Certkiller EIGRP network is displayed below:



All of the routers in the above network segment are configured with EIGRP. If both routers Certkiller 2 and Certkiller 3 were to fail, how would Certkiller 1 react?

- A. Certkiller 1 will automatically route packets via Certkiller 4 to the Certkiller Main Office Network.
- B. Certkiller 1 will place the route via Certkiller 4 into the hold down state.

- C. Certkiller 1 will go into the active state for all routers.
- D. Certkiller 1 will go into the active state for the route to Certkiller Main Office Network.

Answer: D

Explanation:

First, Certkiller 1 it checks the topology table for an alternate route then it sends a query to neighboring routers for other routes to the destination. In this case, since both the optimal route and the feasible successor failed, Certkiller 1 must go into the active state first before routing around the failure via Certkiller 4. The re-route will happen only after Certkiller 1 transitions into the active state and learns about the route through Certkiller 4.

---

**QUESTION 207:**

You are the network engineer at Certkiller . The Certkiller network has five paths from Router CK1 to a given destination. The local EIGRP metric on Router CK1 for these paths are as follows:

CK1

-----  
Path 1: 1500

Path 2: 1500

Path 3: 2000

Path 4: 4000

Path 5: 4000

Variance 3 is configured on Router CK1 . Which paths would be included in Router CK1 's routing table?

- A. path 1, 2, and 4
- B. path 1, 2, and 3
- C. path 1, 2, and 5
- D. path 1, 2, 4 and 5
- E. path 1, 2, 3, 4, and 5
- F. None of the above

Answer: E

Explanation:

If the variance number is higher than the default 1, the EIGRP process multiplies the best (lowest) cost or metric value for a path by the number stated as the variance multiplier. All paths to the same destination that have metrics within this new range are now included in load balancing. Per the question this means:  $1500 * 3 = 4500$ . All available options are within 1500-4500.

Reference: Self-Study CCNP BSCI Exam Certification Guide Third Edition P.485,  
Topic: Load Balancing in EIGRP.

**QUESTION 208:**

Which one of the following commands would configure EIGRP for IP on your Cisco router?

- A. ip eigrp routing
- B. router eigrp process-id
- C. ip eigrp autonomous-system-number
- D. router eigrp autonomous-system-number

Answer: D

Explanation:

Perform the following steps to configure EIGRP for IP:

Step 1 Enable EIGRP and define the autonomous system.

```
routerCK(config)#router eigrp autonomous-system-number
```

Step 2 Indicate which networks are part of the EIGRP autonomous system.

```
routerCK(config-router)#network network-number
```

Step 3 Define bandwidth of a link for the purposes of sending routing update traffic on the link.

```
routerCK(config-if)#bandwidth kilobits. This step is optional.
```

---

**QUESTION 209:**

The syntax for configuring variance on an EIGRP router is:

```
variance number
```

In the EIGRP command sequence above, what is the function of the number variable?

- A. It acts as a multiplier.
- B. It defines the limit for how far the metrics can be separated.
- C. It indicates how many paths can be used for load balancing.
- D. It indicates how many paths can be used for unequal load balancing.

Answer: A

Explanation:

Every routing protocol supports equal cost path load balancing. IGRP and EIGRP also support unequal cost path load balancing, which is known as variance. The variance number command instructs the router to include routes with a metric less than or equal to number times the minimum metric route for that destination. Thus number is used as multiplier, even though it defines a limit for the metrics. This limit is used for unequal load balancing.


Incorrect Answers:

B: Number does not directly define a limit.

C, D: Number does not directly indicate the number of paths that can be used.  
Reference: <http://www.cisco.com/warp/public/103/19.html>

**QUESTION 210:**

The Certkiller EIGRP network is displayed in the diagram below:



You are the systems administrator of the above EIGRP network and you've elected to shut down router B's FDDI interface. After you execute the shutdown, which router will become the feasible successor to network 7 for router A?

- A. B
- B. D
- C. H
- D. There is no feasible successor

Answer: D

Explanation:

Router H will be the successor, and that route will be placed in the Routing table. Router A detects the link failure between Router B and network 7. It checks the topology table for a successor. It finds that H is the successor since the advertised distance for H (30) is less than the feasible distance for B (31). However, there is no next best route - no feasible successor. The candidate route through D has an advertised distance (220) that is higher than the feasible distance of the successor route (40).

Note: Successor is a route selected as the primary route to use to reach a destination. Successors are the entries kept in the routing table.

A feasible successor is a backup route. These routes are selected at the same time the successors are identified, but they are kept in a topology table,

Incorrect Answers:

A: The FDDI interface of B is down.

B: The candidate route through D has an advertised distance (220) that is higher than the

feasible distance of the successor route (40). It cannot be used as a feasible successor.  
C: Router H is the successor, not the feasible successor.

---

**QUESTION 211:**

You are the network engineer at Certkiller . You want to configure EIGRP to run across a 56 Kbps serial PPP link on the Certkiller network. You also want to ensure the proper convergence of EIGRP routes.

What command should you issue on the serial interface?

- A. bandwidth 56
- B. bandwidth 56000
- C. ip bandwidth-percent eigrp 1 56
- D. ip bandwidth-percent eigrp 1 56000

Answer: A

Explanation:

The bandwidth is a logical construct whose value can have wide-reaching implications on the function of your network. It does not affect the actual speed of the link. In fact, it is practical to configure the bandwidth command only on serial lines, where the speed of the link will vary considerably. This command is useful for routing protocols that use bandwidth as part of the metric. For serial interfaces, this value defaults to a full T1, so it may appear to be a better link than what is actually being used.

Router (config) # interface S0

Router (config-if) # bandwidth speed-of-line

The bandwidth command is always specified in kbps.

EIGRP limits itself to 50 percent of the value specified in the bandwidth command, or if the bandwidth command is not set, the interface defaults. If you need to limit this percentage further, the upper limit that EIGRP uses can be stated as a percentage of the bandwidth command.

The ip bandwidth-percent-eigrp command interacts with the bandwidth command on the interface. You would use this command primarily because in your network, the bandwidth command does not reflect the true speed of the link. The bandwidth command might have been altered to manipulate the routing metric and path selection of a routing protocol, such as IGRP or OSPF. It might be better to use other methods of controlling the routing metric and return the bandwidth to a true value. Otherwise, the ip bandwidth-percent eigrp command is available. It is possible to set a bandwidth percent that is larger than the stated bandwidth. This is with the understanding that, although the bandwidth might be stated to be 56 kbps, the link is in fact 256 kbps. The following shows the structure of the ip bandwidth-percent eigrp command:

Router (config) #interface S0


Router (config-if) #ip bandwidth-percent eigrp autonomous-system-number percent

---

**QUESTION 212:**

## 642-801

The Certkiller Network topology is displayed in the diagram below:



Based on the information above, what optional EIGRP configuration will be required in order to achieve full network connectivity?

- A. Use the EIGRP no auto-summary command on CK1 and CK2 .
- B. Use the EIGRP no auto-summary command on CK3 and CK4 .
- C. Use the passive interface on the CK1 and CK2 interface that connects to the 10.1.1.0/24 and 10.1.2.0/24 subnet respectively.
- D. Use the passive interface command between the CK3 and CK1 connection and between the CK3 and CK2 connection.
- E. Use the variance command on CK3 .

Answer: A

Explanation:

A useful feature of EIGRP is automatic route summarization; this summarizes subnets to the classful network boundary. This is enabled by default, you can turn this off per AS by using the following command in router configuration mode:


```
Router(config-router)#no auto-summary
```

In the example above, this is needed, otherwise both CK1 and CK2 will advertise the summarized network 10.0.0.0/8 to router CK3 , making their networks unreachable.

---

### **QUESTION 213:**

You are the network engineer at Certkiller . The Certkiller network is shown in the following graphic:



Router RTX is configured with the following commands:

Interface serial 0

Encapsulation frame relay

You want to complete an EIGRP hybrid multipoint configuration on RTX S0 by configuring it with the correct bandwidth parameter.

Which command should you issue?

- A. bandwidth 56
- B. bandwidth 224
- C. bandwidth 256
- D. bandwidth 768
- E. Non of the above

Answer: B

Explanation:

If the multipoint network has different speeds allocated to the VCs, take the lowest CIR and simply multiply it by the number of circuits. This is applied to the physical interface. In the case of the question  $4 * 56 = 224$  - bandwidth 224

Reference: Self-Study CCNP BSCI Exam Certification Guide Third Edition P. 490.

---

### QUESTION 214:

The Certkiller multipoint frame relay network is using EIGRP for the routing protocol. In this network, how is the bandwidth determined for each multipoint neighbor?

- A. Bandwidth command per neighbor.
- B. The configured CIR per subinterface.

- C. The configured CIR divided by the number of neighbors on that interface.
- D. Bandwidth of the main interface divided by the number of neighbors on that interface.

Answer: D

Explanation:

When configuring multipoint interfaces, especially for Frame Relay, remember that all neighbors share the bandwidth equally, regardless of the actual CIR of each individual PVC. For multipoint interfaces, the bandwidth used by the IOS is taken from the configured bandwidth of the main interface, divided by the total number of neighbors on that interface.

---

**QUESTION 215:**

Your network consists of a router named R1 that's configured in a hub and spoke topology, supporting 24 remote office via a point-to-multipoint Frame Relay EIGRP network deployment. The bandwidth command was not manually entered on the frame relay main interface or the sub-interfaces. What is the perceived bandwidth (as seen by the EIGRP process) of each Frame Relay connection?

- A. 64 kbps
- B. 128 kbps
- C. 1.544 Mbps
- D. 1.536 Mbps
- E. None of the above

Answer: A

Explanation:

The default bandwidth for all serial WAN interfaces is 1.544 Mbps. EIGRP by default assumes the bandwidth of a serial interface is a full T1 (1.544 Mbps) if not specified. This includes sub-interfaces. For multipoint frame relay interfaces, the perceived bandwidth used by the IOS is taken from the configured bandwidth, divided by the number of configured neighbors. Since there are 24 remote offices, there are 24 configured frame relay neighbors, so the bandwidth used by the IOS is the full T1 divided by 24, which is 64kbps.

---

**QUESTION 216:**

EIGRP has been configured on the WAN links of router CK1 . On one of the serial T1 interfaces, the bandwidth statement was not specified. By default, how will the EIGRP process perceive this T1 link?

- A. 256 Kbps
- B. 1.544 Mbps

- C. It depends as it is set by the PVC
- D. It is set by the DLCI

Answer: B

Explanation:

The default EIGRP link speed is 1.544 Mbps for non-high speed WAN links

Note: The enhanced code uses the "bandwidth" subcommand on interfaces and sub-interfaces in order to determine the rate at which to generate EIGRP packets. This parameter is automatically set on fixed-bandwidth interfaces (such as LANs), but defaults to T1 (1544 Kbps) for all serial media.

Incorrect Answers:

A: The default speed is 1.544 Mbps, not 256 Kbps.

C, D: The default speed is not set by the PVC or the DLCI.

---

**QUESTION 217:**

You are in the midst of configuring an NBMA network with EIGRP as the routing protocol. Which of the options below would you use to configure the interface bandwidth for a point-to-point interface?

- A. The DLCI assigns the bandwidth for the interface.
- B. The sliding window size determines interface bandwidth.
- C. You should use the default bandwidth assigned to the interface.
- D. You should manually configure bandwidth as the CIR of the PVC.

Answer: D

Explanation:

The bandwidth can be configured separately on each subinterface. Since this is NBMA we can assume that Frame Relay is used. For Frame Relay on point-to-point the bandwidth should be set it to the CIR of the PVC.

Note 1: NBMA (Non-broadcast Multi-access) supports many (more than two) routers, but have no broadcast capability. Frame Relay and X.25 are example of NBMA.

Note 2: The CIR (Committed Information Rate) is the committed rate (in bits per second) at which the ingress access interface trunk interfaces, and egress access interface of a Frame Relay network transfer information to the destination Frame Relay end system under normal conditions.

Incorrect Answers:

A: Does not apply.

B: Sliding windows does not apply.

C: The bandwidth can and should be configured separately on each subinterface.

Reference: Configuration Notes for the Enhanced Implementation of EIGRP

<http://www.cisco.com/warp/public/103/12.html>

---

**QUESTION 218:**

Router CK2 is configured for EIGRP as shown below:

```
router eigrp 100
network 10.0.0.0
eigrp stub
```

Based on this configuration, which types of routes will be advertised with the EIGRP configuration as shown? (Choose two)

- A. Static
- B. Receive only
- C. Summary
- D. Stub
- E. Connected
- F. dynamic

Answer: C, E

Explanation:

A router that is configured as a stub with the eigrp stub command shares connected and summary routing information with all neighbor routers by default.

Reference:

[http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products\\_feature\\_guide09186a0080087026.html](http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a0080087026.html)

---

**QUESTION 219:**

The Certkiller EIGRP network is utilizing the EIGRP stub routing feature. Which of the following are key concepts that apply when configuring the EIGRP stub routing feature in a hub and spoke network? (Select three)

- A. A hub router prevents routes from being advertised to the remote router.
- B. Only remote routers are configured as stubs.
- C. Stub routers are not queried for routes.
- D. Spoke routers connected to hub routers answer the route queries for the stub router.
- E. A stub router should have only EIGRP hub routers as neighbors.
- F. EIGRP stub routing should be used on hub routers only.

Answer: B, C, E

Explanation:

The Enhanced Interior Gateway Routing Protocol (EIGRP) Stub Routing feature improves network stability, reduces resource utilization, and simplifies stub router configuration.

Stub routing is commonly used in a hub and spoke network topology. In a hub and spoke network, one or more end (stub) networks are connected to a remote router (the spoke) that is connected to one or more distribution routers (the hub). The remote router is

adjacent only to one or more distribution routers. The only route for IP traffic to follow into the remote router is through a distribution router. This type of configuration is commonly used in WAN topologies where the distribution router is directly connected to a WAN. The distribution router can be connected to many more remote routers. Often, the distribution router will be connected to 100 or more remote routers. In a hub and spoke topology, the remote router must forward all nonlocal traffic to a distribution router, so it becomes unnecessary for the remote router to hold a complete routing table. Generally, the distribution router need not send anything more than a default route to the remote router.

When using the EIGRP Stub Routing feature, you need to configure the distribution and remote routers to use EIGRP, and to configure only the remote router as a stub. Only specified routes are propagated from the remote (stub) router. The router responds to queries for summaries, connected routes, redistributed static routes, external routes, and internal routes with the message "inaccessible." A router that is configured as a stub will send a special peer information packet to all neighboring routers to report its status as a stub router.

This feature should only be used on stub routers. A stub router is defined as a router connected to the network core or distribution layer through which core transit traffic should not flow. A stub router should not have any EIGRP neighbors other than distribution routers. Ignoring this restriction will cause undesirable behavior.


Reference:

[http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products\\_feature\\_guide09186a0080087026.html](http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a0080087026.html)

---

**QUESTION 220:**

The Certkiller WAN is displayed in the following diagram:



Traffic from Certkiller 4 to the Certkiller Corporate Network is distributed between the links with unequal costs in the EIGRP network by configuring the variance command on Certkiller 4 to 2. How many paths will participate in the load sharing?

- A. 1
- B. 2
- C. 3
- D. 4

- E. 5
- F. 6

Answer: B

Explanation:

EIGRP automatically load balances across links of equal cost. Whether the traffic is sent on a predestination or round-robin basis depends on the internal switching within the router. It is possible to configure EIGRP to load balance across unequal-cost paths using the variance command.

The variance command allows the administrator to identify the metric scope for including additional paths by the use of a multiplier parameter. The command structure follows:

```
Router(config-router)#variance multiplier
```

The multiplier argument is the metric value used for load balancing. It can be a value from 1 to 128. The default is 1, which means equal-cost load balancing.

Reference: CCNP Self-Study CCNP BSCI Exam certification guide p.485

---

### **QUESTION 221:**

Which configuration command is used to enable EIGRP unequal-cost path load balancing?

- A. maximum-paths
- B. distance
- C. metric
- D. variance
- E. default-metric

Answer: D

Explanation:

EIGRP automatically load balances across links of equal cost. Whether the traffic is sent on a predestination or round-robin basis depends on the internal switching within the router. It is possible to configure EIGRP to load balance across unequal-cost paths using the variance command.

The variance command allows the administrator to identify the metric scope for including additional paths by the use of a multiplier parameter. The command structure follows:

```
Router(config-router)#variance multiplier
```


The multiplier argument is the metric value used for load balancing. It can be a value from 1 to 128. The default is 1, which means equal-cost load balancing.

Reference: CCNP Self-Study CCNP BSCI Exam certification guide p.485

---

### **QUESTION 222:**

The Certkiller WAN is displayed below:



Examine the EIGRP network diagram. All paths should be used for load balancing between Certkiller 4 and the Certkiller Corporate Network. Which value should be assigned to the variance command to accomplish this?

- A. 1
- B. 2
- C. 2.5
- D. 3
- E. 5
- F. 6

Answer: D


Explanation:

Every routing protocol supports equal cost path load balancing. In addition to that, IGRP and EIGRP also support unequal cost path load balancing. Use the variance command to instruct the router to include routes with a metric less than  $n$  times the minimum metric route for that destination, where  $n$  is the number specified by the variance command. The variable  $n$  can take a value between 1 and 128, with the default being 1, which means equal cost load balancing. Traffic is also distributed among the links with unequal costs, proportionately, with respect to the metric. In this example, the worst route is 2.5 times worse than the best route. Since we can not use 2.5 for the variance (the value must be an integer) we must specify a minimum value of 3.

---

### QUESTION 223:

The Certkiller EIGRP network is displayed in the following diagram:



You work as a network technician at Certkiller .com. Study the exhibits carefully. If the following command, variance 2, were added to Certkiller 5, which path or paths would be chosen to route traffic from Certkiller 5 to network X?

- A. Only Certkiller 5- Certkiller 4- Certkiller 1
- B. Only Certkiller 5- Certkiller 2- Certkiller 1
- C. Only Certkiller 5- Certkiller 3- Certkiller 1
- D. Both Certkiller 5- Certkiller 2- Certkiller 1 and Certkiller 5- Certkiller 3- Certkiller 1
- E. Both Certkiller 5- Certkiller 2- Certkiller 1 and Certkiller 5- Certkiller 4- Certkiller 1
- F. All available paths.

Answer: D

Explanation:

Variance is 2 so  $2 * (10+10) = 40$

### QUESTION 224:

Part of the configuration for a Certkiller router is shown below:

```

interface serial0
bandwidth 40
ip bandwidth-percent eigrp 1 200
  
```

In the configuration in the exhibit, how much bandwidth will be used by the EIGRP updates?

- A. 40 kbps
- B. 60 kbps
- C. 80 kbps
- D. 200 kbps
- E. 40 Mbps

- F. 60 Mbps
- G. 80 Mbps
- H. 200 Mbps

Answer: C


Explanation:

By default, a serial interface will assume a bandwidth value of 1544, which is equal to a full T1. This default value can be overridden via the bandwidth configuration statement as shown above. From the file shown in this example, serial 0 is configured as a 40kbps (the bandwidth statement is used in terms of kbps) and EIGRP is configured to use up to twice this amount (200 percent) so up to 80 kbps will be consumed via the EIGRP process.

---

**QUESTION 225:**

Five Certkiller routers are connected together as shown in the diagram below:



Refer to the exhibit. Certkiller 1 and Certkiller 2 are summarizing the 10.0.0.0/8 network and advertising it to all routers. Which two EIGRP configurations will prevent unnecessary EIGRP updates from being sent and possibly causing 10.0.0.0/8 from being SIA on router Certkiller 1 or Certkiller 2? (Select two)


- A. Configure Certkiller 3, Certkiller 4, and Certkiller 5 as EIGRP stub routers.
- B. Configure an ACL on the interface of Certkiller 1 and Certkiller 2.
- C. Configure Certkiller 1 and Certkiller 2 as EIGRP stub routers.
- D. Configure summary routes on Certkiller 1 and Certkiller 2
- E. Configure static routing on Certkiller 3, Certkiller 4, and Certkiller 5

Answer: A, D

Explanation:

Dual-Homed Remote

A dual-homed remote will have two or more distribution (hub) routers. However, the principles of stub routing are the same as they are with a hub and spoke topology.



It is not desirable for traffic from distribution router 2 to travel through any remote router in order to reach network 10.3.1.0/24. If the links are sized to handle the load, it would be acceptable to use one of the backup routes. However, most networks of this type have remote routers located at remote offices with relatively slow links. This problem can be prevented if proper summarization is configured on the distribution router and remote router.

It is typically undesirable for traffic from a distribution router to use a remote router as a transit path. A typical connection from a distribution router to a remote router would have much less bandwidth than a connection at the network core. Attempting to use a remote router with a limited bandwidth connection as a transit path would generally produce excessive congestion to the remote router. The EIGRP Stub Routing feature can prevent this problem by preventing the remote router from advertising core routes back to distribution routers. Routes learned by the remote router from distribution router 1 will not be advertised to distribution router 2. Since the remote router will not advertise core routes to distribution router 2, the distribution router will not use the remote router as a transit for traffic destined for the network core.

Reference:

[http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products\\_feature\\_guide09186a0080087026.html](http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a0080087026.html)

### QUESTION 226:

Which command is used to change the administrative distance for EIGRP?

- A. Router(config-if)# distance eigrp internal-distance external-distance
- B. Router(config-if)# distance weight eigrp internal-distance external-distance
- C. Router(config-router)# distance eigrp internal-distance external-distance
- D. Router(config-router)# distance weight eigrp internal-distance external-distance

Answer: C

Explanation:

To allow the use of two administrative distances-internal and external-that could be a

better route to a node, use the distance eigrp command in router configuration mode. To reset these values to their defaults, use the no form of this command.

distance eigrp internal-distance external-distance

no distance eigrp

Syntax Description

|                   |  |
|-------------------|--|
| internal-distance | Administrative distance for EIGRP internal routes. Internal routes are those that are learned from another entity within the same autonomous system. The distance can be a value from 1 to255.             |
| external-distance | Administrative distance for EIGRP external routes. External routes are those for which the best path is learned from a neighbor external to the autonomous system. The distance can be a value from1to255. |

Defaults

internal-distance: 90

external-distance: 170


Reference:

[http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products\\_command\\_reference\\_chapter09186a00800c](http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products_command_reference_chapter09186a00800c)

---

**QUESTION 227:**

The Certkiller OSPF backbone consists of four routers as shown below:



Assuming that the OSPF router ID was not explicitly defined on router CK2 , what would its router ID default to?

- A. 10.5.0.2
- B. 10.7.0.3
- C. 10.9.0.1
- D. 10.10.10.10
- E. 0.0.0.0
- F. None of the above

Answer: D

Explanation:

The router ID is the tie-breaker for OSPF path selection. The path selection process uses a variety of metrics to select a route. If all other metrics (accessibility, administrative weight, local preference, etc.) are equal, OSPF determines the router ID using the following priority:

1. Use the address configured by the ospf router-id command
2. Use the address of the loopback 0 interface
3. Use the highest IP address of any interface
4. If no interface exists, set the router-ID to 0.0.0.0

---

**QUESTION 228:**

**SIMULATION**

The Berlin and Nuremberg OSPF networks need to be established and configured in the following manner:

1. The router is named Berlin.
  2. The clocking is provided on the Berlin router's serial 0/0 interface.
- \* The secret password in the Berlin router is " Certkiller ".
  - \* You should use OSPF in a single area

1. Area 0 should be used for the routing protocol.
2. Use 1 for the process ID of the routing protocol.
3. The IP addresses and subnet masks are listed in the chart.

Berlin

FA0/0 10.60.2.1/23

S0/0 10.60.4.1/30


Lo 0 10.90.20.10/32

Nuremberg

FA0/0 10.60.6.1/24

S0/0 10.60.4.2/30

Lo 0 10.90.20.20/32



Please provide the complete configuration and also the prompt where the configuration is to be done.

Start by clicking on host that is connected to the router you want to configure.

Lab A

Berlin

FA0/0 10.60.2.1/23

S0/0 10.60.4.1/30

Lo 0 10.90.20.10/32

Secret Password: Certkiller

Lab B

Nuremberg

FA0/0 10.60.6.1/24

S0/0 10.60.4.2/30

Lo 0 10.90.20.20/32

Answer:

Explanation:

<Click on Host2, which is connected to the Lab 1 router>

Berlin>enable <--- notice the prompt

Password: Certkiller

```
Berlin# configure terminal
Berlin(config)#router ospf 1
Berlin(config-router)# network 10.60.2.1 0.0.1.255 area 0 <--- this specific mask should
be use to reflect the subnet mask of 23
Berlin(config-router)# network 10.60.4.1 0.0.0.3 area 0 <--- this specific mask should be
use to reflect the subnet mask of 30
Berlin(config-router)# network 10.90.20.10 0.0.0.0 area 0 <--- this > specific mask
should be use to reflect the subnet mask of 32
Berlin(config-router)#exit
Berlin(config)#^Z
Berlin#copy running-config startup-config
```

---

**QUESTION 229:**

Which command should you use to verify what networks are being routed by a given OSPF process?

- A. show ip ospf
- B. show ip route
- C. show ip protocol
- D. show ip ospf database

Answer: C

Explanation:

The information displayed by the show ip protocols command is useful in debugging routing operations. Information in the Routing Information Sources field of the show ip protocols output can help you identify a router suspected of delivering bad routing information. For OSPF routers, this command will display the routed networks.

Incorrect Answers:

A: To display general information about Open Shortest Path First (OSPF) routing processes, use the show ip ospf command in EXEC mode. This command will display the areas assigned and other useful information, but not the networks being routed.

Example:

```
CK1 # show ip ospf
Routing Process "ospf 201" with ID 192.42.110.200
Supports only single TOS(TOS0) route
It is an area border and autonomous system boundary router
Redistributing External Routes from,
igrp 200 with metric mapped to 2, includes subnets in redistribution
rip with metric mapped to 2
igrp 2 with metric mapped to 100
igrp 32 with metric mapped to 1
Number of areas in this router is 3
Area 192.42.110.0
Number of interfaces in this area is 1
```

Area has simple password authentication

SPF algorithm executed 6 times


B: This will display the active routing table, but not the networks that are being routed.

D: The OSPF database does not display the networks being routed.

---

**QUESTION 230:**

OSPF is being configured over the Certkiller frame relay network as displayed in the diagram below:



Which three of the following commands would you enter on the Certkiller 1 serial 0 interface, if you wanted to implement OSPF on this network? (Select three)

- A. ip ospf network point-to-point
- B. ip ospf network point-to-multipoint
- C. frame-relay map ip 10.1.1.1 200
- D. frame-relay map ip 10.1.1.3 300
- E. frame-relay map ip 10.1.1.1 200 broadcast
- F. frame-relay map ip 10.1.1.3 300 broadcast

Answer: B, E, F

Explanation:

The ip ospf network command, typed under the interface configuration mode, is used to specify the OSPF network configuration and sets the network mode to point-to-multipoint

DLCI- Data-link connection identifier (DLCI) number.

Broadcast - Forwards broadcasts to the specified IP address. This keyword is needed on the frame relay mappings in order to transmit the OSPF information across the frame

relay network.

Reference:

[http://www.cisco.com/en/US/products/sw/iosswrel/ps1830/products\\_feature\\_guide09186a0080087b42.html](http://www.cisco.com/en/US/products/sw/iosswrel/ps1830/products_feature_guide09186a0080087b42.html)

---

**QUESTION 231:**

Router CK1 is configured as shown below:

```
interface serial0
ip address 110.1.1.1 255.255.255.0
encapsulation frame-relay
ip ospf network point-to-multipoint
router ospf 51
network 110.1.1.0 0.0.0.255 area 0
```

Based on the information above, which two of the following statements are true?

(Select two)

- A. DB/BDR elections do not take place.
- B. The router is restricted to a hub and spoke topology.
- C. The area 0 NBMA cloud is configured as more than one subnet.
- D. OSPF neighbor statements are not necessary.
- E. The OSPF hello timers need to be adjusted manually.

Answer: A, D

Explanation:

In an OSPF Point-to-Multipoint environment, DB/BDR elections do not take place.

The

neighbor command became somewhat obsolete with the introduction of the capability to configure other network modes for the interface, regardless of the underlying physical topology.

Reference: Building Scalable Cisco Networks (Cisco Press) page 130 and 181

Point-to-Multipoint Network:

Point-to-multipoint is a single interface that connects to multiple destinations. The underlying network treats the network as a series of point-to-point circuits. It replicates LSA packets for each circuit. OSPF traffic is sent as multicast. There is no DR or BDR election. This technology uses one IP subnet for all endpoints on the network.

By default, the network is considered to be a series of point-to-point interfaces. There is no need to specify neighbors, because the neighbors will see each other and simply become adjacent, with no need for the election of a DR or a BDR. Point-to-multipoint does not try to reduce adjacencies using a DR. Instead, it accepts the extra overhead of having a full set of adjacencies for the sake of stability. Point-to-multipoint forms an adjacency automatically along any PVC, which causes more overhead but is more resilient than NBMA.

---

**QUESTION 232:**

On the Certkiller NBMA Frame Relay network, subinterfaces were configured. In routing OSPF over the frame relay network, what advantage does using subinterfaces provide?

- A. To conserve IP addressing space.
- B. To avoid split-horizon issues with the routing protocol.
- C. Because logical interfaces are more reliable than physical interfaces.
- D. Subinterfaces remain up when the physical interface changes to a down state.
- E. All of the above.

Answer: B

Explanation:

When configuring routers in a NBMA topology, subinterfaces are typically used. A physical interface can be split into multiple logical interfaces, called subinterfaces, with each subinterface being defined as point-to-multipoint interface. Subinterfaces originally were created to better handle issues caused by split horizon over NBMA and distance vector-based routing protocols.

Incorrect Answers:

A: Using subinterfaces creates separate IP subnets for each frame relay link, so actually more IP addresses are used this way.

C: Using logical interfaces has no impact on the stability of any network link.

D: When the physical link goes down, all subinterfaces also go down.

Reference: Building Scalable Cisco Networks (Cisco Press) page 120.

---

### **QUESTION 233:**

You have been assigned the task of connecting two office networks together via a frame relay network, and running OSPF across this network. What kind of configuration structure would you use to accomplish this?

- A. Point-to-point over sub-interfaces.
- B. Point-to-multipoint star configuration.
- C. Point-to-multipoint using a single subnet.
- D. Point-to-multipoint nonbroadcast using a single subnet.
- E. None of the above.

Answer: A

Explanation:

Using point to point subinterfaces will eliminate the issues surrounding split horizons over NBMA networks. Configuring the network in this way is the Cisco recommended solution.

Reference: RFC1586, Guidelines for Running OSPF over Frame Relay Network.

**QUESTION 234:**

A Certkiller router is configured as shown below:

```
interface serial 0
ip address 164.67.36.1 255.255.255.224
encapsulation frame-relay
ip ospf network non-broadcast
!
router ospf 1
network 164.67.36.0 0.0.0.31 area 0
neighbor 164.67.36.2
neighbor 164.67.36.3
```

Based on this configuration, which of the following statements are true? (Select two)

- A. There can be no DR or BDR in this configuration.
- B. This is a point-to-point configuration over Frame Relay.
- C. The network mode is nonbroadcast multiaccess (NBMA).
- D. The DR and BDR need a static list of neighbors due to non-broadcast.

Answer: C, D

Explanation:

The configuration printout shows an OSPF configuration across a frame relay network that is not configured for subinterfaces. With this setup you are required to specify neighbors so that a DR and BDR can be elected. In addition, all locations are part of the same IP subnet so the OSPF neighbors need to be explicitly assigned at the host location of an NBMA network.


Note: Had subinterfaces been used, the network would be seen as a series of logical point to point links, so the OSPF neighbors would not need to be specified.

---

**QUESTION 235:**

**SIMULATION**

Network topology exhibit



Certkiller .com is planning to provide network connectivity for two its sites:

## 642-801

Certkiller 1 and Certkiller 2. Each site is to have one LAN. You have been asked to configure the first router at the Certkiller 1 location. The Certkiller 1 router has been completely configured except for the routing protocol. Configure OSPF as the routing in a single area to allow a host on the LAN of the Certkiller 1 router to communicate with a host on the LAN of the Certkiller 2 router.

Due to the fact that adjacent subnets are in use or are planned for future growth on other routers in the AS, make sure you use specific subnet information in your configuration. Please note, the Certkiller 2 router will be installed at a later time. The Certkiller 1 router has been configured with the following specifications:

The router is named Certkiller 1.

The clocking is provided on the Certkiller 1's serial 0/0 interface.

The secret password on the Certkiller 1 router is ' Certkiller '.

Name: Certkiller 1

FA0/0 172.10.143.1/24

S0/0 172.20.44.1/30

Lo 0 172.30.10.197/32

Name: Certkiller 2

FA0/0 172.10.144.1/24

S0/0 172.20.44.2/30

Lo 0 172.30.10.198/32

Secret password: Certkiller

Answer:

Explanation:

Click on Host 2

en

Certkiller

config t

router ospf 1

network 172.10.143.0 0.0.0.255 area 0

network 172.20.44.0 0.0.0.3 area 0

network 17230.10.197 0.0.0.0 area 0

exit

ctrl z

copy run start

---

### **QUESTION 236:**

In a network running OSPF, what is the term used to describe the administrative process of dividing a large area into smaller areas?

- A. interior areas
- B. OSPF subarea
- C. link-state protocol
- D. hierarchical routing

Answer: D

Explanation:

OSPF's ability to separate a large internetwork into multiple areas is referred to as hierarchical routing. The backbone area, area 0, is considered to be at the top of the hierarchical chain, with the other areas lying below it.

---

**QUESTION 237:**

Assuming that you are configuring an ABR in an OSPF area, which IOS command would you execute if your goal was to summarize the networks advertised out of the area?

- A. summary-address address mask
- B. area area-id range address mask
- C. auto-summary address mask area area-id
- D. network network-number wildcard mask area area-id

Answer: B

Explanation:

The area area-id range address mask command consolidates IA (intra-area) routes on an ABR. This command instructs the ABR to summarize routes for a specific area before injecting them into a different area.

Incorrect Answers:

- A: The summary-address address mask command consolidates external routes (inter-area) on an ASBR
- C: Auto-summarization is not useful here.
- D: The network command cannot be used for this purpose. This command is simply used to add a network to the OSPF routing process.

---

**QUESTION 238:**

Which two of the following characteristics are defined by the network command? (Select two)

- A. The OSPF area ID
- B. The OSPF router ID
- C. The OSPF process ID
- D. Which interface belongs to which OSPF area

Answer: A, D

Explanation:

The network command designates the OSPF area for an interface with the specified IP

address.

Syntax: network address wildcard-mask area area-id

---

**QUESTION 239:**

Router CK1 is configured for OSPF as shown below:

```
router ospf 76
network 172.22.23.0 0.0.0.0 area 1
network 172.18.0.0 0.0.255.255 area 0
area 0 range 172.18.0.0 255.255.0.0
area 1 range 172.22.23.0 255.255.255.0
```

Which of the following statements are true regarding the above configuration?  
(Select three)

- A. The OSPF router ID is 76.
- B. This is an area border router.
- C. The designated router priority is 76.
- D. This router connects area 1 to the backbone area.
- E. Any router interfaces with an address of 172.18.x.x are in area 0.

Answer: B, D, E

Explanation:

An interface may belong to only one area. If a router has multiple interfaces and if any of those interfaces belong to different areas, the router is considered as an area border router. The networks that follow the network command are connected to each other and Area 0 is always the backbone area. Finally, the command "network 172.18.0.0 0.0.255.255 area 0" identifies that all interfaces with IP address of 172.18.0.0 area within area 0.

Incorrect Answers:

A, C: The command router ospf 76 identifies the process ID as 76 and not the router ID nor the router priority.

Catherine Paquet and Daine Teare, "Building Scalable Cisco Networks" (Cisco Press 2001), pp 178-181.

---

**QUESTION 240:**

When configuring a multi-area OSPF network to summarize routes, what additional command is required by ASBR's that is not needed by ABR's?

- A. area range command
- B. ospf summarize command
- C. aggregate-route command
- D. summary-address command
- E. None of the above

Answer: D

Explanation:

Using this command for OSPF causes an OSPF autonomous system boundary router (ASBR) to advertise one external route as an aggregate for all redistributed routes that are covered by the address. For OSPF, this command summarizes only routes from other routing protocols that are being redistributed into OSPF.

Incorrect Answers:

A: The area range command is used only with area border routers (ABRs). It is used to consolidate or summarize routes for an area. The result is that a single summary route is advertised to other areas by the ABR.

B, C: There are no such commands for OSPF.

Reference: OSPF Commands

---

**QUESTION 241:**

When designing OSPF networks, maintaining stability in an area is important. Which of the following describes a reason for this?

- A. Instability causes more LSAs to be sent, requiring more CPU to recalculate routes.
- B. Convergence cannot happen until holddown timers expire, so routing loops can occur.
- C. Flooding the area topological database instances consumes excessive bandwidth.
- D. Summary link LSAs cannot be sent until all routers in the OSPF area have the same topological database.

Answer: A

Explanation:

Instability of routes (links) would force sending of LSAs (Link State Advertisements), and CPU time would be required to recalculate the routes.

Incorrect Answers:

B: OSPF does not use hold down timers. Holddown timers are used by distance vector protocols, such as RIP.

C: The topological database is not distributed, only link changes.

D: Summary-link LSAs are not sent to all routers.

Summary-link LSAs originate from area border routers, and flood throughout the LSA's associated area. Each summary-LSA describes a route to a destination outside the area but within the AS.

Reference: <http://www.faqs.org/rfcs/rfc2328.html>

---

**QUESTION 242:**

A Certkiller OSPF router is configured in the following manner:

```
router ospf 200
network 203.42.67.0 0.0.0.255 area 7
network 203.42.68.0 0.0.0.255 area 0
```

```
area 7 stub no-summary  
area 7 default-cost 30
```

Which of the following statements are true regarding this configuration? (Select two)

- A. Area 7 is a totally stubby area.
- B. If the backbone becomes discontinuous, traffic can be routed through area 7.
- C. Redistribution of other routing protocols takes place at the area designated router.
- D. Area 7 non-ABR routers contain only intra-area routing information and a default route.

Answer: A, D

Explanation:

An extension to stub areas is what is called "totally stubby areas". Cisco indicates this by adding a "no-summary" keyword to the stub area configuration. A totally stubby area is one that blocks external routes and summary routes (inter-area routes) from going into the area.

Note: The area stub command is used to define an area as a stub area.

Syntax: area area-id stub [no-summary]

The no-summary optional parameter prevents an ABR from sending summary link advertisements into the stub area.

Reference: OSPF Design Guide

<http://www.cisco.com/warp/public/104/3.html>

---

### QUESTION 243:

Router CK1 is an ABR in the Certkiller OSPF network. What does an ABR connect to?

- A. Multiple OSPF areas
- B. OSPF and RIP networks
- C. Multiple designated routers
- D. Multiple OSPF autonomous systems
- E. Multiple redistributed networks.

Answer: A

Explanation:

An ABR (Area Border Router) shares an interface with at least one other OSPF area. A router that contains network/area statements for two or more different areas is an ABR.

---

### QUESTION 244:

You are a systems administrator of a large multi-area OSPF network, and you've just created a new area for an upcoming remote network. Ordinarily OSPF areas

are be connected to the backbone, area 0. However, circumstances dictate you to connect it to the existing area 2 at this time.

Which conditions have to be met in order to make this configuration work? (Select three)

- A. There must be a virtual link.
- B. Area 2 must be a stub area.
- C. Area 2 cannot be a stub area.
- D. Area 2 must attach directly to area 0.
- E. Network summary link LSAs must be disabled.

Answer: A, C, D

Explanation:

A: There must be a virtual link from the new area to the backbone, area 0. The virtual link provides the disconnected area a logical path to the backbone.

C: The area through which you configure the virtual link, known as a transit area (here area 2), must have full routing information. It cannot be a stub area.

D: The transit area, area 2, must attach directly to area 0.

Reference: OSPF Virtual Link

<http://www.cisco.com/warp/public/104/ospfdb7.html>

Incorrect Answers:

B: The transit area, area 2, cannot be a stub area.

E: This is not a requirement.

---

**QUESTION 245:**

Which of the following is an OSPF configuration parameter that is used on an ABR, but not on an internal router?

- A. A virtual link to area 0.
- B. OSPF summarization command.
- C. default-cost extension to the area command.
- D. no-summary extension to the area stub command.
- E. None of the above

Answer: D

Explanation:

The no-summary extension of the area stub command is used only for ABRs connected to totally stubby areas. It prevents an ABR from sending summary link advertisements into the stub area. This option is used for creating a totally stubby area.

Incorrect Answers:

A: For a virtual link to work both ends need to be configured.

B, C: These commands are not specific to an ABR only.

**QUESTION 246:**

A new location is being added to the Certkiller network and a new OSPF area is being created for it. What kind of OSPF router would you use to connect this new OSPF area to the backbone area?

- A. ABR
- B. Stub router
- C. Internal router
- D. Backbone router

Answer: A

Explanation:

ABRs are used to connect non-backbone areas to the backbone area (area 0).

Incorrect Answers:

B: Stub routers do not apply. An area is stub, not a router and not stub routers are typically a type of ABR, but not every ABR is configured as a stub.

C: Internal routers are only used within an area.

D: Backbone routers sit on the perimeter of the backbone area. They have at least one interface connected to area 0. However, backbone routers do not necessarily connect to other areas.

---

**QUESTION 247:**

Router CK1 is an OSPF ABR that connects area 3 to the network. You wish to configure area 3 to summarize the IP address range 172.16.20.192 - 172.16.20.223. Which two of the following IOS commands would you execute to accomplish this? (Select two)

- A. network 172.16.20.192 0.0.0.31 area 3
- B. area 3 range 172.16.20.192 172.16.20.223
- C. area 3 range 172.16.20.192 255.255.255.224
- D. network 172.16.20.192 255.255.255.224 area 3

Answer: A, C

Explanation:

A: When configuring multiple OSPF areas, make sure to associate the correct network addresses with the desired area ID. Syntax: network address wildcard-mask area area-id.

We must use a wildcard mask (0.0.0.31) and not a network mask (255.255.255.224)


C: We must instruct the ABR to summarize routes for a specific area before injecting them into a different area. Syntax: area area-id range address mask

Incorrect Answers:

B: This is the wrong syntax. We should use a network mask to specify the address range.  
D: We should use a wildcard mask, not a network mask with the network command.

**QUESTION 248:**

The Certkiller OSPF network is displayed in the diagram below:



Based on the information above, what kind of router is RTD in OSPF terminology?

- A. Designated router
- B. Internal router
- C. Backbone router
- D. Area border router
- E. Autonomous system boundary router
- F. None of the above.

Answer: B

Explanation:

An area is a set of networks and hosts within an AS that have been administratively grouped together. We recommend that you configure an area as a collection of contiguous IP subnetted networks. Routers that are wholly contained within an area are called internal routers. All interfaces on internal routers are directly connected to networks within the area.

Incorrect Answers:

- A: Designated routers are defined on multi-access networks. In a network consisting of point to point circuits, the DR/BDR election process does not take place.
- C: Backbone routers are routers that are contained within area 0.

D: Routers that belong to more than one area are called area border routers.

E: Routers that exchange routing information with routers in other ASs are called AS boundary routers.

---

**QUESTION 249:**

Router CK1 is an ABR that is part of the Certkiller OSPF network. CK1 has been configured with the "area range" command. What does this command provide?

- A. It provides the range of areas in the OSPF internetwork.
- B. It provides the range of areas in which this ABR participates.
- C. It provides a summary of networks to be advertised outside the area.
- D. It provides a summary of networks outside the area and inside the AS.

Answer: C

Explanation:

The area range command creates/deletes an area address range entry and optionally specifies whether to advertise the addresses. It is used for summarizing routes into an area.

Syntax: [no] area area-id range address mask [do-not-advertise]

Description:

Address ranges are used to aggregate address ranges from within the area into one single advertisement sent into the adjacent areas, or to prevent the advertisement of networks. An OSPF area is defined as a list of address ranges. Each address range consists of an address and mask pair. Area range entries are used to aggregate network information before the advertisements are flooded into the backbone. A single summary LSA is originated for each range.

---

**QUESTION 250:**

**SIMULATION**

You are the network administrator at Certkiller . The Certkiller network consists of a single Windows 2000 Active Directory domain Certkiller .com. Windows XP is used on all servers and client computers in the network.


Only Cisco routers are used throughout the company. The routers are named Certkiller 1, Certkiller 2, and Certkiller 3.

The network is characterized by:

- \* Certkiller 2 and Certkiller 3 are configured with OSPF.
- \* Certkiller 3's Loopback0 interface is in Area 2.
- \* Certkiller 3's S0/1 interface and Certkiller 2's S0/0 interface are both in Area 0.

The Certkiller network is shown in the following exhibit:

## 642-801



You are required to configure the network as follows:

- \* Certkiller 2's S0/1 interface in Area 1.
- \* Certkiller 1's S0/0 interface in Area 1.
- \* Certkiller 1's S0/0 and Certkiller 2's S0/1, and only these interfaces, should be in Area 1. Use the appropriate mask!
- \* Area 1 should not receive any external routes.
- \* Area 1 should not receive any inter-area routes, except the default route.

Answer:

Explanation:

```
Click on host 2 on Certkiller 1
enab
config t
router ospf 1
network 192.168.1.4 0.0.0.3 area 1
area 1 stub
Click on host 4 Certkiller 2
enab
config t
router ospf 1
network 192.168.1.4 0.0.0.3 area 1
network 192.168.2.4 0.0.0.3 area 0
area 1 stub no-summary
```

Note: No need to configure Certkiller 3

Variation #1

Certkiller 1:

S0/0: 192.168.18.5/30

Certkiller 2:

S0/0: 192.168.40.6/30

S0/1: 192.168.18.6/30


Certkiller 3:

S0/1: 192.168.40.5/30  
Loopback0: 200.200.200.200/32

**QUESTION 251:**

**SIMULATION**

The Certkiller WAN is depicted below:



You work as a network engineer at Certkiller .com. OSPF has been configured on routers Certkiller 1, Certkiller 2, and Certkiller 3. Certkiller 2's S0/0 interface and Certkiller 3's S0/1 interface are in Area 0. Certkiller 3's Loopback0 interface is in Area2.

Your task is to configure the following:

Certkiller 1's S0/0 interface in Area 1.

Certkiller 2's S0/1 interface in Area 1.

Use the appropriate mask such that ONLY Certkiller 1's S0/0 and Certkiller 2's S0/1 could be in Area 1.

Area 1 should not receive any external or inter-area routes (except the default route).

Answer:

Explanation:

Configuration for Certkiller 1

```
Certkiller 1#config
```

```
Certkiller 1(config)#int s0/0
```

```
Certkiller 1(config-if)#ip address 192.168.34.5 255.255.255.252
```

```
Certkiller 1(config-if)#no shutdown
```

```
Certkiller 1(config-if)#exit
```

```
Certkiller 1(config)#router ospf 1
```

```
Certkiller 1(config-router)#network 192.168.34.4 0.0.0.3 area 1
```

```
Certkiller 1(config-router)#area 1 stub
```

```
Certkiller 1(config-router)#^Z
```

Configuration for Certkiller 2

```
Certkiller 2#config
```

```
Certkiller 2(config)#interface serial 0/0
```

```
Certkiller 2(config-if)#ip address 192.168.75.6 255.255.255.252
Certkiller 2(config-if)# no shutdown
Certkiller 2#config
Certkiller 2(config)#interface serial 0/1
Certkiller 2(config-if)#ip address 192.168.34.6 255.255.255.252
Certkiller 2(config-if)# no shutdown
Certkiller 2(config-if)#exit
Certkiller 2(config)#router ospf 1
Certkiller 2(config-router)#network 192.168.75.4 0.0.0.3 area 0
Certkiller 2(config-router)#network 192.168.34.4 0.0.0.3 area 1
Certkiller 2(config-router)#area 1 stub no-summary
Certkiller 2(config-router)#^z
Certkiller 2#copy running-config startup-config
OR for Certkiller 2
Certkiller 2#config terminal
Certkiller 2(config)#router ospf 1
Certkiller 2(config-router)#network 192.168.34.4 0.0.0.3 area 1
Certkiller 2(config-router)#area 1 stub no summary
Certkiller 2(config-router)#network 192.168.75.4 0.0.0.3 area 0
Certkiller 2(config-router)#^z
Certkiller 2#copy running-config startup-config
```

Variation #1:


```
Certkiller 1:
S0/0: 192.168.11.5/30
Secret password: Certkiller
Certkiller 2:
S0/0: 192.168.44.6/30
S0/1: 192.168.11.6/30
Secret password: Certkiller
Certkiller 3:
S0/0: 192.168.44.5/30
Loopback0: 221.221.221.221
Secret password: Certkiller
```

Variation #2:

```
Certkiller 1:
S0/0: 192.168.16.5/30
Secret password: Certkiller
Certkiller 2:
S0/0: 192.168.51.6/30
S0/1: 192.168.16.6/30
Secret password: Certkiller
Certkiller 3:
S0/0: 192.168.51.5/30
Loopback0: 55.55.55.55
Secret password: Certkiller
```

**QUESTION 252:**

The Certkiller OSPF network is depicted below:



Which two statements concerning Certkiller 2 are true about the above configuration? (Choose two.)

- A. The interfaces S0/1 and Fa0/0 will not process any routing update that the router receives.
- B. The interfaces S0/1 and Fa0/0 are configured as passive interfaces.
- C. The S0/0 interface will not process any routing update that it receives.
- D. The S0/0 interface will not sending out routing updates.
- E. Certkiller 3 will hear about the 171.69.232.0 and 172.24.101.0 networks through OSPF update.
- F. Certkiller 1 and Certkiller 4 will hear about the 140.10.0.0 network through OSPF update.


Answer: E, F

Explanation:

This is a bit of a trick question, since the statement "no passive interface default" command will effectively cancel out the previous command, which will result in normal OSPF operation. Due to this, all networks will be advertised normally via the OSPF process.

**QUESTION 253:**

The Certkiller OSPF network is shown below:



Based on the info shown above, which command will summarize area 15?

- A. on the ABR E0/0 interface: area 15 range 10.10.8.0 255.255.248.0
- B. on the ABR S0/0 interface: area 15 range 10.10.8.0 255.255.252.0
- C. on the ABR in the OSPF config-router mode: area 15 range 10.10.12.0 255.255.252.0
- D. on all routers in area 15 in the OSPF config-router mode: area 15 range 10.10.12.0 255.255.248.0
- E. on the ABR S0/0 interface: summary-address 10.10.8.0 255.255.248.0
- F. on the ABR in the OSPF config-router mode: summary-address 10.10.12.0 255.255.252.0

Answer: C


Explanation:

To consolidate and summarize routes at an area boundary, use the area range command in router configuration mode. The area range command is used only with Area Border Routers (ABRs). It is used to consolidate or summarize routes for an area. The result is that a single summary route is advertised to other areas by the ABR. Routing information is condensed at area boundaries. External to the area, a single route is advertised for each address range. Only Choice C specifies the correct syntax and route summarization network mask.

---

**QUESTION 254:**

Exhibit, Network Topology



You work as a network technician at Certkiller .com. Study the exhibits carefully. Which command will Certkiller 2 use to summarize routes for the 192.168.16.0/22 supernet before injecting them into Area 0?

- A. area 10 range 192.168.16.0 255.255.252.0
- B. summary-address 192.168.16.0 255.255.252.0
- C. ip summary-address ospf 101.192.168.16.0 255.255.252.0
- D. area 0 range 192.168.16.0 255.255.252.0
- E. ip summary-address area 0 192.168.16.0 255.255.252.0

Answer: A

Explanation:

The area range command is used only with Area Border Routers (ABRs) which is router Certkiller 2 in this example. It is used to consolidate or summarize routes for an area. The result is that a single summary route is advertised to other areas by the ABR. Routing information is condensed at area boundaries. External to the area, a single route is advertised for each address range. Only Choice A specifies the correct syntax and route summarization network mask.

### QUESTION 255:

In order to configure ISIS to route IP traffic, the following command was entered into router CK1 :

```
ip router isis
```

In which IOS configuration mode was this command typed in?

- A. Line configuration mode
- B. Router configuration mode
- C. Global configuration mode
- D. Interface configuration mode
- E. Root configuration mode

Answer: D

Explanation:

To configure an IS-IS routing process for IP on an interface, use the ip router isis interface configuration command.

Note: To enable IS-IS, perform the following tasks starting in global configuration mode:

Step 1: router isis

Enable IS-IS routing and specify an IS-IS process for IP, which places you in router configuration mode.

Step 2: net network-entity-title

Configure NETs for the routing process; you can specify a name for a NET as well as an address.

Step 3: interface-type number

Enter interface configuration mode.

Step 4: ip router isis [tag]

Specify the interfaces that should be actively routing IS-IS.

Incorrect Answers:

A, B, C: The ip router isis cannot be used in either line, router or Global configuration mode.

E: This mode does not exist.

---

**QUESTION 256:**

The Certkiller IS-IS network has been configured with mesh groups. Which of the following describe an advantage of utilizing mesh groups?

- A. Mesh groups optimize LSP flooding.
- B. Mesh groups keep the routers more secure.
- C. Mesh groups help the routers form adjacencies.
- D. Mesh group speed the flow of data across WAN links.
- E. None of the above.

Answer: A

Explanation: The mesh group feature is a mechanism to reduce flooding of LSPs in nonbroadcast multi-access (NBMA) networks with highly meshed, point to point topologies.

Reference: RFC 2973, IS-IS Mesh Groups

---

**QUESTION 257:**

When configuring IS-IS, why is it important to manually configure the link costs for each of the associated links within the network?

- A. Because by default, all links have a cost of 10 regardless of the bandwidth.
- B. It is not important because the default provides for optimal routing.
- C. Because there is no default link cost. The link cost must be configured on each

interface.

D. Because by default, all LAN interfaces have a cost of 10 and all WAN interfaces have a cost of 50.

Answer: A

Explanation:

The original IS-IS specification defines four different types of metrics: cost, delay, expense, and error. The Cisco implementation uses cost only. All links use the metric of 10 by default.

The original IS-IS specification defines four different types of metrics. Cost, being the default metric, is supported by all routers. Delay, expense, and error are optional metrics. The delay metric measures transit delay, the expense metric measures the monetary cost of link utilization, and the error metric measures the residual error probability associated with a link.

The Cisco implementation uses cost only. If the optional metrics were implemented, there would be a link-state database for each metric and SPF would be run for each link-state database.

IS-IS uses a single required default metric with a maximum path value of 1024. The metric is arbitrary and typically is assigned by a network administrator. Any single link can have a maximum value of 64, and path links are calculated by summing link values. Maximum metric values were set at these levels to provide the granularity to support various link types while at the same time ensuring that the shortest-path algorithm used for route computation will be reasonably efficient. IS-IS also defines three optional metrics (costs): delay, expense, and error.

Reference: Introduction to Intermediate System-to-Intermediate System Protocol, [http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys\\_wp.htm](http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys_wp.htm)

Incorrect Answers:

B: While some routing protocols calculate the link metric automatically based on bandwidth (OSPF) or bandwidth/delay (Enhanced Interior Gateway Routing Protocol [EIGRP]), there is no automatic calculation for IS-IS.

C: The default link cost is 10.

D: There is no such distinction between LAN and WAN interfaces.

---

### **QUESTION 258:**

Which of the following best describe the addresses that are used at the network layer of the OSI model?

- A. Internet Protocol address
- B. Media Access Control address
- C. Packet Layer Protocol address
- D. Network Service Access Point address
- E. Authority and Format Identifier address

Answer: D

Explanation:

The term "network address" is used to refer to the Network Service Access Point (NSAP) at which the OSI Network Service is made available to a Network Service user by the Network Service provider.

Incorrect Answers:

A: IP is a protocol that is based on the OSI model.

B: MAC addresses operate at a lower level in the OSI model.

C, E: Do not apply.

Reference: RFC 941, Addendum to the Network Service Definition Covering Network Layer Addressing

---

**QUESTION 259:**

The OSI IS-IS NSAP address 47.040C.0061.040C.0056.0D12.00 is applied to the Ethernet interface of router CK1 . What is the area ID of the address?


- A. 00
- B. 47
- C. 47.040C
- D. 47.040C.0061
- E. 040C.0056.0D12
- F. None of the above

Answer: D

Explanation:

An NSAP address (figure 7) has two major parts: the initial domain part (IDP) and the domain specific part (DSP) (Figure 7). The IDP consists of a 1-byte authority and format identifier (AFI) and a variable-length initial domain identifier (IDI), and the DSP is a string of digits identifying a particular transport implementation of a specified AFI authority. Everything to the left of the system ID can be thought of as the area address of a network node.

Figure 7 The NSAP address



- An NSAP Consists of 3 Parts: Area Address, System ID, and NSAP Selector
- Total Length Between 8 and 20 Bytes
- Example: 49.0001.0000.0000.0007.00

Incorrect Answers:

A: 00 is the SEL.

B, C: 47 and 47.040C is just a part of the Domain

E: 040C.0056.0D12 is the SystemID.

Reference:

[http://www.cisco.com/en/US/tech/CK3\\_65/CK3\\_81/technologies\\_white\\_paper09186a00800a3e6f.shtml](http://www.cisco.com/en/US/tech/CK3_65/CK3_81/technologies_white_paper09186a00800a3e6f.shtml)

---

**QUESTION 260:**

Which of the answer choices below represents a valid private, locally administered NSAP address?

- A. 39.0f01.0002.0000.0c00.1111.00
- B. 48.0f01.0002.0000.0c00.1111.00
- C. 49.0004.30ac.0000.3090.c7df.00
- D. 52.0f01.0002.0000.0c00.1111.00

Answer: C

Explanation:

The private NSAP addresses have AFI beginning with 49, as these are locally administered addresses..

AESA Network Service Access Point (NSAP) ATM Addresses

There are 3 types of private ATM addresses:

1. NSAP encoding format for E.164 addresses - The authority and format identifier (AFI) is 45. These addresses are used in establishing ISDN calls by public networks, and they are normally used in public telephony.
2. Data Country Code (DCC) AESA - The AFI is 39. These addresses are to be used in public networks. For example, the initial domain identifier (IDI) value 0x84.0f identifies the United States.
3. International Code Designator (ICD) AESA - The AFI is 47. These addresses are used in private organizations, and the ICD field indicates the code set or organization. Cisco uses by default ICD addresses.

Incorrect Answers:

- A: Addresses beginning with 39 is administered by the ISO for ISO Data Country Code.  
B, D: 48 and 52 are unassigned AFI addresses.

Reference:

[http://www.cisco.com/en/US/tech/CK3\\_9/CK4\\_9/technologies\\_tech\\_note09186a00800c9761.shtml](http://www.cisco.com/en/US/tech/CK3_9/CK4_9/technologies_tech_note09186a00800c9761.shtml)

---

**QUESTION 261:**

You are a guest lecturer at the Certkiller Academy and are preparing a lesson on the Cisco IS-IS. Which of the following points can you include in your discussion of the Cisco IS-IS NSAP address System IDs? (Select three)

- A. System IDs can vary in size within a domain.
- B. The System ID identifies a node in an IS-IS network.
- C. The System ID must be unique within a Level-1 area.
- D. The System ID must be unique within a Level-2 area.
- E. The System ID must be the MAC address of the router.

Answer: B, C, D

Explanation:

B: Each system ID within an area must be unique. It is used to identify a IS-IS node.

C: All Level 1 routers and hosts in an area must have an NSAP with the same area address.

D: Level 2 routers advertise their own area addresses (NSAP) to the other Level 2 routers in the backbone.

Incorrect Answers:

A: All ISs and ESs in a routing domain must have system IDs of the same length.

Furthermore, Cisco implements a fixed length of 6 bytes for the system ID.

E: There are several techniques for creating unique system IDs

\* Start numbering 1, 2, 3, 4, and so on.

\* Use Media Access Control (MAC) addresses.

\* Convert and use the loopback IP address: 192.168.11.1 --> 192.168.011.001--> 1921.6801.1001.

Reference: Introduction to Intermediate System-to-Intermediate System Protocol

[http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys\\_wp.htm](http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys_wp.htm)

---

### **QUESTION 262:**

When comparing and contrasting the NET and NSAP used in IS-IS routing, which of the following statements is true?

A. A NET is an NSAP address with the N-selector byte set to 00.

B. Network Entity Titles do not have to start and stop on byte boundaries.

C. The System ID field of the NSAP address does not uniquely identify a node.

D. To identify a Domain, an NSAP address must be used because a NET can only identify an Area.

E. A NET is a special version of an NSAP address restricted to 8 bytes for the Area Address, System ID and the N-Selector byte.

Answer: A

Explanation:

NETs and NSAPs are ISO addresses. The differences between the NET and NSAP addresses are subtle. The NET address is the address of the host, where the value in the NSEL field is set to 0x00. Therefore, there is no upper-layer protocol identified within the host. With no application identified within the end host, the packet can be routed to the destination, but it cannot be handed off to a process after it has been delivered.

However, routers do not have upper-layer protocols to identify because they are transitory ISs. Therefore, the NSAP of the router is referred to as a NET because the NSEL field is set to 00. The NSAP is the full ISO address. It not only defines the area and destination host within the area, but also specifies where to send the incoming packet after it has reached the host. The NSEL field at the end of the ISO address specifies the

upper-layer protocol and is similar to the Protocol field of the IP header.

Network Entity Title (NET) In IS-IS, this is the ISO address of the system, but not to the process destination within the system. The NET describes both the area and system ID of a system in the IS-IS network but excludes the NSEL, which is set to 0x00. If the NSEL identifies the process within the system, the ISO address is called the "NSAP address."

Network Service Access point (NSAP) In IS-IS, this describes a service at the network layer to which the packet is to be directed. The NSAP is the NET address with the NSEL field set to a positive value, a value other than 0x00.

Network entity title (NET)

The NET describes both the area and system ID of a system in the IS-IS network but excludes the NSEL, which defines the NSAP address of the system.

Network service access point (NSAP)

Describes a service at the network layer to which the packet is to be directed. The NSAP is the NET address with the SEL field set to a value other than 0x00.

*An ISO Address*

| IDP              |     | DSP               |                           |                   |
|------------------|-----|-------------------|---------------------------|-------------------|
| AFI<br>(1 octet) | IDI | High Order<br>DSP | System ID<br>(1-8 octets) | NSEL<br>(1 octet) |
| AREA             |     |                   | ID                        | SEL               |

---

**QUESTION 263:**

What are the basic configuration steps needed to enable IS-IS?

- A. Configure the net system-id command under router isis and enable IS-IS on each interface with the ip router isis command.
- B. Configure the network net-id command(s) under router isis and enable IS-IS on each interface with the ip router isis command.
- C. Configure the network net-id command(s) and the is-type level-1-2 command under router isis.
- D. Configure the net system-id and the network net-id commands under router isis.
- E. Configure the net system.-id and the network net-id commands under router isis and enable IS-IS on each interface with the ip router isis command.

Answer: A

Explanation:

In order to enable IS-IS for IP on a Cisco router and have it exchange routing information with other IS-IS enabled routers, you must perform the following two tasks:

1. Enable the IS-IS process and assign area
2. Enable IS-IS for IP routing on an interface

The sample configuration below configures an IS-IS router with the following parameters:

1. Area 49.0001

2. Level 1 (L1) and Level 2 (L2) routers (this is the default unless otherwise specified)
3. No optional parameters
4. Running IS-IS for IP only
5. Loopback interfaces (loopbacks are advertised by IS-IS, not IS-IS enabled)

Router CK1

!

```
interface Loopback0
```

```
ip address 172.16.1.1 255.255.255.255
```

!--- Creates loopback interface and assigns

!--- IP address to interface Loopback0.

!

```
interface Ethernet0
```

```
ip address 172.16.12.1 255.255.255.0
```

```
ip router isis
```

!--- Assigns IP address to interface Ethernet0

!--- and enables IS-IS for IP on the interface.

!

```
router isis
```

```
passive-interface Loopback0
```

```
net 49.0001.1720.1600.1001.00
```

!

!--- Enables the IS-IS process on the router,

!--- makes loopback interface passive

!--- (does not send IS-IS packets on interface),

!--- and assigns area and system ID to router.

Reference:

[http://www.cisco.com/en/US/tech/CK365/technologies\\_configuration\\_example09186a0080093f38.shtml](http://www.cisco.com/en/US/tech/CK365/technologies_configuration_example09186a0080093f38.shtml)

---

### **QUESTION 264:**

An administrator has decided to configure Integrated IS-IS on the network. What must the administrator do to enable the interfaces to distribute IP information using IS-IS?

- A. The networks configured on each interface must be associated to the IS-IS routing protocol using the network router configuration command.
- B. Each interface must be enabled to support IS-IS with the ip router isis interface configuration command.
- C. All configured IP networks will automatically be enabled when IS-IS is configured with the router isis global configuration command.
- D. Each network must be identified with the net interface configuration command.

Answer: B

Explanation:

The preparation for configuring any routing protocol requires a thorough understanding

of the network topology and a coherent addressing scheme. When you have these, the basic configuration of the Integrated IS-IS is as follows:

- 1) Enable the routing process Integrated IS-IS with the router isis command.
- 2) Configure the Network Entity Title (NET) address, thus assigning the area with the net network-address router subcommand.
- 3) Enable Integrated IS-IS for IP on the relevant interfaces with the IP router isis interface subcommand.

Reference: Cisco press 642-801 Study Guide p.397

---

**QUESTION 265:**

You have a router that's been configured with multiple IP routing protocols, and you're interested in checking out if and how inbound and outbound routing updates are being filtered. Which one of the following commands would list these filters?

- A. show ip
- B. show ip route
- C. show ip protocols
- D. show ip interface
- E. show protocol filters

Answer: C

Explanation:

The command show ip protocols command is used to display the parameters and current state of the active routing protocol process. The information is presented on a routing protocol basis and includes applied inbound and outbound filters.

Incorrect Answers:

- A: Show ip is an incomplete command.
  - B: The show ip route command displays active routes, not information on filters.
  - D: The show ip interface command lists a summary of an interface's IP information and status. However, it does not list the filters applied on a routing protocol basis.
  - E: This is an invalid command.
- 

**QUESTION 266:**

You have a low speed serial connection that is being used only as a backup to a frame relay network. Which routing method could you use on your internal routers to minimize the bandwidth wasted for frivolous routing updates?

- A. Use a distance vector routing protocol.
- B. Use private IP addresses.
- C. Use dial-on-demand routing.
- D. Use route summarization.
- E. Use a routing protocol that tolerates route flapping.

Answer: D

Explanation:

An advantage to using route summarization in a large complex network is that it can isolate topology changes from other routers. That is, if a specific link in the domain were flapping (going down and up rapidly), the summary route would not change, so no router external to the domain would need to keep modifying its routing table due to this flapping activity.

Reference: Building Scalable Cisco Networks (Cisco Press) page 76

---

**QUESTION 267:**

You want to decrease the amount of routing query traffic across your EIGRP network. Which of the following is an effective means of controlling EIGRP query traffic?

- A. Route summarization.
- B. Configuring route filters.
- C. Using a hierarchical addressing scheme.
- D. Establishing separate autonomous systems.
- E. None of the above.

Answer: A

Explanation:

The best solution to control EIGRP queries is to reduce the range of queries. The most effective method to restrict the range of queries is the establishment of route summarization boundaries. Since route summarization decreases the total number of routing table entries, fewer routing queries are needed.

---

**QUESTION 268:**

DRAG DROP

You are a systems administrator of one of the Certkiller EIGRP networks. Your goal is to suppress Certkiller 's router from sending updates out of its Serial 0/0 interface without compromising its ability to form and maintain neighbor adjacencies on the same interface.

The router has already been configured as follows:


```
Router(config)#access-list 20 deny any
Router(config)#access-list 100 permit ip any any
Router(config)#router eigrp 1
```

Drag the correct answers on the left to the correct command line space on the right, to finish configuring the router.

## 642-801



Answer:



Explanation:

We do not want to allow the routing updates along serial 0/0 to go out. This distribute list will deny all IP networks from being advertised out the interface.

---

### QUESTION 269:

Your OSPF router has one serial interface, and one Ethernet LAN interface. The sub-interface is configured in the following manner:

```
interface serial 0.122 point-to-point
ip address 192.168.1.1 255.255.255.0
encapsulation frame-relay
frame-relay interface-dlci 122
```

You want to allow hosts on your LAN to send and receive data, but you don't want routing traffic to go through that interface. Which of the following commands should you use to complete your configuration?

- A. interface serial 0.122 point-to-point  
passive-interface ethernet 0
- B. interface ethernet 0  
ip address 192.168.12.1 255.255.255.0  
passive-interface
- C. router ospf 172  
area 1 nssa  
network 192.168.1.0 0.0.0.255 area 0  
network 192.168.12.0 0.0.0.255 area 1
- D. router ospf 172  
passive-interface ethernet 0  
network 192.168.1.0 0.0.0.255 area 0  
network 192.168.12.0 0.0.0.255 area 1

Answer: D

Explanation:

We use the passive-interface command to configure the ethernet interface to be passive. The passive-interface router configuration command is used to disable sending routing updates on an interface.

Incorrect Answers:

A: We are not configuring the serial interface. Furthermore, the passive-interface command is a router configuration command, not an interface configuration command.

B: We cannot use the passive-interface command like this.

C: We should configure the Ethernet interface as passive, not the area as a not-so-stubby area (NSSA).

---

**QUESTION 270:**

**DRAG DROP**

You have a named Router CertK running EIGRP that has already been configured with the following:

```
Router CertK (config)#access-list 30 deny any
```

```
Router CertK (config)#access-list 40 permit ip any
```

```
Router CertK (config)#router eigrp 1
```

To complete your configuration you want to configure the serial 0/0 interface to prevent routing updates from going out that interface while still allowing the formation and upkeep of neighbor adjacencies on that interface. Drag the correct command phrase from the bottom and place it in the right command line sequence above.

```
RouterCertK(config-router)#   
```

Select from these

|            |                    |                    |     |
|------------|--------------------|--------------------|-----|
| in         | distribute-list 40 | distribute-list 30 | out |
| serial 0/0 | interface          | passive interface  |     |

Answer:

```
RouterCertK(config-router)#   
```

Select from these

|           |                    |
|-----------|--------------------|
| in        | distribute-list 40 |
| interface | passive interface  |

Explanation:

We use distribute list 30, which stops IP traffic, on outward traffic on the serial 0/0 interface.

RouterCK continues receiving routing updates from its neighbor, but the distribute-list prevents routes from being advertised out of serial 0. Furthermore, neighbor adjacencies are allowed to be formed between RouterCK and its neighbor on serial 0/0.

Incorrect Answers:

passive interface serial 0/0:

On EIGRP, passive interface causes the router to stop sending and receiving hello packets. This will prevent the interface from maintaining neighbor adjacencies.

distribute-list 40 out serial 0/0:

We must stop outgoing traffic, not allow it.

distribute-list 30 in serial 0-0:

We must stop outgoing traffic, not incoming traffic.


Reference: How Does the Passive Interface Feature Work in EIGRP?

<http://www.cisco.com/warp/public/103/16.html>

---

### QUESTION 271:

The Certkiller network is displayed below:



172.27.6.0/24

You need to configure a static route so that users on the corporate network can reach the 172.27.6.0/24 LAN. Based on the diagram above, which of the following is the correct way to do this?

- A. Certkiller 1(config)#ip route 172.27.6.0 255.255.255.0 172.17.8.2
- B. Certkiller 2(config)#ip route 172.27.6.0 255.255.0.0 172.17.8.2
- C. Certkiller 2(config)#ip route 172.27.6.0 255.255.255.0 172.17.8.2
- D. Certkiller 1(config)#ip route 172.27.6.0 255.255.0.0 172.17.8.1
- E. Certkiller 2(config)#ip route 172.27.6.0 255.255.255.0 172.17.8.1

Answer: E

Explanation:

Only choice E provides the correct network subnet mask and next hop IP address.

IP route prefix mask {address|Interface} [distance] [tag tag] [permanent]

Prefix 172.27.6.0 mask 255.255.255.0 address 172.17.8.2

Address - The IP address of the next hop router that can be used to reach that network.

Reference: Building Scalable Cisco Networks (Cisco Press) page 464

---

### QUESTION 272:

The Certkiller network consists of a main office and a single remote office. You need to configure the remote office in using the following guidelines:

1. The main office must learn all of its routes from the regional office.
2. The regional office must not learn routes from the main office.
3. The most scalable solution should be used.

Which of the following answer choices best describe what should be done to satisfy these requirements?

- A. Configure static routes pointing to the network behind the central office router
- B. Configure a default route pointing to the networks behind the central office router
- C. Make the interface that is connected to the central office a passive interface to block incoming updates
- D. Enable route update filtering on the interface that is connected to the central office to block incoming updates
- E. None of the above will work

Answer: D

Explanation:

The only viable solution here is to enable the main office to dynamically learn about all of the regional office routes, while still not receiving any routes is by enabling incoming route filters.

Incorrect Answers:

A, B: Configuring static routes or default routes on the regional office will not enable the main office to learn about any of the regional office routes.

C: A passive interface would prevent all route updates to the main office. Passive interfaces do not filter incoming routing updates, so the regional office could still learn routes from the main office.

Note: Configuring the main office using a passive interface would satisfy the conditions here, but this question asks us how to configure the regional office, not the main office.

---

### **QUESTION 273:**

If a router already has a route to the 172.27.0.0/16 network in the routing table, which of the following commands would cause EIGRP to flag that network as the candidate default route?

- A. ip default-network 172.27.0.0
- B. ip route 0.0.0.0 0.0.0.0 172.27.0.0
- C. ip default-network 0.0.0.0
- D. default-information originate
- E. ip classless

Answer: A

Explanation:

Unlike the ip default-gateway command, you can use ip default-network when ip routing is enabled on the Cisco router. When you configure ip default-network the router considers routes to that network for installation as the gateway of last resort on the router.


Reference:

[http://www.cisco.com/en/US/tech/ CK3 65/ CK5 54/technologies\\_tech\\_note09186a0080094374.shtml](http://www.cisco.com/en/US/tech/ CK3 65/ CK5 54/technologies_tech_note09186a0080094374.shtml)

---

**QUESTION 274:**

Two Certkiller routers are connected together as shown in the diagram below.



The networks in the above diagram are all configured with RIP. On router CK1 , the following configuration change was made:

```
CK1 (config)#ip default-network Network CK3
```

Which router will receive a default route as a result of this change?

- A. Router CK1 only-
- B. Router CK2 only.
- C. Both routers.
- D. Neither router.

Answer: C

Explanation:

The ip default-network command is used as a method of distributing default route information to other routers. When running RIP, you can create the default route by using the ip default-network command. If the router has a directly connected interface onto the network specified in the ip default-network command, RIP will generate (or source) a default route to its RIP neighbor routers.

Unlike the ip default-gateway command, you can use ip default-network when ip routing is enabled on the Cisco router. When you configure ip default-network the router considers routes to that network for installation as the gateway of last resort on the router. In this example, it will create a default network for CK1 , since it knows about the route to CK3 , and CK1 will advertise this default route to CK2 .

Reference: <http://www.cisco.com/warp/public/105/default.html#ipnetwork>

---

**QUESTION 275:**

Router CK1 is configured as follows:

```
interface serial 0
ip address 185.64.1.1 255.255.255.0
!
interface ethernet 0
ip address 15.10.10.1 255.255.255.0
!
```

```
router rip
network 15.10.10.0
!
ip route 0.0.0.0 0.0.0.0 serial0
```

What result would the above commands accomplish?

- A. RIP updates are sent and received on interface serial0 of the router.
- B. A default route is sent to neighbors on interface serial0 of the router.
- C. A default route is sent to neighbors on interface ethernet0 of the router.
- D. RIP updates are sent and received on interfaces serial0 and ethernet0 of the router.
- E. None of the above.

Answer: E

Explanation:

Since RIP is configured for only the network on the Ethernet segment, RIP updates will only be sent and received on this network, not on the serial interface. In addition, static routes are not sent via RIP, unless the "redistribute static" command is configured. This includes default static routes.

Note: Had the configuration command "redistribute static" been placed under the RIP routing process, then choice C would have been correct.

---

### **QUESTION 276:**

Which of the following commands ensures that permanently created static route entries are dynamically added into the routing process?

- A. inject static
- B. inject permanent
- C. redistribute all
- D. redistribute static

Answer: D

Explanation:

If you define a static route to an interface that is not one of the networks defined in a network command, no dynamic routing protocols will advertise the route unless a redistribute static command is specified for these protocols.

Incorrect Answers:

A: The inject static is used to configure legacy DECnet systems. It is not a valid IP routing command.

B: There is no such command.

C: The command is not valid.

Reference: Cisco, Configuring IP Routing Protocol-Independent Features

---

**QUESTION 277:**

Which IOS command below would be useful if you had to create a static list of the other routers in your NBMA cloud?

- A. network
- B. neighbor
- C. ip route
- D. router ospf

Answer: B

Explanation:

The neighbor command is used to statically map the IP address of the neighbor. This is useful in multi-point networks where the neighbors may not be able to be found dynamically due to the nature of the NBMA network.

---

**QUESTION 278:**

The Certkiller router had the following configuration line added:

```
Certkiller (config)#ip route 172.27.6.0 255.255.255.0 s0/0
```

What information can be gathered from this configuration change? (Select two)

- A. This is a route to a public network.
- B. There is only one path to this network from Certkiller .
- C. This is a route to interface s0/0 on the next hop router.
- D. Packets destined for this network are sent via interface s0/0 on Certkiller .
- E. Packets destined for this network enter router Certkiller through interface s0/0.

Answer: B, D

This command places a static route to the 172.27.6.0 network. All traffic destined to this network from this router will be sent across the s0/0 interface.

Incorrect Answers:

A: The 172.27.6.0 IP address range is part of the private IP address space as specified by RFC 1918.

C: The serial 0/0 indicates the interface to use going out, not the next hop.

E: Packets destined for this network will exit through serial 0/0, not enter through it.

Reference: Building Scalable Cisco Networks (Cisco Press) page 464

---

**QUESTION 279:**

The Certkiller 2 router is running RIP version 2 and is configured as shown below:

```
Certkiller 2(config)#router rip
```

```
Certkiller 2(config-router)#version 2
```

```
Certkiller 2(config-router)#no auto-summary
```

Based on this information, which of the following is true?

- A. Subnets are summarized at the network boundary.
- B. Subnets are advertised across network boundaries.
- C. Subnet mask information is not passed in the routing updates.
- D. Subnets are made discontinuous.
- E. None of the above

Answer: B

Explanation:

By default, RIP version 2 summarizes subnets across network boundaries. To restore the default behavior of automatic summarization of subnet routes into network-level routes, use the auto-summary router configuration command. To disable this feature and transmit subnet routing information across classful network boundaries, use the no form of this command.

Reference:

[http://www.cisco.com/en/US/products/sw/iosswrel/ps1826/products\\_command\\_summary\\_chapter09186a00800d](http://www.cisco.com/en/US/products/sw/iosswrel/ps1826/products_command_summary_chapter09186a00800d)

---

**QUESTION 280:**

What would you find in an EIGRP routing table if route summarization is configured on the routers Serial0 interface, and it summarizes routes learned from its Ethernet0 interface?

- A. A summary route pointing to the Null0 interface.
- B. A summary route pointing to the Serial0 interface.
- C. A summary route pointing to the Ethernet0 interface.
- D. A summary route pointing to the Loopback0 interface.

Answer: A

Explanation:

With EIGRP, when summarization is configured on any interface, the IOS immediately creates and installs the summarized route pointing to null 0 in the routing table.

---

**QUESTION 281:**

With regard to controlling routing updates with the use of route filtering, which of the following statement is true?

- A. Only inbound routes can be filtered.
- B. Only outbound routes can be filtered.
- C. Routes to be filtered are selected by using distribute lists.
- D. Routes to be filtered are selected using only extended access lists.

E. Routes to be filtered are selected by using the distribute-group command.

Answer: C

Explanation:

The syntax for the distribute-list in/out command is:

```
distribute-list access-list-number in/out [interface-name]
```

where access-list-number is the standard IP access-list against which the contents of the incoming or outgoing routing update are matched. The [interface-name] argument is optional and specifies the interface on which the update is expected. It is important to note that the access-list referred to in access-list-number is applied to the contents of the update, not to the source or destination of the routing update packets. The router decides whether or not to include the contents in its routing table based on the access-lists. For example:

```
access-list 1 permit 1.0.0.0 0.255.255.255router ripdistribute-list 1 in!--- The distribute-list command is given!--- under the router configuration mode.
```

Incorrect Answers:

A, B: Distribute lists can be used to filter traffic based on incoming as well as outgoing routes.

D: Standard access-lists are used with distribute lists, not extended access lists.

E: This is an invalid command. The correct syntax is "distribute-list"

---

### **QUESTION 282:**

Router CK1 has been configured to filter routes. Which of the following are reasons to control routing updates via route filtering? (Choose three)

- A. To hide certain networks from the rest of the organization.
- B. For easier implementation.
- C. To control network overhead on the wire.
- D. For simple security.
- E. To prevent adjacencies from forming.

Answer: A, C, D

Explanation:

Route filtering is used to filter out routing updates from other parts of the network, making certain networks unreachable from other parts of the organization. This can be done to simply decrease the amount of traffic overhead associated with routing updates over parts of the network, or for security reasons.


Incorrect Answers:

B: Route filtering requires a more complex router configuration, and can also lead to difficult troubleshooting.

E: Route filtering will only block certain routes from incoming or outgoing, but the routing protocol adjacencies will still be formed between neighboring routers.

**QUESTION 283:**

You are the network engineer at Certkiller . The router topology for the Certkiller network is shown in the following graphic:



Router RTA is configured as follows:

```

RTA(config)#router rip
RTA(config-router)#network 10.0.0.0
RTA(config-router)#distribute-list 44 in interface BRI0
RTA(config-router)#exit
RTA(config)#access-list 44 deny 172.16.1.0 0.0.0.255
RTA(config)#access-list 44 permit any
  
```

What are the effects of this RIP configuration on router RTA? (Choose two)

- A. No routing updates will be sent from interface BRI0 on router RTA to router RTX.
- B. Router RTA will not advertise the 10.0.0.0 network to router RTX.
- C. The route network 172.16.1.0 will not be entered into the routing table on router RTA.
- D. User traffic from the 172.16.1.0 network is denied by access-list 44.
- E. The routing table on router RTA will be updated with the route to router RTW.

Answer: C, E

Explanation:

Answer C is correct as the distribute list calls ACL 4 which denies network 172.16.1.0.

Answer E is correct because RTW networks are not being blocked.

Incorrect Answers:

A, B: Only incoming routes are filtered, not outgoing routes. All known routing entries will be advertised from router A.

D: This is incorrect because this ACL is being called by a distribute-list which only affects route table updates, not user traffic.

**QUESTION 284:**

Router CK1 is configured as follows:

```
interface serial 0
ip policy route-map force
ip route-cache policy
route-map force permit 10
match ip address 1
set ip next-hop 172.20.16.5
access-list 1 permit 10.0.0.5 0.0.0.0
```

Based on the information above, which two of the following statements correctly describe what happens when a packet is received on serial 0 of CK1 ? (Choose two)

- A. If a packet is destined to 10.0.0.5, it is a candidate for fast-switch policy routing.
- B. If a packet was sourced from 10.0.0.5, it is a candidate for fast-switch policy routing.
- C. The route-map policy is incomplete because it does not account for traffic that does not match 10.0.0.5.
- D. If a packet was sourced from 10.0.0.5, it will be sent to 172.20.16.5, regardless of the desired destination IP address.

Answer: B, D

Explanation:

Policy routing is used to override the normal routing behavior of a router. In this example, the router is configured to send all traffic that matches access-list 1 to the router with the next hop IP address of 172.20.16.5.

"ip policy route-map map-tag"

map-tag is the name of the route map to use for policy routing. This must match a map tag

specified by a route-map command.

- When policy routing is configured, turn on the fast switching with this interface command:  
`ip route-cache policy`
- Fast-switched policy routing supports all of the **match** commands and most of the **set** commands, except for the following restrictions:
  - The **set ip default** command is not supported.
  - The **set interface** command is supported only over point-to-point links, unless a route-cache entry exists using the same interface specified in the **set interface** command in the route map. In addition, when process switching, the routing table is consulted to determine a path to the destination. During fast switching, the software does not make this check because fast switching is a cache of the process switch lookup. Instead, if the packet matches, the software blindly forwards the packet to the specified interface. This is a similar situation to the one described in reference to load balancing earlier.

---

**QUESTION 285:**

Which external EIGRP administrative distance should be used in order to allow external EIGRP routes to be more preferred over IS-IS routes but less preferred over internal EIGRP routes?

- A. 89
- B. 114
- C. 119
- D. 121
- E. 131
- F. 141

Answer: B

Explanation:

To meet the conditions of this task the Administrative Distance should be more than 90 and less than 115. B is our only valid option.

The following table displays the default AD for all routing protocols:

| Default Administrative Distances |                  |
|----------------------------------|------------------|
| Route Source                     | Default Distance |
| Connected interface              | 0                |
| Static route                     | 1                |

|                             |     |
|-----------------------------|-----|
| Enhanced IGRP summary route | 5   |
| External BGP                | 20  |
| Internal Enhanced IGRP      | 90  |
| IGRP                        | 100 |
| OSPF                        | 110 |
| IS-IS                       | 115 |
| RIP                         | 120 |
| EGP                         | 140 |
| EIGRP external route        | 170 |
| Internal BGP                | 200 |
| Unknown                     | 255 |

---

**QUESTION 286:**

With the exception of EIGRP and BGP, which type of administrative distance is changed by the distance router configuration command?

- A. External
- B. Internal
- C. Local
- D. Default

Answer: B

Explanation:

To ensure that the optimal path is chosen, it is sometimes necessary to change the administrative distance to make it less favorable. The command structure is protocol-dependent, in that EIGRP requires a separate command. The following command syntax is used for EIGRP:


Router(config)#distance eigrp internal-distance external-distance

Internal-distance is administrative distance for EIGRP internal router. These are routes learned from another entity within the same autonomous system.

---

**QUESTION 287:**

The Certkiller WAN is depicted below:



Exhibit

```
hostname Certkiller1
!
router rip
version 2
network 10.0.0.0
distribute-list 10 in Serial10
```

```
access-list 10 permit 10.0.1.0.255.255
```

Based on the information above, what is the effect of the distribute list command in the Certkiller 1 configuration?

- A. Certkiller 1 will filter only the 172.24.1.0/24 route from the Certkiller 2 RIP updates.
- B. Certkiller 1 will permit only the 10.0.0.0/24 route in the Certkiller 2 RIP updates.
- C. Certkiller 1 will filter 10.1.0.0/24 and the 172.24.1.0/24 routes from the Certkiller 2 RIP updates.
- D. Certkiller 1 will not filter any routes because there is no exact prefix match.

Answer: C


Explanation:

The distribute list is used to define the routing updates to be permitted through the interface. In the example above, the filter list specifies that only the 10.0.1.0/16 networks will be permitted via incoming RIP updates through interface serial 0. All other incoming RIP updates, including the 10.1.0.0/24 and 172.24.1.0/24 will be filtered. Since the 10.0.1.0/24 is part of the 10.0.1.0/16 network, it will be allowed.

---

**QUESTION 288:**

Five Certkiller routers are connected as shown in the following exhibit:



Refer to the exhibit. Assume all routers have been preconfigured to advertise their connected networks in the indicated IGP. Which three configurations would correctly propagate the RTA E0 LAN segment to RTC, but not to RTD or RTE? (Choose three)

- A. RTB(config)# router rip  
RTB(config-router)# version 1  
RTB(config-router)# passive-interface default  
RTB(config-router)# no passive-interface s0/0  
RTB(config-router)# no passive-interface Fa0/0
- B. RTB(config)# router rip  
RTB(config-router)# version 2  
RTB(config-router)# passive-interface s1/0  
RTB(config-router)# passive-interface s1/1
- C. RTB(config)# router igrp 123  
RTB(config-router)# passive-interface default  
RTB(config-router)# no passive-interface s0/1
- D. RTB(config)# router eigrp 123  
RTB(config-router)# passive-interface default  
RTB(config-router)# no passive-interface s0/0  
RTB(config-router)# no passive-interface s1/1
- E. RTB(config)# router eigrp 123  
RTB(config-router)# passive-interface default  
RTB(config-router)# no passive-interface s0/0  
RTB(config-router)# no passive-interface s0/1
- F. RTB(config)# router ospf 1  
RTB(config-router)# passive-interface s0/0  
RTB(config-router)# passive-interface s0/1

Answer: B, C, E

Explanation:

B: In this answer choice, the connections to RTD and RTE are specified as passive, so no updates will be sent to these routers from router B. The link to RTC, however, is not listed as passive so the routes will be propagated to it.

C, D: In this answer choice, the passive-interface default statement is configured, meaning that any links not specified individually as not being passive will be. Since both

of these choices leave the links to RTD and RTE as passive, and RTB explicitly specifies the link to RTC as not being passive, the route will propagate to RTC but not RTD and RTE.

---

**QUESTION 289:**

A Certkiller router was configured with BGP as shown below:

```
router bgp 6500
redistribute static
ip route 164.20.0.0 255.255.0.0 null 0
```

Based on this information, is this a recommended router configuration?

- A. Yes. It allows BGP to advertise the 164.20.0.0 /16 network.
- B. Yes. It results in all traffic for all subnets of 172.16.0.0 being dropped at this router.
- C. No. Cisco prefers that you use the aggregate-address command to distribute IGP routes into BGP.
- D. Yes. Cisco prefers this method of redistributing IGP routes into BGP over using the network command.

Answer: C

Explanation:

Redistribution of static routes configured to the null 0 interface into BGP is done to advertise aggregate routes rather than specific routes from the IP table. However, Cisco recommends the use of the aggregate-address-command instead.

---

**QUESTION 290:**

You are a network analyst at Certkiller and are currently in the process of analyzing two autonomous systems, each running a different routing protocol, but connected between each other with redundant paths. The junior administrator wants to prevent routing loops between the two autonomous systems and asks you if there's an IOS feature that could help him. How would you answer?

- A. Route filtering.
- B. Passive interfaces.
- C. Static redistribution.
- D. Two-way redistribution.

Answer: A

Explanation:

Multiple autonomous systems or routing domains can share route information through the redistribution process. Proper implementation of redistribution requires route filters to prevent feedback loops from forming. It is strongly recommended that redistribution between multiple ASs and multiple routing protocols be accompanied by route filters.

Reference: CCNP #640-503 Building Scalable Cisco Networks (Cisco Press), More EIGRP Scalability Rules

---

**QUESTION 291:**

Routing loops are becoming a problem on a particular network. What can an administrator do to reduce the occurrence of routing loops on a redistributed network?

- A. Use multiple default gateways for redundancy.
- B. Use one-way redistribution for greater stability.
- C. Use two-way redistribution for greater stability.
- D. Use overlapped routing protocols for redundancy.
- E. None of the above.

Answer: B

Explanation:

One-way redistribution would help avoiding the routing loops problem.

Incorrect Answers:

- A: Multiple default gateways would increase the risk of routing loops.
  - C: Two-way redistribution would increase the risk of routing loops.
  - D: Using several routing protocols would increase the risk of routing loops.
- 

**QUESTION 292:**

You are in the process of redistributing EIGRP into another routing protocol on your network. Which of the following commands would you execute if your intention was to use altered administrative distance (AD) parameters for redistributing EIGRP into the second routing protocol?

- A. default-metric eigrp metric
- B. distance eigrp administrative-weight
- C. distance eigrp internal-distance external-distance
- D. distance eigrp external-distance internal-distance

Answer: C

Explanation:

The distance eigrp command is used to allow the use of two administrative distances---internal and external---that could be a better route to a node.

Syntax: distance eigrp internal-distance external-distance

Note: Use the distance eigrp command if another protocol is known to be able to provide a better route to a node than was actually learned via external Enhanced IGRP or if some internal routes should really be preferred by Enhanced IGRP.

Incorrect Answers:

A: Incorrect usage of the default-metric command. Furthermore, this command would not be of use here.

B, D: This is the incorrect use of the distance eigrp command.

---

**QUESTION 293:**

Which three of the following IOS commands could an administrator use to verify route redistribution? (Select three)

- A. debug
- B. traceroute
- C. show summary
- D. show ip route
- E. ipconfig

Answer: A, B, D

Explanation:

A: The debug command can be used to debug redistribution.

B: We can verify connectivity, and the presence of a route, with the traceroute command.

D: We can verify that the routes have been redistributed with the show ip route command. The routes will be shown.

Incorrect Answers

C: The show summary command displays a summary of relationships among owners, content rules, and services. It is not of help in this scenario.

E: This is a command used by PC hosts to verify that the IP stack is working properly.

---

**QUESTION 294:**

A Certkiller OSPF router has been configured in the following manner:

```
router ospf 1
```

```
redistribute eigrp 1 metric 33 subnets
```

Based on this information, what function does the 33 parameter in the redistribute command serve?

- A. It specifies the metric cost to be applied to the redistributed routes.
- B. It specifies the administrative distance on the redistributed routes.
- C. It specifies the metric limit to 33 subnets in each OSPF route advertisement.
- D. It specifies the process-id for the pseudo process that injects the EIGRP routes into OSPF.

Answer: A

Explanation: It specifies the metric cost to be applied to the redistributed routes. In this example, a metric of 33 will be applied to all of the EIGRP routes when redistributed into OSPF. The "subnets" keyword specifies that subnet mask

information is to be preserved during this process.

Reference: Redistributing Routing Protocols

<http://www.cisco.com/warp/public/105/redist.html#examples>

---

**QUESTION 295:**

If you had a static route configured on a router, and that static route had to be advertised to other routers in the network; which one of the following statements would be true?

- A. The router automatically advertises static routes to RIP routers.
- B. You should configure redistribution using the redistribute command.
- C. You should enable static advertisements using the static routes advertise command.
- D. You should include the static route in a distribution list using the distribute-list command.

Answer: B.

Explanation:

If you want a router to advertise a static route in a routing protocol, you will need to redistribute it into a dynamic routing protocol.

Incorrect Answers:

A: Static routes are not automatically redistributed when they are configured. Static routes must be redistributed manually.

C: There is no static route advertise command.

D: This is used for filtering inbound and outbound routes.

Reference: Catherine Paquet and Diane Teare, "Building Scalable Cisco Networks" (Cisco Press 2001), p 465.

---

**QUESTION 296:**

On the subject of redistributing EIGRP between other routing protocols, which three of statements below are correct? (Select three)

- A. IPX RIP redistribution with IPX EIGRP is enabled by default.
- B. AppleTalk EIGRP and RTMP redistribution is enabled by default.
- C. EIGRP for IPX automatically redistributes route information with Novell RIP.
- D. Redistribution between EIGRP and IGRP is always enabled by default regardless of the AS number used.

Answer: A, B, C

Explanation:

A: IPX RIP redistribution with Enhanced IGRP is enabled by default.

B: AppleTalk RTMP redistribution is enabled by default.

C: EIGRP automatically redistributes route information with Novell RIP.

Note: Using a routing protocol to advertise routes that are learned by some other means, such as by another routing protocol, static routes, or directly connected routes, is called redistribution.

Incorrect Answers:

D: Redistribution between EIGRP and IGRP in the same autonomous system is automatically done. Manual configuration is required if different autonomous systems are used.

---

**QUESTION 297:**

**DRAG DROP**

You are the network administrator of an antiquated RIP network that's become too inefficient for its size, so you want to make the transition to OSPF. Your goal is to redistribute RIP routes to OSPF. Your metric will be 80, and you are using the private IP network of 10.0.0.0.

Choose the correct command phrases on the left and drag them onto the correct command sequence on the left.

|                                  | Place here                              | Select from these                                     |
|----------------------------------|---|---|
| RouterCertkiller(config)#        | <input type="text" value="Place here"/> | <input type="text" value="router ospf 13"/>           |
| RouterCertkiller(config-router)# | <input type="text" value="Place here"/> | <input type="text" value="router rip"/>               |
| RouterCertkiller(config-router)# | <input type="text" value="Place here"/> | <input type="text" value="redistribute ospf 13"/>     |
|                                  |   | <input type="text" value="redistribute rip"/>         |
|                                  |   | <input type="text" value="redistribute rip network"/> |
|                                  |   | <input type="text" value="redistribute rip subsets"/> |
|                                  |   | <input type="text" value="metric 80"/>                |
|                                  |   | <input type="text" value="default-metric 80"/>        |

Answer:

|                                  | Place here  | Select from these                                     |
|----------------------------------|---|---|
| RouterCertkiller(config)#        | <input type="text" value="router ospf 13"/>           |   |
| RouterCertkiller(config-router)# | <input type="text" value="redistribute rip subnets"/> | <input type="text" value="router rip"/>               |
| RouterCertkiller(config-router)# | <input type="text" value="default-metric 80"/>        | <input type="text" value="redistribute ospf 13"/>     |
|                                  |   | <input type="text" value="redistribute rip"/>         |
|                                  |   | <input type="text" value="redistribute rip network"/> |
|                                  |   | <input type="text" value="metric 80"/>                |

Explanation:

Step 1: router ospf 13

We are distributing into ospf. 13 denotes the AS (autonomous system) in use (which we only can assume be 13).

Step 2: redistribute rip subnets

The RIP routes are redistributed into OSPF. The subnets keyword tells OSPF to redistribute all subnet routes. Without the subnets keyword, only networks that are not subnetted will be redistributed by OSPF.

Step 3: default-metric 80


Finally the default metric is set.

Reference: RIP and OSPF Redistribution

<http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/cs001.htm>

**QUESTION 298:**

The Certkiller network is displayed in the diagram below:



Assuming that CK2 is configured for mutual redistribution; which of the routes below should be present on Router CK3 's routing table? (Select four)

- A. 172.16.211.0/24
- B. 172.16.213.0/24
- C. 172.16.214.0/24

- D. 172.16.215.0/24
- E. 172.16.210.64/26
- F. 172.16.212.16/28

Answer: A, B, C, D

Explanation:

- A: This is a classful route.
- B: This route is redistributed from OSPF into IGRP.
- C: This route is learned through IGRP.
- D: This is a directly connected route.

Incorrect Answers:

E, F: These are subnetted routes. IGRP is a classful routing protocol. These routes will not be redistributed into IGRP.

Reference: <http://www.cisco.com/warp/public/105/52.html>

Redistributing Between Classful and Classless Protocols: EIGRP or OSPF into RIP or IGRP

---

**QUESTION 299:**

In order to enable a RIP network to communicate with an OSPF network, redistribution is configured on the Certkiller network. What kind of router would be used to redistribute RIP into an OSPF network?

- A. ABR
- B. ASBR
- C. Internal router
- D. Backbone router

Answer: B

Explanation:

External route summarization is specific to external routes that are injected into OSPF via redistribution. Only ASBRs can summarize external routes. These types of routes cannot be summarized by any other router type.

---

**QUESTION 300:**

One of the Certkiller routers is configured to redistribute the interior IP networks into BGP. What are two potential consequences that can occur when dynamically learned routes from IGP protocols get redistributed into BGP? (Select two)

- A. Routing loops can occur.
- B. The IGP routing table is reduced.
- C. External IGP learned routes might not necessarily have originated in this AS.
- D. Route processing is done using process switching instead of cache switching.

Answer: A, C

Explanation:

It is fairly common to redistribute IGP routes (such as Enhanced IGRP, IGRP, IS-IS, OSPF, and RIP routes) into BGP. However, precautions should be made when doing so. Some of your IGP routes might have been learned from BGP (C), so you need to use access lists to prevent the redistribution of routes back into BGP, or else routing loops can occur (A). In addition, when these IGRP routes are redistributed, all of the networks in the IGP routing table will be redistributed, including those that were learned externally.

Reference: Using the Border Gateway Protocol for Interdomain Routing  
<http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/icsbgp4.htm>

---

**QUESTION 301:**

What's the preferred method of route redistribution, when two routes of different protocols get exchanged? (Select two)

- A. Use one way route redistribution when there is one path.
- B. Use one way route distribution when there are multiple paths.
- C. Use static routes when there are multiple paths.
- D. Use two way route distribution when there is one path.
- E. Use two way route redistribution where there are multiple paths.
- F. Use static routes when there is one path.

Answer: B, C

Explanation:

B. One way redistribution- To avoid routing loops and problems with varying convergence times, allow routes to be exchanged in only one direction, not both directions. In the other direction, you should consider a default route.

C. When you want to prevent routing loops - Many companies have large enough networks that redundant paths are prominent. In some cases, for example, when a path to the same destination is learned from two different routing protocols, you may want to filter the propagation of one of the paths.

Reference: Building Scalable Cisco Networks (Cisco Press) page 472

---

**QUESTION 302:**

You are the administrator of a network and are planning on configuring route redistribution. The network you plan on working on is running different routing protocols, and the routers are connected together with redundant links. If your goal was to allow traffic to take the best route, what would you do first?

- A. Define the default metric.
- B. Identify the boundary router(s).
- C. Determine which routing protocol is running in the core router(s).
- D. Determine which routing protocol is running in the edge router(s).
- E. None of the above

Answer: A

Explanation:

Because different protocols use different metrics for finding the best path to a destination, the first step in configuring mutual redistribution is to define the default metric that will be used as a baseline standard.

---

**QUESTION 303:**

You have a network with multiple routing protocols running in different Autonomous Systems, redistributed together with two-way redistribution. What could you do to prevent routing loops?

- A. Manually configuring the static routes.
- B. Manually configuring the default gateway.
- C. Manually configuring the administrative k-value.
- D. Manually configuring the administrative distance.

Answer: D

Explanation:

If you manually configure the administrative distance, you can easily decide which protocol the router will choose for each route, thus preventing the possibility of routing loops.

Incorrect Answers:

A: Manually configuring static routes could cause additional routing issues, since static routes are preferred over the dynamically learned routes they could override the information provided by the redistribution.

B: Manual routing configurations should be avoided when used with dynamically learned routes that are being redistributed.

C: These values are used only by IGRP and EIGRP, so they may not apply at all when redistributed into other routing protocols.

---

**QUESTION 304:**

Router CK1 is configured as shown below:

```
router igrp 100
network 197.135.20.0
network 197.135.24.0
network 197.135.27.0
```

```
redistribute rip
default-metric 10 100 255 1 1500
distance 140 0.0.0.0 255.255.255.255 9
access-list 9 permit 197.135.20.0
access-list 9 permit 197.135.24.0
access-list 9 permit 197.135.27.0
```

Which of the statements below correctly describe the configuration above? (Select two)

- A. Networks 197.135.20.0, 197.135.24.0, and 197.135.27.0 are allowed into the routing table.
- B. The RIP learned routes to networks 197.135.20.0, 197.135.24.0, and 197.135.27.0 will be assigned an administrative distance of 140.
- C. The IGRP learned routes to networks 197.135.20.0, 197.135.24.0, and 197.135.27.0 will be assigned an administrative distance of 140.
- D. Changing the administrative distance to a number larger than the default value makes networks 197.135.20.0, 197.135.24.0, and 197.135.27.0 unreachable.

Answer: A, C

Explanation:

We are redistributing RIP into IGRP. The redistribute rip command specifies that routes learned via RIP will be advertised in the IGRP updates. All of the routes specified in the network statements will be allowed into the routing table. No route filtering is taking place with this command. The access lists are used only to define which networks will have their administrative distances changed when redistributed. In this example, the 3 networks in access list 9 will have their AD changed to 140 when redistributed.

We examine the following command:

```
distance 140 0.0.0.0 255.255.255.255 9
```

The 140 defines the administrative distance that specified routes will be assigned.

The 0.0.0.0 255.255.255.255 defines the source address of the router supplying the routing information, in this case any router.

The 9 defines the access-list to be used to filter incoming routing updates to determine which will have their administrative distance changed.

And one of the access-list statements:

```
access-list 9 permit 197.135.27.0
```

9 is the access-list number.

```
permit
```

allows all networks that match the address to be permitted, in this case to have their administrative distance changed.

197.135.27.0 A network to be permitted, in this case to have its administrative distance changed.

Incorrect Answers:

B: RIP routes are redistributed into IGRP; not vice versa. However, it is not clear that RIP even knows about this route, since the RIP configuration was omitted. We can, however, be certain that IGRP knows about it since these networks were added to the

IGRP routing process.

D: In this case we are changing the AD to 140 for the specified networks. Although the AD has been changed to a value that is higher than the default, they are still reachable.

---

**QUESTION 305:**

**DRAG DROP**

You are a network administrator at Certkiller and you're in the process of migrating their networks from RIP to OSPF (you're using the private address of 10.0.0.0 for routing). You want to redistribute the RIP routes into OSPF and assign a metric of 60 to them. Drag the commands on the left onto the corresponding boxes on the right.

|  |   |
|--|---|
| <div style="border: 1px solid black; background-color: #00FFFF; padding: 2px; margin-bottom: 2px;">router ospf 1</div> <div style="border: 1px solid black; background-color: #00FFFF; padding: 2px; margin-bottom: 2px;">router rip</div> <div style="border: 1px solid black; background-color: #00FFFF; padding: 2px; margin-bottom: 2px;">redistribute ospf 1</div> <div style="border: 1px solid black; background-color: #00FFFF; padding: 2px; margin-bottom: 2px;">redistribute rip</div> <div style="border: 1px solid black; background-color: #00FFFF; padding: 2px; margin-bottom: 2px;">redistribute rip networks</div> <div style="border: 1px solid black; background-color: #00FFFF; padding: 2px; margin-bottom: 2px;">redistribute rip subnets</div> <div style="border: 1px solid black; background-color: #00FFFF; padding: 2px; margin-bottom: 2px;">metric 60</div> <div style="border: 1px solid black; background-color: #00FFFF; padding: 2px;">default-metric 60</div> | <div style="margin-bottom: 20px;">Router(config)# <div style="border: 1px solid black; background-color: #FFFF00; padding: 5px; display: inline-block;">Drop answer here</div></div> <div style="margin-bottom: 20px;">Router(config-router)# <div style="border: 1px solid black; background-color: #FFFF00; padding: 5px; display: inline-block;">Drop answer here</div></div> <div>Router(config-router)# <div style="border: 1px solid black; background-color: #FFFF00; padding: 5px; display: inline-block;">Drop answer here</div></div> |
|--|---|

Answer:

|  |   |
|--|---|
| <div style="border: 1px solid black; background-color: #00FFFF; padding: 2px; margin-bottom: 2px;">router rip</div> <div style="border: 1px solid black; background-color: #00FFFF; padding: 2px; margin-bottom: 2px;">redistribute ospf 1</div> <div style="border: 1px solid black; background-color: #00FFFF; padding: 2px; margin-bottom: 2px;">redistribute rip</div> <div style="border: 1px solid black; background-color: #00FFFF; padding: 2px; margin-bottom: 2px;">redistribute rip networks</div> <div style="border: 1px solid black; background-color: #00FFFF; padding: 2px; margin-top: 20px;">metric 60</div> | <div style="margin-bottom: 20px;">Router(config)# <div style="border: 1px solid black; background-color: #00FFFF; padding: 5px; display: inline-block;">router ospf 1</div></div> <div style="margin-bottom: 20px;">Router(config-router)# <div style="border: 1px solid black; background-color: #00FFFF; padding: 5px; display: inline-block;">redistribute rip subnets</div></div> <div>Router(config-router)# <div style="border: 1px solid black; background-color: #00FFFF; padding: 5px; display: inline-block;">default-metric 60</div></div> |
|--|---|

---

**QUESTION 306:**

A Certkiller router is configured for route redistribution as shown below:

```
ip route 30.0.0.0 255.0.0.0 172.16.1.2
ip route 192.168.1.0 255.255.255.0 172.10.1.2
router eigrp 100
redistribute ospf 100
network 172.19.1.0
```

```
router ospf 108
redistribute static
redistribute eigrp 100
network 172.16.1.0 00.0.255 area 0
distribute-list 5 out static
access-list 5 permit 30.0.0.0 0.255.255.255
```

In the above configuration, what is the function of the command line distribute-list 5 out static?

- A. It denies the route to 30.0.0.0 via OSPF
- B. It denies the route to 30.0.0.0 via EIGRP
- C. It propagates the route to 30.0.0.0 via OSPF
- D. It propagates the route to 30.0.0.0 via EIGRP

Answer: C

Explanation:

The "distribute-list 5 out static" command filters routes learned from static entries by using access list 5, before those routes are passed to the OSPF process. In this example, the static route to the 30.0.0.0/8 network matches access list 5, so it is permitted to be redistributed into the OSPF protocol.

Incorrect Answers:

- A, B: This access-list permits, not denies, routes.
- D: The route is propagated via OSPF, not via EIGRP.

---

### **QUESTION 307:**

You are configuring redistribution to advertise EIGRP routes into OSPF on a boundary router. Given the configuration:

```
router ospf 1
redistribute eigrp 1 metric 25 subnets
```

Which is the function of the subnets in the redistribute command?

- A. It specifies subnetted routes should be advertised into OSPF.
- B. It specifies subnetted routes should be advertised out of OSPF.
- C. It specifies routes that will be summarized on the 25-bit boundary.
- D. It specifies a limit of 25 subnets for each OSPF route advertisement.
- E. None of the above.

Answer: A

Explanation:

The subnets keyword is used to specify that subnet mask information used by classless routing protocols should be preserved when redistributed into the other routing protocol. In this case, the EIGRP routes, including subnet masks, will be redistributed into OSPF.

Incorrect Answers:

B: Here the routes are being placed into the OSPF routing process, not advertised out of.  
C, D: In this configuration, the value of "25" means that the metric will be 25 when placed into OSPF. It has nothing to do with the bit boundary or the number of subnets that can be advertised.

---

**QUESTION 308:**

The Certkiller network is in the process of changing the routing protocol from EIGRP to OSPF, as well as changing the IP network to the private 10.0.0.0/8 network. To do this, router CK1 has been configured as shown below:

```
router ospf 100
redistribute eigrp 100 metric 100 metric-type 1
network 172.16.0 0.0.0.255.255
```

All other configurations use the default values.

You want to successfully redistribute all networks and subnets on the Certkiller network. What can you do to accomplish this goal? (Choose two)

- A. Change the OSPF process-id number from 100 to 1 in the router ospf command.
- B. Configure the redistribute command under router eigrp 1 instead.
- C. Change the EIGRP AS number from 100 to 1 in the redistribute command.
- D. Add the subnets option to the redistribute command.
- E. Add the network 10.0.0.0 0.255.255.255 command under router ospf 100.
- F. Change the metric to an EIGRP compatible metric value for Bandwidth, Delay, Reliability, Load, MTU, (such as 64 1000 100 1 1500) in the redistribute command.

Answer: D, E

Explanation:

The use of the "subnets" keyword will ensure that the subnet mask information will be redistributed into the OSPF process. Since both EIGRP and OSPF support the use of VLSM, this is recommended.

Since the network is being migrated to OSPF, it would be best to begin advertising the network via the OSPF process, as shown in answer choice E.

---

**QUESTION 309:**

You are using multiple protocols in different Autonomous Systems (AS). You need to redistribute between the systems. You are using two-way redistribution.

Which action should help you avoid routing loop issues?

- A. Manually configuring the static routes.
- B. Manually configuring the default gateway.
- C. Manually configuring the administrative K-factor.
- D. Manually configuring the administrative distance.

Answer: D

Explanation:


If you manually configure the administrative distance, you can manually control which router will choose which route thus preventing the possibility of routing loops.

---

**QUESTION 310:**

**SIMULATION**

Network topology exhibit



Certkiller .com recently completed a merger with Foo Inc. The two companies have been using separate routing protocols on their companies corporate networks, and an immediate solution is required for the two companies to begin sharing data. A boundary router, Certkiller 2, has been established to perform mutual redistribution of route information between the two networks. Configure route redistribution from EIGRP into IS-IS and from IS-IS into EIGRP on the boundary router per the following requirements.

1. Seed metric for EIGRP must have the following characteristics:

Bandwidth=64 Kbps

2. Delay=100

3. Reliability=255

4. Load=1

5. MTU=1500

Seed metric for IS-IS must be set to 50.

Only redistribute Level-1 IS-IS routes into EIGRP.

Redistribute EIGRP routes into IS-IS as Level-1 routes.

Name: Certkiller 1

S0/0: 10.10.10.5/30

Loopback1: 1.1.1.1/32

Secret Password: Certkiller

Name: Certkiller 2

S0/0: 192.168.1.6/30

Loopback1: 2.2.2.2/32

Secret Password: Certkiller

Name: Certkiller 3

S0/0: 192.168.1.5/30

Loopback1: 3.3.3.3/32

Secret Password: Certkiller

Answer:

Explanation:

Certkiller 2> en

Certkiller 2# config t

Certkiller 2(config-router)# router isis

Certkiller 2(config-router)# redistribute eigrp 100 level-1 metric 50

Certkiller 2(config-router)# exit

Certkiller 2(config)# router eigrp 100

Certkiller 2(config-router)# redistribute isis level-1 metric 64 10 255 1 1500

Certkiller 2(config-router)# redistribute connected

Certkiller 2(config-router)# ^z

Certkiller 2# copy running-config start up-config.

Alternative #1:

1. Seed metric for EIGRP must have the following characteristics:

Bandwidth=64 Kbps

2. Delay=2 ms

3. Reliability=255

4. Load=1

5. MTU=1500

Name: Certkiller 1

S0/0: 10.197.197.5/30

Loopback1: 1.1.1.1/32

Secret Password: Certkiller

Name: Certkiller 2

S0/0: 192.168.191.6/30

Loopback1: 2.2.2.2/32

Secret Password: Certkiller

Name: Certkiller 3

S0/0: 192.168.191.5/30

Loopback1: 3.3.3.3/32

Secret Password: Certkiller

Alternative #2:

1. Seed metric for EIGRP must have the following characteristics:

Bandwidth=512 Kbps

2. Delay=5 ms

3. Reliability=255

4. Load=1

5. MTU=1500

6. Seed metric for IS-IS must be set to 24

Name: Certkiller 1

S0/0: 10.117.117.5/30

Loopback1: 1.1.1.1/32

Secret Password: Certkiller

Name: Certkiller 2  
S0/0: 192.168.72.6/30  
S0/1: 10.117.117.6/30  
Secret Password: Certkiller  
Name: Certkiller 3  
S0/0: 192.168.72.5/30  
Loopback1: 3.3.3.3/32  
Secret Password: Certkiller

---

**QUESTION 311:**

Which three statements are true when configuring redistribution for OSPF? (Select three)

- A. The default metric is 10.
- B. The default metric is 20.
- C. The default metric type is 2.
- D. The default metric type is 1.
- E. Subnets do not redistribute by default.
- F. Subnets redistribute by default.

Answer: B, C, E

Explanation:

If a metric is not specified, OSPF puts a default value of 20 when redistributing routes from all protocols except Border Gateway Protocol (BGP) routes, which get a metric of 1.

For all redistributed OSPF routes, the default metric type is 2.

When there is a major net that is subnetted, you need to use the keyword subnet to redistribute protocols into OSPF. Without this keyword, OSPF only redistributes major nets that are not subnetted.

Reference:

[http://www.cisco.com/en/US/tech/ CK3 65/technologies\\_tech\\_note09186a008009487e.shtml#ospf](http://www.cisco.com/en/US/tech/ CK3 65/technologies_tech_note09186a008009487e.shtml#ospf)

---

**QUESTION 312:**

Router CK1 is configured with OSPF and EIGRP as shown in the diagram below:

```
router ospf 5
network 10.0 0 0 0.0.255.255 area 0
```

```
router eigrp 100
network 172.10.0.0
redistribute ospf 5
```

You work as a network technician at Certkiller .com. You are troubleshooting a redistribution of OSPF routes into EIGRP. Study the exhibit carefully.

Which statement is true?

- A. Redistributed routes will have an external type of 1 and a metric of 1.
- B. Redistributed routes will have an external type of 2 and a metric of 20
- C. Redistributed routes will maintain their original OSPF routing metric.
- D. Redistributed routes will have a default metric of 0 and will be treated as reachable and advertised.
- E. Redistributed routes will have a default metric of 0 but will be treated as unreachable and not advertised.

Answer: E

Explanation:

A default metric is required to redistribute a protocol into IGRP or EIGRP. Only connected routes and interface static routes can be redistributed without a default metric. If a default metric is not specified, it will default to a metric of 0 (infinity) and will not be advertised.

Reference:

[http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products\\_command\\_reference\\_chapter09186a008009](http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_command_reference_chapter09186a008009)

---

### **QUESTION 313:**

Which three statements are true concerning redistributed routes when a default metric is not configured? (Select three)

- A. RIP, IGRP, and EIGRP assign a default metric of 0 (infinity) to redistributed routes and will advertise these routes accordingly.
- B. RIP, IGRP, and EIGRP assign a default metric 0 (infinity) to redistributed routes and will only advertise these routes if a valid seed metric is configured.
- C. IS-IS assigns a default metric of 0 to redistributed routes.
- D. IS-IS assigns a default metric of 10 to redistributed routes.
- E. OSPF assigns a default metric of 1 for routes from all protocols except BGP, which gets a metric of 20.
- F. OSPF assigns a default metric of 20 for routes from all protocols except BGP, which gets a metric of 1.

Answer: B, C, F

Explanation:

A default metric is required to redistribute a protocol into IGRP or EIGRP. Only connected routes and interface static routes can be redistributed without a default metric. If a default metric is not specified, it will default to a metric of 0 (infinity) and will not be advertised

The IS-IS metric must be between 1 and 63. There is no default-metric option in IS-IS-you should define a metric for each protocol, as shown in the example above. If no

metric is specified for the routes being redistributed into IS-IS, a metric value of 0 is used by default.


If a metric is not specified, OSPF puts a default value of 20 when redistributing routes from all protocols except Border Gateway Protocol (BGP) routes, which get a metric of 1.

Reference:

[http://www.cisco.com/en/US/tech/CK365/technologies\\_tech\\_note09186a008009487e.shtml#isis](http://www.cisco.com/en/US/tech/CK365/technologies_tech_note09186a008009487e.shtml#isis)

### QUESTION 314:

The Certkiller WAN is shown in the following diagram:



Exhibit, Router Configuration

Certkiller2# **show ip route**

Gateway of last resort is not set

```

    2.0.0.0/24 is subnetted, 1 subnets
C       2.2.2.0 is directly connected, Ethernet0/0
C       3.0.0.0/8 is directly connected, Serial1/0
O E2 200.1.1.0/24 [110/20] via 2.2.2.2, 00:16:17, Ethernet0/0
O E1 200.2.2.0/24 [110/104] via 2.2.2.2, 00:00:41, Ethernet0/0
    131.108.0.0/24 is subnetted, 2 subnets
O   131.108.2.0 [110/74] via 2.2.2.2, 00:16:17, Ethernet0/0
O IA 131.108.1.0 [110/84] via 2.2.2.2, 00:16:17, Ethernet0/0

```

Which command should be added to Certkiller B under router bgp 100 to allow only the external OSPF routes to be redistributed to Certkiller C?

- A. redistribute ospf 1
- B. redistribute ospf 1 match external 1
- C. redistribute ospf 1 match external 2
- D. redistribute ospf 1 match external 1 external 2

Answer: D

Explanation:

### Redistribution of Only OSPF External (Type 1 and 2) Routes into BGP

Use the external keyword along with the redistribute command under router bgp to redistribute OSPF external routes into BGP. With the external keyword, you have three choices:

1. redistribute both external type-1 and type-2 (Default)
2. redistribute type-1
3. redistribute type-2

Enter the commands in the configuration mode as described here:

```
RTB(config-router)# router bgp 100
```

```
RTB(config-router)# redistribute ospf 1 match external
```

Note: The exact scenario and network diagram shown in this question can be found in the reference link below.


Reference:

[http://www.cisco.com/en/US/tech/CK365/technologies\\_tech\\_note09186a00800943c5.shtml](http://www.cisco.com/en/US/tech/CK365/technologies_tech_note09186a00800943c5.shtml)

---

### QUESTION 315:

The Certkiller routed network is shown below:



The relevant portion of the configuration of router Certkiller 2 is shown below:

```
hostname Certkiller2
!  
router igrp 7  
  network 181.16.0.0  
  redistribute rip metric 1 1 1 1  
  distribute-list 1 in serial 1  
!  
router rip  
  network 178.1.0.0  
  redistribute igrp 7 metric 2  
!  
access-list 1 deny 192.168.1.0  
access-list 1 permit any
```

You work as a network technician at Certkiller .com. Study the exhibits carefully. Both Certkiller 5 and Certkiller 2 are performing two-way IGRP/RIP route redistribution.

What behavior will Certkiller 2 use with the given Certkiller 2 configurations?

- A. Updates about 192.168.1.0 that are learned through RIP will be ignored.
- B. Updates about 192.168.1.0 will be advertised out the s0 interface.
- C. Updates about 192.168.1.0 that are heard through the s1 interface will be ignored.
- D. Updates about 192.168.1.0 that are learned through IGRP will be added to the routing table, but not advertised to Certkiller 1.

Answer: C

Explanation:

When a distribute list is defined, the access list that it is tied to is used to define the routing updates that are permitted. In our example, the access list specifies that only the 192.168.1.0 network is to be filtered, while all other network advertisements are allowed. Since the distribute list specifies that this filter is to be used for incoming updates on interface serial 1, choice C is correct.

---

### **QUESTION 316:**

Router CK1 is configured as shown below:

```
router rip  
network 10.0.0.0  
!  
router ospf 5  
network 172.10.0.0 0.0.255.255 area 0  
redistribute rip
```

Based on the exhibited configuration commands, which statement is true?

- A. Redistributed routes will be tagged as external type 1 (E1) with a metric of 30.

- B. Redistributed routes will be tagged as external type 2 (E2) with a metric of 20.
- C. Redistributed routes will maintain their original RIP routing metric.
- D. Redistributed routes will have a default metric of 0 and will be treated as unreachable and not advertised.
- E. Redistributed routes will have a default metric of 0 but will not be treated as reachable and will not be advertised.

Answer: B

Explanation:


If a metric is not specified, OSPF puts a default value of 20 when redistributing routes from all protocols except Border Gateway Protocol (BGP) routes, which get a metric of 1. The default metric type for OSPF is external, type 2.

Reference:

[http://www.cisco.com/en/US/tech/CK365/technologies\\_tech\\_note09186a008009487e.shtml#ospf](http://www.cisco.com/en/US/tech/CK365/technologies_tech_note09186a008009487e.shtml#ospf)

---

**QUESTION 317:**



During the redistribution process configured on Certkiller A, some of the EIGRP routes, such as 10.1.1.0/24 and 10.2.2.0/24 are not being redistributed into the OSPF routing domain. Which two items could be a solution for this problem? (Choose two.)

- A. Change metric-type to 2 in the redistribute command.
- B. Configure the redistribute command under router eigrp 1 instead.
- C. Change the EIGRP AS number from 1 to 100 in the redistribute command.
- D. Add the subnets option to the redistribute command.
- E. Change the metric to an EIGRP compatible metric value (bandwidth, delay, reliability, MTUs) in the redistribute command.

Answer: C, D

---

**QUESTION 318:**

Your network appears to be experiencing some issues relating to policy based routing. Which troubleshooting command could you use to view the current route

maps configured on a specific interface?

- A. show interface
- B. show route-map
- C. show ip policy
- D. show ip route map
- E. None of the above

Answer: C

Explanation:

The show ip policy command is used to display which route map is associated with which interface.

Incorrect Answers:

A: The show interface command shows detailed information about the Cisco router/switch. However, the output does not include information on configured route maps.

B: The show route-map command displays the contents of all route maps or the specified route map.

D: This is an invalid command.

---

**QUESTION 319:**

A router is configured for policy based routing as shown below:

```
interface serial 0
ip policy route-map demo
route-map demo permit 10
match ip address 4
set interface serial2 serial3
access list 4 permit 10.3.3.2 0.0.0.0
```

Based on the information above, which of the following statements is true?

- A. If the packet is sourced from 10.3.3.2, it is a candidate for fast-switched policy routing.
- B. If the packet is destined for 10.3.3.2, it will be routed out interface serial 2 and interface serial 3 in a load-sharing fashion.
- C. If the packet is sourced from 10.3.3.2, it will be routed out interface serial 2, unless it is not up, in which case it will be routed out interface serial 3.
- D. If the packet is sourced from 10.3.3.2, the flow will be routed out interface serial 2 and interface serial 3 in a load-sharing fashion for the duration of the flow.

Answer: C

Explanation:

The set command specifies the list of interfaces which may be used to forward the traffic. If the first interface goes down, the second interface is being used. In this example,

packets sourced from 10.3.3.2 (standard access list matches based on the source, not the destination) will be sent out the serial 2 interface. These packets will be sent to the serial 3 interface only if the serial 2 interface is unreachable (down).

---

**QUESTION 320:**

When configuring policy based routing using the "route-map" command, what is the final command you have to enter to complete the configuration?

- A. Set
- B. Match
- C. Map-list
- D. IP policy

Answer: A

Explanation:

Route maps are complex access lists: A collection of route-map statements that have the same route-map name are considered one route-map. Route maps are configured by using match commands to match the criteria for the action, followed by a set command to define the action to take. To configure route maps, perform the following:

Step 1: Router Certkiller (config)# route-map map-tag [permit | deny] [ sequence-number]

First we define the conditions for policy routing.

Step2: Router Certkiller (config-route-map)#match { conditions }

Then we define the conditions to match

Step 3: Router Certkiller (config-route-map)# set { actions }

Finally we define the action to be taken on a match.

---

**QUESTION 321:**

If you're using policy-based routing, what could you do to prevent packets with no match in the route map from being returned to the normal forwarding process?

- A. Set the next-hop metric to 255 for packets without a match.
- B. Use a set statement to route packets to the null0 interface.
- C. Use a "deny any" statement as the last statement in the route map.
- D. Use a logical "OR" in the match statement to send packets to the null0 interface.

Answer: B

Explanation:

If it is desired not to revert to normal forwarding and to drop a packet that does not match the specified criteria, then a set statement to route the packets to interface null 0 should be specified as the last entry in the route-map.

Incorrect Answers:

A: Setting a metric of 255 will not necessarily make the route unreachable, since this is a

valid metric for OSPF, IGRP, and EIGRP.

C: The "deny all" statement is already implicitly enabled at the end of each access list that is used in a route map. As opposed to regular access lists, however, traffic that does not match the permit statements are routed using the normal forwarding process.

D: It is not possible to use a logical OR statement. Doing so would mean that all traffic could be routed to null 0.

---

**QUESTION 322:**

At the end of every route map, there's a command to implicitly "deny any" What is the end result of this rule?

- A. Packets that reach the end of the route map are discarded.
- B. Packets are forwarded to the null interface for special handling.
- C. Packets that reach the end of the route map are routed in a normal fashion.
- D. Packets that reach the end of the route map are returned in the originating interface.

Answer: C


Explanation:

The implicit deny any in every route map means that packets not meeting any of the matching criteria are routed in the normal fashion. This is in contrast to regular access lists, where packets that fail to match any of the permit statements are discarded.

---

**QUESTION 323:**

The Certkiller network is displayed below:



You need to implement a route map on the E0 interface of router Certkiller 1 (name it ISPA). Which command would be issued first to do this?

- A. Certkiller 1(config)#interface e0

- Certkiller 1(config-if)ip route-map ISPA
- B. Certkiller 1(config)#interface e0
- Certkiller 1(config-if)#ip policy route-map ISPA
- C. Certkiller 1(config)#interface e0
- Certkiller 1(config-if)#policy route-map ISPA
- D. Certkiller 1(Config)#interface e0
- Certkiller 1(Config-if)#policy route map ISPA

Answer: B

Explanation:

The correct syntax to enable an interface for policy based routing is "ip policy route-map name." One interface can only have one route-map tag, but you can have multiple route map entries with different sequence numbers. These entries are evaluated in sequence number order until the first match. If there is no match, packets will be routed as usual.

Reference:

[http://www.cisco.com/en/US/products/hw/switches/ps4324/products\\_configuration\\_guide\\_chapter09186a008019](http://www.cisco.com/en/US/products/hw/switches/ps4324/products_configuration_guide_chapter09186a008019)

---

### QUESTION 324:

#### DRAG DROP

You have the following access list statement configured on your router:


```
access-list 5 permit 176.234.5.0
```

Your goal is to configure a route map to modify the metric to 30 for the network in the access list above. Drag the proper commands on the right side to its proper command sequence on the left side. (Hint: you aren't required to apply the route map yet.)

| Place here  | Select from these  |
|---|--|
| Router(config)# <input type="text" value="Place here"/>           | <input type="text" value="Set distance 30"/>                 |
| Router(config-route-map)# <input type="text" value="Place here"/> | <input type="text" value="Set administrative- distance 30"/> |
| Router(config-route-map)# <input type="text" value="Place here"/> | <input type="text" value="Set metric 30"/>                   |
|   | <input type="text" value="Match network 5"/>                 |
|   | <input type="text" value="Match ip-class 5"/>                |
|   | <input type="text" value="Match ip address 5"/>              |
|   | <input type="text" value="Route-map permit ip 10 CertK"/>    |
|   | <input type="text" value="Route-map CertK permit 10"/>       |

Answer:

## 642-801



Explanation:

Step 1: First we must enter Route-Map Configuration mode. We issue the following command

```
route-map CertK permit 10
```

CertK is the tag, we permit (the alternative is to deny), and we use the sequence number 10.

Step 2: We define the match condition.

```
Match ip address 5
```

We match the IP address to the IP address of access-list 5, namely 176.234.5.0.

Step 3: We apply the set statement.

```
set metric 30
```

We should simply change the metric.

Note: Syntax route-map

```
route-map map-tag [[permit | deny] | [sequence-number]]
```

Each route map will consist of a list of match and set configuration. The match will specify a match criteria and set specifies a set action if the criteria enforced by the match command are met.

The related commands for match are: match as-path, match community, match cns, match interface, match ip address, match ip next-hop, match ip route-source, match metric, match route-type, match tag.

The related commands for set are: set as-path, set cns, set automatic-tag, set community, set interface, set default interface, set ip default next-hop, set level, set local-preference, set metric, set metric-type, set next-hop, set origin, set tag, set weight.

Incorrect Answers:

```
Route-map permit ip 10 CertK
```

The command has the wrong syntax. The tag (here: word), must be immediately after the keyword route-map.

```
Match ip-class 5 and Match ip-class 5
```

There are no such commands.

```
Set distance 30 and Set administrative- distance 30
```

There are no such commands.

**QUESTION 325:**

The Certkiller router is running on Cisco IOS version 11.2. On a router running Cisco IOS 11.2 or later, what is the default switching mode for forwarding packets that coincide with the established routing policy?

- A. fast
- B. slow
- C. NetFast
- D. process

Answer: A

Explanation:

Starting with IOS 11.2, the default switching mode is fast. Prior to this, it was process switching. The following describes these two different methods.

Process Switching

In process

switching the first packet is copied to the system buffer. The router looks up the Layer 3 network address in the routing table and initializes the fast-switch cache. The frame is rewritten with the destination address and sent to the outgoing interface that services that destination. Subsequent packets for that destination are sent by the same switching path.

Fast Switching

When packets are fast switched, the first packet is copied to packet memory and the destination network or host is found in the fast-switching cache. The frame is rewritten and sent to the outgoing interface that services the destination. Subsequent packets for the same destination use the same switching path.

---

**QUESTION 326:**

DRAG DROP

You are the network engineer at Certkiller . You want to configure a route map that will modify the metric for the network in the following access list:

```
access-list 1 permit 192.168.1.0
```

Drag the commands in the proper order to configure the route map. You are not required to apply the route map at this time.

## 642-801

|                                |                                       |                  |
|--------------------------------|---------------------------------------|------------------|
| set distance 50                | Router Certkiller(config)#            | Drop answer here |
| set administrative-distance 50 | Router(config-route-map)#             | Drop answer here |
| set metric 50                  | Router Certkiller (config-route-map)# | Drop answer here |
| set network 1                  |                                       |                  |
| match network 1                |                                       |                  |
| match ip-class 1               |                                       |                  |
| match ip address 1             |                                       |                  |
| route-map permit 10 word       |                                       |                  |
| route-map word permit 10       |                                       |                  |

Answer:

|                                |                                       |                          |
|--------------------------------|---------------------------------------|--------------------------|
| set distance 50                | Router Certkiller(config)#            | route-map word permit 10 |
| set administrative-distance 50 | Router(config-route-map)#             | match ip-address 1       |
| set network 1                  | Router Certkiller (config-route-map)# | set metric 50            |
| match network 1                |                                       |                          |
| match ip-class 1               |                                       |                          |
| route-map permit 10 word       |                                       |                          |

Explanation:

route-map word permit 10

match ip address 1

set metric 50

The route-map command is shown here:

```
Router(config)#route-map map-tag [{permit | deny}
sequence-number]
```

The match commands used in policy-based routing are summarized in Table the table.

These match commands are used to determine whether the packet is to be policy-routed, as opposed to being forwarded simply by destination. If it is to be policy-routed, the packet is sent down a different path, typically one less traveled.

The set Commands for Policy-Based Routing with Route Maps

The set commands used in policy-based routing are summarized in Table 18-4. These set commands are used after the match criteria has been satisfied. Whereas the match parameter determines whether the packet will be policy-routed, the set parameter determines how the packet is to be policy-routed.

Table 18-3 The match Commands used in Policy Based Routing

| Command                        | Description  |
|--------------------------------|--|
| match ip address               | This states the number or name of a  |
|                                | standard or extended access list that  |
| [access-list-number   name]    | will be used to examine incoming   |
|                                | packets. A standard IP access list is  |
| [...access-list-number   name] | used to match criteria for the source address of the packet. An extended IP access list is used to specify criteria based on source and destination, |
|                                | application, protocol type, TOS, and precedence. If multiple access lists are specified, matching any one will result                                |

---

**QUESTION 327:**

By default, what actions will a router take when the configuration command "route-map cisco" is applied? (Choose two)

- A. Any matches are considered to be permitted
- B. Any matches are considered to be denied
- C. The default sequence number will be 1.
- D. The default sequence number will be 5.
- E. The default sequence number will be 10.

Answer: A, E

Explanation:

By default, when a route-map is applied the initial sequence number that is enabled by the router is 10. In addition, all matching traffic will be permitted by default.


---

**QUESTION 328:**

```

hostname CertkillerB
!
interface serial 0/0
ip policy route-map SNOW
!
route-map SNOW permit 5
match ip address 1
set ip next-hop 172.16.4.2
!
route-map SNOW deny 10
match ip address 2
set ip next-hop 172.16.4.2
!
access-list 1 permit 172.16.6.5 0.0.0.0
access-list 2 permit 172.16.6.10 0.0.0.0

```



Refer to the diagram above. Policy-based routing is enabled for the incoming traffic on Certkiller B.

Given the above configuration, how is the packet that comes from 172.16.6.10 routed?

- A. The packet is denied access and is dropped.
- B. The packet is routed through the normal destination-based routing process.
- C. The packet is sent to Certkiller C.
- D. The packet is sent back to Certkiller A.


Answer: B

Explanation:

When policy routing is used, the packets will go through the route-map process until a permitting match is found. If no match is found, then the normal routing behavior will occur. In the case here, the packet does indeed appear to match the second policy statement, but statement 10 is a deny statement, so the next policy statement will be looked at. Since there are no other policy statements, none will match and the normal routed behavior will occur.

**QUESTION 329:**

The Certkiller WAN is displayed in the diagram below:



Which configuration will force traffic to use the NY link to reach the 192.168.10.0/24 and 192.168.11.0/24 networks?

- A. Certkiller C(config)# router bgp 100  
 Certkiller C(config-router)# neighbor 192.168.5.2 route-map LOCAL in  
 Certkiller C(config)# route-map LOCAL permit 10  
 Certkiller C(config)# set local-preference 300
- B. Certkiller B(config)# router bgp 100  
 Certkiller B(config-router)# neighbor 172.16.20.2 route-map LOCAL in  
 Certkiller B(config)# route-map LOCAL permit 10  
 Certkiller B(config)# set local-preference 300
- C. Certkiller C(config)# router bgp 100  
 Certkiller C(config-router)# neighbor 192.168.5.2 route-map LOCAL out  
 Certkiller C(config)# route-map LOCAL permit 10  
 Certkiller C(config)# set local-preference 50
- D. Certkiller B(config)# router bgp 100  
 Certkiller B(config-router)# neighbor 172.16.20.2 route-map LOCAL out  
 Certkiller B(config)# route-map LOCAL permit 10  
 Certkiller B(config)# set local-preference 300

Answer: A

**Explanation:**

The default local preference value of any link is 100, and with BGP a higher local preference is preferred over a lower value. In choice A, the local preference is being set to 300, so it will be preferred.

**Incorrect Answers:**

- B. This will set the local preference to 300 for the link via AS300, but we want to prefer the path via AS200.
- C. This will set the preference to 50, but the default value of is 100 so this will force the NY link to be less desirable than the other.
- D. We wish to apply the higher preference map to the 192.168.5.2 peer, not 172.16.20.1

---

**QUESTION 330:**

Router CK1 was configured as shown below:

```
router bgp 50001
network 10.0.0.0
network 172.16.0.0
neighbor 10.1.1.1 remote-as 5003
neighbor 10.1.1.1 route-map test out
!
access-list 1 permit 10.0.0.0/8
access-list 2 permit any
!
route-map test permit 10
match ip address 1
set metric 200
!
route-map test permit 20
match ip address 2
```

Based on the information shown above, what does the route map named test accomplish?

- A. Marks all prefixes advertised to the 10.1.1.1 neighbor with a MED of 200.
- B. Marks all prefixes received from the 10.1.1.1 neighbor with a MED of 200.
- C. Marks the 10.0.0.0/8 prefix advertised to the 10.1.1.1 neighbor with a MED of 200.
- D. Marks the 10.0.0.0/8 prefix received from the 10.1.1.1 neighbor with a MED of 200.
- E. Permits only the 10.0.0.0/8 prefix to be advertised to the 10.1.1.1 neighbor.
- F. Permits only the 10.0.0.0/8 prefix to be received from the 10.1.1.1 neighbor.

Answer: C

Explanation:

Since the route map test is applied in the outbound direction, it will be used for advertised updates to the 10.1.1.1 neighbor. The first statement matches, meaning that the 10.0.0.0 network matches access list 1, so the MED value of 200 will be placed when advertising the network out to the BGP peer.

---

**QUESTION 331:**

Router CK1 was configured as shown below:

642-801

```
!  
access-list 100 permit ip host 100.100.100.3 host 200.200.200.4  
!  
route-map ddt permit 10  
match ip address 100  
set ip default next-hop 10.10.10.2  
!
```

You work as a network administrator at Certkiller .com. Study the network topology exhibit carefully. Which statement describes the result of using the set ip default next-hop command?

- A. If a packet received by the router matches the access list, it will always be forwarded to 10.10.10.2
- B. If a packet received by the router matches the access list, it will be forwarded to 10.10.10.2 only if a path to 200.200.200.0 is present in the routing table.
- C. If a packet received by the router matches the access list, it will be forwarded to 10.10.10.2 only if a path to 200.200.200.0 is not present in the routing table.
- D. If a packet received by the router matches the access list, it will be forwarded to 200.200.200.4 only if a path to 100.100.100.0 is present in the routing table.
- E. If a packet received by the router matches the access list, it will be forwarded to 200.200.200.4 only if a path to 100.100.100.0 is not present in the routing table.

Answer: C

Explanation:

The set ip default next-hop command verifies the existence of the destination IP address in the routing table, and...

1. if the destination IP address exists, the command does not policy route the packet, but forwards the packet based on the routing table.
2. if the destination IP address does not exist, the command policy routes the packet by sending it to the specified next hop.


Reference:

[http://www.cisco.com/en/US/tech/CK364/technologies\\_configuration\\_example09186a00801f3b54.shtml](http://www.cisco.com/en/US/tech/CK364/technologies_configuration_example09186a00801f3b54.shtml)

---

**QUESTION 332:**

The Certkiller WAN is shown in the diagram below:



Exhibit

```
Certkiller1(config)# route-map REDIST-MAP permit 10  
Certkiller1(config-route-map)# match ip address 101  
Certkiller1(config-route-map)# set metric 100  
Certkiller1(config-route-map)# set metric-type type-1  
Certkiller1(config-route-map)# exit
```

```
Certkiller1(config)# route-map REDIST-MAP deny 20  
Certkiller1(config-route-map)# match ip address 110  
Certkiller1(config-route-map)# exit
```

```
Certkiller1(config)# route-map REDIST-MAP permit 30  
Certkiller1(config-route-map)# set metric 5000  
Certkiller1(config-route-map)# set metric-type type-2  
Certkiller1(config-route-map)# exit
```

```
Certkiller1(config)# router ospf 10
```

Based on the above information, which statement is true regarding the Certkiller .com network?

- A. There is a match entry missing in the route-map REDIST-MAP permit 30 statement.
- B. All routes not matching access list 101 or 110 will be flagged as an OSPF external type 2 with a metric of 5000.
- C. There is a set entry missing in the route-map REDIST-MAP deny 20 statement.
- D. All routes matching access list 110 will be forwarded.
- E. The set metric-type entry is not a valid route-map command.

Answer: B

Explanation:

With policy routing, the rules are looked at in order based on the permit or deny statements at the end of the route map name. The first statement, numbered 10, will match all packets that match access list 101. For packets not matching access list 101, the next rule is looked at, which is rule number 20. This rule matches access list 110. Finally, for routes not matching either of these, the last rule is consulted. Since no access list is used, all remaining routes will match, and their metric types will be set to 2, with a metric of 5000.

---

### QUESTION 333:

Your network has multiple connections to the Internet via multiple ISPs. Which of the following are common practices for configuring the network to receive BGP routes from these Internet Service Providers?

- A. Accept full routes from the ISPs.
- B. Accept only IGP routes from the ISPs
- C. Accept an external route from the ISPs.
- D. Accept only redistributed routes from the ISPs.
- E. All of the above

Answer: A

Explanation:

The configuration of the multiple connections to the ISPs can be classified depending on the routes that are provided to the AS from the ISPs. Three common ways of the configuring the connections are:

1. All ISPs pass only default routes to the AS.
2. All ISPs pass default routes, and selected specific routes (for example, from customers with who the AS exchanges a lot of traffic) to the AS.
3. All ISPs pass all routes to the AS (A).

---

**QUESTION 334:**

Regarding the use of BGP prefix lists, which one of the statements below is true?

- A. They start at 10 by default.
- B. They automatically increment by ten by default.
- C. They are displayed using the show ip prefix-list command.
- D. They are displayed using the show bgp prefix-list command.
- E. None of the above.

Answer: C

Explanation:

The show ip prefix-list command is used to display information about a prefix list or prefix list entries.

Note: Filtering by prefix list involves matching the prefixes of routes with those listed in the prefix list, similar to using access lists.

Incorrect Answers:

- A, B: By default, the entries of a prefix list will have sequence values of 5, 10, 15 etc. They start at 5 and increment by 5.
- D: This is an invalid command.

---

**QUESTION 335:**

BGP is being used on the Certkiller network for Internet routing.

In EBGP, which of the following configuration lines advertises the subnet 154.2.1.0 255.255.255.0 to EBGP neighbors?

- A. Router (config-router)#network 154.2.1.0
- B. Router (config-router)#network 164.2.1.0
- C. Router (config-router)#network-advertise 154.2.1.0
- D. Router (config-router)#network 154.2.1.0 mask 255.255.255.0

Answer: D

Explanation:


When configuring BGP, we specify the mask keyword also along with the network command. If the network mask is omitted, the IP network is taken at the network class boundary.

---

**QUESTION 336:**

Routers A, B, and C are all part of the same autonomous system and are configured for BGP routing. The relevant configurations for RTR A and RTR B are shown below:

```
Hostname RTR A
!
interface serial 0/0
ip address 140.140.13.2 255.255.255.252
description Connection to ISP A
!
interface 10/0
ip address 10.10.10.2 255.255.0
description Connection to RTR B
!
interface f0/1
ip address 172.16.30.1 255.255.255.0
description Connection to RTR C
router bgp 5201
neighbor 10.10.10.3 remote-as 5201
neighbor 140.140.13.2 remote-as 5203
Hostname RTR B
!
interface 10/0
ip address 10.10.10.3 255.255.255.0
description Connection to RTR A
!
interface 10/1
ip address 10.10.20.2 255.255.255.0
description Connection to RTR C
!
router bgp 5201
neighbor 10.10.10.2 remote-as 5201
neighbor 10.10.20.1 remote-as 5201
```



Based on the information provided, what will RTR A do with the updates it receives from ISP A?

- A. Send the update to RTR B and C.
- B. Send the update to RTR B only.
- C. Ignore and discard the packet.
- D. Update its tables and discard the packet.

Answer: C

Explanation:

When specifying the remote external AS in the BGP peer statement, the IP address of the neighbor router should be used as the peer. In this example, the remote peer specified in RTR A is 140.140.13.2, but as we can see from the configuration file this is the IP address of its own serial interface. The correct statement should have been "neighbor 140.140.13.1 remote-as 5203" under the BGP process.

Because of this, the EBGP peering session will not be established and all BGP updates from ISP A will be ignored.

---

**QUESTION 337:**

Your network has 2 separate T1 connections with two different providers. TO support this, BGP is being configured. The required information regarding the ISPs is shown below:

Your network: 164.67.36.0/24

Your AS number: 300

AS number of ISP1: 1005

AS number of ISP2: 1010

Based on the information provided, which of the following commands would you enter to advertise your network to ISP1 and ISP2 via BGP?

- A. router bgp 1005  
network 164.67.36.0 mask 255.255.255.0

```
neighbor 15.1.1.1 remote-as 1005
neighbor 25.1.1.1 remote-as 1010
B. router bgp 300
network 164.67.36.0 mask 255.255.255.0
neighbor 15.1.1.1 remote-as 1005
neighbor 25.1.1.1 remote-as 1010
C. router bgp 300
network 164.67.36.0
neighbor 15.1.1.1 remote-as 1005
neighbor 25.1.1.1 remote-as 1010
D. router bgp 1010
network 164.67.36.0
neighbor 15.1.1.1 remote-as 1005
neighbor 25.1.1.1 remote-as 1010
```

Answer: B

Explanation:

Step 1: router bgp 300

The router bgp command is used to activate the BGP protocol and identify the local autonomous system.

Step 2: network 164.67.36.0 mask 255.255.255.0

The network command controls which networks are originated by this router.

Syntax: network network-number network-mask

Step 3:

```
neighbor 15.1.1.1 remote-as 1005
```

```
neighbor 25.1.1.1 remote-as 1010
```

The neighbor remote-as command to identify a peer router with which the local router will establish a session.

Incorrect Answers:

A, D: We must specify the local autonomous system in the router bgp command. We must use AS 300, not AS 1005 or AS 1010

C: Both the network number and the network mask must be specified with the network command.

---

### **QUESTION 338:**

The Certkiller Internet router is configured as follows::

```
router bgp 64000
network 17.0.0.0
neighbor 178.5.1.1 remote-as 64000
neighbor 197.4.1.2 remote-as 64100
```

Based on this router configuration, which of the following statements are true?  
(Select two)

A. The router with IP address 178.5.1.1 runs IBGP with RouteCK.

- B. The router with IP address 178.5.1.1 runs EBGP with RouterCK.
- C. The router with IP address 197.4.1.2 runs EBGP with RouterCK.
- D. The router with IP address 178.5.1.1 runs as a community member with RouterCK.
- E. The router with IP address 197.4.1.2 runs as a peer group member with RouteCK.

Answer: A, C

Explanation:

The router bgp command is used to activate the BGP protocol and identify the local autonomous system.

The neighbor command activates a BGP session with another router using either IBGP or EBGP.

Syntax: neighbor { ip-address|peer-group-name }remote-as autonomous-system

The value placed in the autonomous system field of the neighbor command determines whether the communication with the neighbor is an EBGP or an IBGP session.

A: If the autonomous system field configured in the router bgp command is identical to the field in the neighbor remote-as command, then BGP will initiate an internal session (IBGP). Here both the local AS and the neighbor 178.5.1.1 are in the AS 64000.

C: If the field values are different, then BGP will initiate an external session (EBGP). The network 197.4.1.2 has a different AS number.

---

### **QUESTION 339:**

You wish to permit all BGP prefixes between /10 and /18 for the 207.0.0.0 network. How should a BGP prefix list be configured to do this?

- A. ip prefix-list 207.0.0.0/8 ge 10 le 18
- B. ip prefix-list 207.0.0.0/8 ge 18 le 10
- C. ip prefix-list 207.0.0.0/24 ge 10 le 18
- D. ip prefix-list 207.0.0.0/24 ge 18 le 10
- E. None of the above

Answer: A

Explanation:

With prefix lists, the "ge" means "greater than or equal to" while the le means "less than or equal to." Here, we want to specify all prefixes greater than or equal to 10 and less than or equal to 18 as specified in choice

A. In addition, we must specify an 8 bit network mask, not a 24 bit network mask.

---

### **QUESTION 340:**

Routers CK1 and CK2 are configured for BGP as shown below:

```
RTR CK1
router bgp 200
```

```
neighbor 183.215.22.1 remote-as 200
neighbor 183.215.22.1 update-source loopback 1
RTR CK2
router bgp 200
```

```
neighbor 147.229.1.1 remote-as 200
```

These two routers currently have an active BGP peering session between them and they are able to pass routes to each other. Based on these configuration files, which of the following are true? (Select three)

- A. RTR CK1 and RTR CK2 are running IBGP inside AS 200
- B. The IP address of RTR CK1 's Loopback 1 interface is 147.229.1.1.
- C. The IP address of RTR CK1 's Loopback 1 interface is 183.215.22.1.
- D. RTR CK1 and RTR CK2 are running EBGP between the autonomous systems.
- E. RTR CK1 has forced BGP to use the loopback IP address as the source in the TCP neighbor connection.

Answer: A, B, E

Explanation:

A: The two statements that define a remote AS (autonomous) use the same AS number. Both routers must belong to the same AS and there IBGP is used.

B: The IP address of the CK1 loopback address is 147.229.1.1 as specified in the neighbor statement of the CK2 router: neighbor 147.229.1.1 remote-as 200

E: The "neighbor 183.215.22.1 update-source loopback 1" command issued at CK1 makes CK1 use this loopback for the BGP connection to CK2 . See Note 2 below.

Note 1: Syntax of neighbor command:

```
Router(config-router)#neighbor {ip-address | peer-group-name}remote-asas-number
```

This command specifies a BGP neighbor.

Note 2: For iBGP, you might want to allow your BGP connections to stay up regardless of which interface is used to reach a neighbor. To enable this configuration, you first configure a loopback interface and assign it an IP address (neighbor 183.215.22.1 update-source loopback 1). Next, configure the BGP update source to be the loopback interface (we have to assume this step - it is not indicated by the exhibit). Finally, configure your neighbor to use the address on the loopback interface (neighbor 147.229.1.1 remote-as 200).

Incorrect Answers:

C: 183.215.22.1 is the IP address of CK2 , not CK1 .

D: There is only one AS and the routers belong to that AS. Within an AS IBGP is used, not EBGP.

---

### **QUESTION 341:**

The Certkiller BGP router is configured as shown below:

```
router bgp 64000
neighbor 172.16.1.1 remote-as 64000
neighbor 10.1.1.2 remote-as 64550
```

network 200.52.1.192 mask 255.255.255.224

no synchronization

aggregate-address 200.52.1.0 255.255.255.0

The BGP peers 172.16.1.1 and 10.1.1.2 in the above configuration are active, as well as the interface with IP address 200.52.1.192. Based on this information, which statement below is true about your configuration?

- A. Router CertK has an IBGP connection with neighbor 10.1.1.2.
- B. Router CertK has an EBGP connection with the neighbor 172.16.1.1.
- C. Router CertK advertises only a route 200.52.1.0 255.255.255.0 in BGP.
- D. Router CertK advertises only a route 200.52.1.192 255.255.255.224 in BGP.
- E. Router CertK advertises both the routes to 200.52.1.0 255.255.255.0 and 200.52.1.192 255.255.255.224 in BGP.

Answer: E

Explanation:

The aggregate route, 200.52.1.0 255.255.255.0, and the more specific route, 200.52.1.192 255.255.255.224, will both be advertised. To only advertise the aggregate route the summary-only option of the aggregate-address command must be used.

Note: The aggregate-address command is used to create an aggregate, or summary, entry in the BGP table.

Incorrect Answers:

A: The neighbor 10.1.1.2 belongs to another autonomous system and is a EBGP neighbor, not an IBGP neighbor.

B: Router A and neighbor 172.16.1.1 belong to the same autonomous system and is an IBGP neighbor, not an EBGP neighbor.


C: The summary-only option of the aggregate-address command is used to only advertise the summary and not the specific routes.

D: The aggregate route, 200.52.1.0 255.255.255.0, will be advertised as well.

---

### QUESTION 342:

The Certkiller network consists of two separate autonomous systems as shown below:



You need to configure Router CK2 as a BGP route reflector and Router CK1 as the client. Assuming that Router CK3 isn't running BGP, which two of the commands

below would you enter on CK2 to satisfy your goals? (Select two)

- A. neighbor 165.50.12.1 remote-as 65100
- B. neighbor 165.50.12.2 remote-as 64000
- C. neighbor 165.50.12.1 route-reflector-client
- D. neighbor 165.50.12.2 route-reflector-client

Answer: B, D

Explanation:

B: Router CK2 (config-router)# neighbor 165.50.12.2 remote-as 64000

We configure router CK1 (165.50.12.2) as a neighbor in AS 64000.

D: Router CK2 (config-router)# neighbor 165.50.12.2 route-reflector-client

Configures the router CK2 as a BGP route reflector and configures the specified neighbor CK1 (165.50.12.2) as its client.

Incorrect Answers:

A: We must specify router CK1 as neighbor, not CK2 itself (165.50.12.1). Furthermore, we should use the local AS (64000), not the remote AS 65100.

C: We must specify router CK1 as route reflector client, not CK2 itself (165.50.12.1).

---

**QUESTION 343:**

The Certkiller network consists of a series of routers that are all configured for IBGP. Which one of the following IBGP characteristics is true?

- A. The IBGP routers must always be fully meshed.
- B. The IBGP routers can be in a different AS.
- C. The IBGP routers must be directly connected.
- D. The IBGP routers do not need to be directly connected.
- E. None of the above are true.

Answer: D

Explanation:

The IBGP routers do not have to be directly connected. The remote IBGP peers need only be reachable via a TCP connection. For example, if the network is also running an interior routing protocol such as EIGRP or OSPF, the remote IBGP router could be many hops away, as long as it is reachable via the IGP that is being used.

Incorrect Answers:

A: Using route reflectors or confederations a full mesh topology is not necessary.

B: The IBGP routers must be placed in the same AS. Peers that are in different autonomous systems are using EBGP, not IBGP.

C: The IBGP routers do not have to be directly connected.

---

**QUESTION 344:**

Which of the following statements regarding BGP peer groups are true? (Select two)

- A. Peer members inherit all options of the peer group.
- B. Peer groups can be used to simplify BGP configurations.
- C. Peer groups are optional non-transitive attributes for BGP.
- D. A peer group allows options that affect outbound updates to be overridden.
- E. A common name should be used on all routers because this information is passed between neighbors.

Answer: A, B

Explanation:

A BGP peer group is a group of BGP neighbors with the same update policies.

A: By default members of the peer group inherit all of the configuration options of the peer group.

Note: Members can also be configured to override these options

B: Peer groups are normally used to simplify router configurations when many neighbors have the same policy.

Note: BGP neighbors who share the same outbound policies can be grouped together in what is called a BGP peer group. Instead of configuring each neighbor with the same policy individually, Peer group allows to group the policies which can be applied to individual peer thus making efficient update calculation along with simplified configuration.

Incorrect Answers:

C: The only Optional non-transitive attribute in BGP is MED.

D: Does not apply.

E: Using a common router name provides no BGP benefits at all.

Reference: BGP Peer Groups, <http://www.cisco.com/warp/public/459/29.html>

---

### **QUESTION 345:**

A BGP router is configured as shown below:

```
interface ethernet 0
ip address 10.10.10.1 255.255.0.0
!
int serial 0
ip address 172.16.1.1 255.255.255.252
!
router bgp 65001
neighbor 192.168.1.1 remote-as 65002
```

Based on the above configuration, which of the following BGP statements would inject the 10.10.0.0/16 prefix into the BGP routing table?

- A. network 10.0.0.0
- B. network 10.10.0.0 mask 255.255.0.0

- C. network 10.10.10.1 mask 255.255.255.255
- D. network 10.10.10.0 mask 255.255.255.0
- E. network 10.0.0.0 mask 255.255.0.0

Answer: B

Explanation:

The /16 mask is equal to 255.255.0.0, so answer choice B matches the address and the mask. To specify the route as classless, the mask keyword should be included or the network will be summarized at the network boundary.

---

**QUESTION 346:**

Why would a systems administrator enter in the network command when configuring BGP?

- A. Local routes matching the network command are filtered from the BGP routing table.
- B. Local routes matching the network command can be installed into the BGP routing table.
- C. Routes matching the network command will be filtered from BGP routing updates.
- D. External routes matching the network command will be installed into the BGP routing table.
- E. None of the above.

Answer: B

Explanation:


Use the network router configuration command to permit BGP to advertise a network if it is present in the IP routing table.

Reference: Building Scalable Cisco Networks (Cisco Press) page 342

---

**QUESTION 347:**

The Certkiller network is depicted in the following diagram:



Which of the following command sets would you use if you want Certkiller 1 and Certkiller 2 to exchange BGP routing information? (Select two)

- A. Certkiller 2(config)#router bgp 200  
Certkiller 2(config-router)#neighbor 10.1.1.1 remote-as 100
- B. Certkiller 1(config)#router bgp 100  
Certkiller 1(config-router)#exit  
Certkiller 1(config)#interface S0  
Certkiller 1(config-if)#neighbor 10.1.1.2 remote-as 200
- C. Certkiller 1(config)#router bgp 100  
Certkiller 1(config-router)#neighbor 10.1.1.2 remote-as 200
- D. Certkiller 2(config)#router bgp 100  
Certkiller 2(config-router)#exit  
Certkiller 2(config)#interface S0  
Certkiller 2(config-if)#neighbor remote-as 100
- E. Certkiller 1(config)#router bgp 100  
Certkiller 1(config-router)#network 10.0.0.0  
Certkiller 1(config-router)#neighbor 10.1.1.2 remote-as 100
- F. Certkiller 2(config)#router bgp 200  
Certkiller 2(config-router)#network 10.0.0.0  
Certkiller 2(config-router)#network 10.1.1.1 remote-as 200

Answer: A, C

Explanation:

Choices A and C give the correct syntax for configuring EBGP peering sessions.

Incorrect Answers:

B, D: The BGP peering configurations are placed under the BGP routing process, not in interface configuration mode.

E, F: The network 10.0.0.0 command is not required here, since this is an EBGP peering session using the directly connected interface. Since each neighbor shares the 10.0.0.0 network they each already know how to reach this network.

**QUESTION 348:**

Router CK-1 is configured for BGP routing as shown below:

```
router bgp 65300
network 27.0.0.0
neighbor 192.23.1.1 remote-as 65300
```

From the perspective of router CK-1, what kind of router is the router with IP address 192.23.1.1?

- A. A peer router running IBGP
- B. A peer router running EBGP
- C. A community member running IBGP
- D. A peer group member running IBGP
- E. A peer group member running EBGP

Answer: A

Explanation:

Both the local and remote router is configured with the same autonomous system number so they are peer routers running IBGP.

---

**QUESTION 349:**

When you're configuring BGP on a Cisco router, what is true of the command "network"?

- A. Local routes matching the network command are filtered from the BGP routing table.
- B. Local routes matching the network command can be installed into the BGP routing table.
- C. Sending and receiving BGP updates is controlled by using a number of different filtering methods.
- D. The route to a neighbor autonomous system must have the correct MED applied to be installed into BGP routing table.
- E. None of the above.

Answer: B

Explanation:

The network command allows BGP to advertise an IGP route if it is already in the IP table. A matching route must exist in the routing table before the network is announced. The network command is used to permit BGP to advertise a network if it is present in the IP routing table.

---

**QUESTION 350:**

Router CK1 needs to be configured to advertise a specific network. Which of the following commands would you use if you wanted to advertise the subnet 154.2.1.0 255.255.255.0 to the EBGp neighbors on your subnet?

- A. Router (config-router)#network 154.2.1.0
- B. Router (config-router)#network 164.2.1.0
- C. Router (config-router)#network-advertise 154.2.1.0
- D. Router (config-router)#network 154.2.1.0 mask 255.255.255.0
- E. None of the above

Answer: D

Explanation:

The network command is used to specify the networks to be advertised by the Border Gateway Protocol (BGP) and multiprotocol BGP routing processes.

Syntax: network network-number [masknetwork-mask] [route-map map-name]

Mask and route-map are optional. If the

mask keyword is configured, then an exact match must exist in the routing table.

Incorrect Answers:

A: If we do not specify the subnet mask then additional networks are allowed to be advertised. The classful subnet mask of 154.2.1.0 is 255.255.0.0 - a Class B network.

B: This is using the incorrect IP address, as well as a missing subnet mask.

C: The network-advertise is an invalid command.

---

### **QUESTION 351:**

You are the administrator of a company with BGP connections to multiple ISP's. How could you configure BGP to make it favor one particular ISP for outbound traffic?

- A. Configure weight
- B. Enable route reflector
- C. Create a distribute list
- D. Enable the Longer Autonomous System path option.
- E. All of the above.

Answer: A

Explanation:

If the router learns about more than one route to the same destination, the route with the highest weight will be preferred. Weight is a Cisco BGP parameter that is local to the router. When terminating multiple ISP connections into the same router, weight can be used to affect which path is chosen for outbound traffic.

Incorrect Answers:

B: A route reflector cannot be used to influence outbound traffic. A route reflector modifies the BGP split horizon rule by allowing the router configured as the route

reflector to propagate routes learned by IBGP to other IBGP peers. This saves on the number of BGP TCP sessions that must be maintained, and also reduces the BGP routing traffic.

C: Distribute lists restrict the routing information that the router learns or advertises. By itself a distribute list cannot make routes from one ISP be preferred to routers from another ISP.

D: This choice describes ASD path pre-pending, which would be used to influence the path that incoming traffic takes, not outgoing.

---

**QUESTION 352:**

What are two solutions to overcome the full mesh requirement with iBGP? (Choose two)

- A. BGP groups
- B. BGP local preference
- C. Route reflector
- D. Confederation
- E. Aggregate-address

Answer: C, D

Explanation:

In general, all IBGP peers must be configured to be fully meshed. If they are not, then all of the IBGP routers will not have the updated information from the external BGP routers. There are two ways to overcome the scalability issues of a full IBGP mesh: route reflectors and confederations. With route reflectors, internal BGP routers peer only with the route reflector. With confederations, the AS is broken up into smaller, more manageable sub autonomous systems.

---

**QUESTION 353:**

An ISP is running a large IBPG network with 25 routers. The full mesh topology that is currently in place is inefficiently using up bandwidth from all of the BGP traffic.

What can the administrator configure to reduce the number of BGP neighbor relationships within the AS?

- A. Route reflectors
- B. Route maps
- C. Route redistribution
- D. Peer groups
- E. Aggregate addresses

Answer: A

Explanation:

In general, all IBGP peers must be configured to be fully meshed. If they are not, then all of the IBGP routers will not have the updated information from the external BGP routers. There are two ways to overcome the scalability issues of a full IBGP mesh: route reflectors and confederations. With route reflectors, internal BGP routers peer only with the route reflector, and then the route reflectors connect with each other. This can considerably reduce the number of IBGP sessions. Another solution to the scalability problem of IBGP is the use of confederations. With confederations, the AS is broken up into smaller, more manageable sub autonomous systems.

---

**QUESTION 354:**

**DRAG DROP**

Arrange the BGP attributes on the left in order or priority.

Your boss at Certkiller.com asks you to place the BGP attributes in the correct order used for determining a route.

BGP Attributes, select from these

Place here

- Originate route
- AS\_Path
- weight
- Local preference
- MED

- Place first here
- Place second here
- Place third here
- Place fourth here
- Place fifth here

Answer:

Your boss at Certkiller.com asks you to place the BGP attributes in the correct order used for determining a route.

BGP Attributes, select from these

Place here

- weight
- Local preference
- Originate route
- AS\_Path
- MED

Explanation:

How the Best Path Algorithm Works:

BGP assigns the first valid path as the current best path. It then compares the best path with the next path in list, until it reaches the end of the list of valid paths. The following

is a list of rules used to determine the best path.

1. Prefer the path with the highest WEIGHT.

Note: WEIGHT is a Cisco-specific parameter, local to the router on which it's configured.

2. Prefer the path with the highest LOCAL\_PREF. Note the following:

3. 1. Path without LOCAL\_PREF is considered as having the value set with the bgp default local-preference command, or 100 by default.

2. Prefer the path that was locally originated via a network or aggregate BGP subcommand, or through redistribution from an IGP. Local paths sourced by network or redistribute commands are preferred over local aggregates sourced by the aggregate-address command.

3. Prefer the path with the shortest AS\_PATH. Note the following:

4. 1. The AS\_CONFED\_SEQUENCE and AS\_CONFED\_SET are not included in the AS\_PATH length.

2. Prefer the path with the lowest origin type: IGP is lower than EGP, and EGP is lower than INCOMPLETE.

3. Prefer the path with the lowest multi-exit discriminator (MED). Note the following:

4. 1. This comparison is only done if the first (neighboring) AS is the same in the two paths; any confederation sub-ASs are ignored. In other words, MEDs are compared only if the first AS in the AS\_SEQUENCE is the same for multiple paths. Any preceding AS\_CONFED\_SEQUENCE is ignored.

2. If bgp always-compare-med is enabled, MEDs are compared for all paths. This option needs to be enabled over the entire AS, otherwise routing loops can occur.

3. If bgp bestpath med-confed is enabled, MEDs are compared for all paths that consist only of AS\_CONFED\_SEQUENCE (paths originated within the local confederation).

4. Paths received from a neighbor with a MED of 4,294,967,295 will have the MED changed to 4,294,967,294 before insertion into the BGP table.

5. Paths received with no MED are assigned a MED of 0, unless bgp bestpath missing-as-worst is enabled, in which case they are assigned a MED of 4,294,967,294.

6. The bgp deterministic med command can also influence this step as demonstrated in the How BGP Routers Use the Multi-Exit Discriminator for Best Path Selection.

7. Prefer external (eBGP) over internal (iBGP) paths. If bestpath is selected, go to Step 9 (multipath).

Note: Paths containing AS\_CONFED\_SEQUENCE and AS\_CONFED\_SET are local to the confederation, and therefore treated as internal paths. There is no distinction between Confederation External and Confederation Internal.

8. Prefer the path with the lowest IGP metric to the BGP next hop. Continue, even if bestpath is already selected.

9. Check if multiple paths need to be installed in the routing table for BGP Multipath. Continue, if bestpath is not selected yet.

10. When both paths are external, prefer the path that was received first (the oldest one). This step minimizes route-flap, since a newer path will not displace an older one, even if it would be the preferred route based on the next decision criteria (Steps 11, 12, and 13). Skip this step if any of the following is true:

1. The bgp best path compare-routerid command is enabled.

Note: This command was introduced in Cisco IOS(r) Software Releases 12.0.11S,

12.0.11SC, 12.0.11S3, 12.1.3, 12.1.3AA, 12.1.3.T, and 12.1.3.E.

10. 1. The router ID is the same for multiple paths, since the routes were received from the same router.

2. There is no current best path. An example of losing the current best path occurs when the neighbor offering the path goes down.

3. Prefer the route coming from the BGP router with the lowest router ID. The router ID is the highest IP address on the router, with preference given to loopback addresses. It can also be set manually using the `bgp router-id` command.

Note: If a path contains route-reflector (RR) attributes, the originator ID is substituted for the router ID in the path selection process.

12. If the originator or router ID is the same for multiple paths, prefer the path with the minimum cluster list length. This will only be present in BGP route-reflector environments. It allows clients to peer with RRs or clients in other clusters. In this scenario, the client must be aware of the RR-specific BGP attribute.

13. Prefer the path coming from the lowest neighbor address. This is the IP address used in the BGP neighbor configuration, and corresponds to the remote peer used in the TCP connection with the local router.

---

### QUESTION 355:

The Certkiller BGP routing table consists of the following network routes:

```
192.168.12.0/24
192.168.13.0/24
192.168.14.0/24
192.168.15.0/24
```

What is the correct command to summarize these prefixes into a single summary prefix of 192.168.12.0/22 while also allowing for the advertisement of the more specific prefixes?

- A. `network 192.168.12.0 mask 255.255.252.0`
- B. `network 192.168.12.0 mask 0.0.3.255`
- C. `network 192.168.12.0`
- D. `aggregate-address 192.168.12.0 255.255.252.0`
- E. `aggregate-address 192.168.12.0 255.255.252.0 summary-only`
- F. `aggregate-address 192.168.12.0 255.255.252.0 as-set`


Answer: D

Explanation:

To summarize BGP prefixes into one aggregated route, use the "aggregate-address" command. When used alone, this will advertise the aggregate route, along with the individual specific routing entries. To advertise only the aggregated route, use the "summary-only" keyword, as specified in choice E.

**QUESTION 356:**

The Certkiller BGP network is displayed below:




Certkiller 3 is the hub router and Certkiller 1 and Certkiller 2 are the spokes. There are no virtual circuits between the spoke locations. Each router is in a separate AS. What is needed to successfully route traffic to the 11.11.11.0/24 network from Certkiller 1?

- A. The neighbor 10.10.10.1 next-hop-self command on Certkiller 1.
- B. The neighbor 10.10.10.1 next-hop-self command on Certkiller 2.
- C. The neighbor 10.10.10.1 next-hop-self command on Certkiller 3.
- D. Nothing is required. This is the default behavior on this topology

Answer: C

Explanation:

The following example illustrates the issue:

**BGP Next Hop (NBMA)**

If the common media as you see in the shaded area above is a frame relay or any NBMA cloud then the exact behavior will occur as if we were connected via Ethernet. RTC will advertise 180.20.0.0 to RTA with a next hop of 170.10.20.3.

The problem is that RTA does not have a direct PVC to RTD, and cannot reach the next hop. In this case routing will fail.

In order to remedy this situation a command called next-hop-self is created.

The next-hop-self Command

Because of certain situations with the next hop as we saw in the previous example, a command called next-hop-self is created. The syntax is:

```
neighbor {ip-address|peer-group-name} next-hop-self
```

The next-hop-self command allows us to force BGP to use a specified IP address as the next hop rather than letting the protocol choose the next hop.

In the previous example, the following configuration solves our problem:

```
RTC#
```

```
router bgp 300
```

```
neighbor 170.10.20.1 remote-as 100
```

```
neighbor 170.10.20.1 next-hop-self
```

```
RTC advertises 180.20.0.0 with a next hop = 170.10.20.2
```

```
Reference:http://www.cisco.com/warp/public/459/bgp-toc.html#bgpnexthop
```

---

### QUESTION 357:

The following was seen on router CK1 :

```
Certkiller1#show ip bgp summary
BGP table version is 8, main routing table version 8
4 network entries (8/12 paths) using 832 bytes of memory
5 BGP path attribute entries using 576 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
2 received paths for inbound soft reconfiguration
Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
10.1.1.1 4 50001 80 81 8 0 0 04:15:10 2
10.2.2.2 4 50002 79 81 0 0 0 00:00:15 Active
10.3.3.3 4 50003 82 82 0 0 0 02:00:00 Idle
```

Based on the show ip bgp summary output in the exhibit, which two statements are true? Select two?

- A. The BGP session to the 10.1.1.1 neighbor is established.
- B. The BGP session to the 10.2.2.2 neighbor is established.
- C. The BGP session to the 10.3.3.3 neighbor is established.
- D. The router is attempting to establish a BGP peering to the 10.1.1.1 neighbor.
- E. The BGP session to the 10.3.3.3 neighbor is established, but the router has not received any BGP routing updates from the 10.3.3.3 neighbor.
- F. The router is attempting to establish BGP peering with the 10.2.2.2 neighbor.

Answer: A, F

Explanation:

In BGP, established valid neighbors will have a number under the state/pfxRcd heading. Neighbors that show an active state mean that the router is actively trying to establish a peering relationship. Idle peers are not active, and not actively trying to peer. In this case, only the peer to 10.1.1.1 is currently operational with two route prefixes received, while the peer to 10.2.2.2 is trying to become established.

**QUESTION 358:**

You want to hide some of your internal IP subnets from outside networks. By what means can you conceal the details of your IP addressing scheme to the outside world?

- A. Subnetting
- B. Supernetting
- C. Challenge Handshake Protocol
- D. Usernames and passwords

Answer: B

Explanation:

In large internetworks, hundreds or even thousands of networks can exist. In these environments, it is often not desirable for routers to maintain all these routes in their routing table. Route summarization (also called route aggregation or supernetting) can reduce the number of routes that a router must maintain because it is a method of representing a series of network numbers in a single summary address. By advertising only one large supernet to the outside world, the details of your IP network scheme can remain hidden.

---

**QUESTION 359:**

**SIMULATION**

The Certkiller network will be using a new ISP for their Internet connection. The new Internet provider will be allocating a new registered class C IP address subnet to use. In order to not have to change out the entire internal IP addressing scheme on your network, you plan to use NAT to translate all of the internal IP address to the new IP addresses that are being assigned to Certkiller . You have been assigned the task of making changes so that the following requirements are met:

1. Maintain the current IP addressing scheme
2. Configure Certkiller 's router for network address translation (NAT) so all the internal computers use the same external IP address
3. Configure a static translation so Certkiller employees can access their email from the internet.

Privileged mode password: Certkiller

The IP Addresses are shown below:

Name: Certkiller NAT

SO 192.168.15.1/24

E0 10.100.5.1/24

E-mail Server's External Address


192.168.15.5/24

E-mail Server's Internal Address

10.100.5.5/24

## 642-801

To configure the router click on a host icon that is connected to a router by a serial console cable.



Answer:

Explanation:

```
Certkiller NAT#Conf t
Certkiller NAT(config)#Access-list 5 permit 10.100.5.0 0.0.0.255
Certkiller NAT(config)#Ip nat pool lan 192.168.15.1 192.168.15.1 netmask 255.255.255.0
Certkiller NAT(config)#Ip nat inside source list 5 pool lan overload
Certkiller NAT(config)#Ip nat inside source static 10.100.5.5 192.168.15.5
Certkiller NAT(config-if)#Int S0
Certkiller NAT(config-if)#Ip nat outside
Certkiller NAT(config-if)#Int E0
Certkiller NAT(config-if)#Ip nat inside
Certkiller NAT(config-if)#exit
Certkiller NAT(config)#exit
Certkiller NAT#copy run start
```


Note; alternative suggestion:

```
Certkiller Nat(Config)# ip nat inside source static 10.100.5.5 192.168.15.1
Certkiller Nat(Config)#ip nat inside source list 10 interface S0 overload
Certkiller Nat(Config)#access-list 10 permit 10.100.5.0 0.0.0.255
Certkiller Nat(Config)#int S0/0
Certkiller Nat(Config)#ip nat outside
Certkiller Nat(Config)#exit
Certkiller Nat(Config)#int e0/0
Certkiller NAT(Config)#ip nat outside
Certkiller Nat(Config)#exit
```

---

### **QUESTION 360:**

The Certkiller network is using RIP and IGRP for routing as shown below:



You have been assigned the task of configuring the routers to meet the following requirements:

1. In the unlikely event that a link goes down, you have to ensure a backup link exists
  2. The proper metrics must be distributed through the IGRP portion of the network
- Based on this information, which of the following should you do? (Select two)

- A. Apply the distance command to Main
- B. Apply the distance command to CK2
- C. Apply the distribute-list command to Main
- D. Apply the distribute-list command to CK2

Answer: B, C

Explanation:

B: Most routing protocols, including IGRP and RIP, use the distance command to modify the administrative distance. We need to change the metric of the RIP routes learned by CK2 .

C: To control the advertising and processing of routes in routing updates the distribute-list command is used.

Incorrect Answers:

A: We change the administrative distance of RIP routes on CK2 .


D: We should apply the distribute-list command to Main, not to CK2 .

### QUESTION 361:

SIMULATION

Network topology exhibit

## 642-801



Certkiller .com is changing ISPs. As a result, they will need to install an Email server and utilize network address translation (NAT) for the internal network.

Certkiller .com does not want to change the IP Addresses on all of the internal routers and servers. The new ISP, Foo Services, will allocate a registered class C address for Certkiller to use. The internal IP Address scheme will remain the same. Configure the router to provide NAT so that all internal Certkiller PCs will use the single external IP Address assigned to the router interface. Configure a static translation so that Certkiller .com's Email server will be accessible from the Internet.

Name: Certkiller NAT

S0: 192.168.17.1/24

E0: 10.0.216.1/24

Secret Password: Certkiller

Answer:

Explanation:

```
Certkiller 2> en
```

```
Certkiller 2# config t
```

```
Certkiller nat(config)#access-list 2 deny host 10.0.216.5
```

```
Certkiller nat(config)#access-list 2 permit 10.0.216.0 0.0.0.255
```

```
Certkiller nat(config)#ip nat pool nat 192.168.17.1 192.168.17.1 netmask 255.255.255.0
```

```
Certkiller nat(config)#ip nat inside source list 2 pool nat overload
```

```
Certkiller nat(config)#ip nat inside source static 10.0.216.5 192.168.17.5
```

```
Certkiller nat(config)#interface ethernet 0
```

```
Certkiller nat(config-if)#ip address 10.0.216.1 255.255.255.0
```

```
Certkiller nat(config-if)#ip nat inside
```

```
Certkiller nat(config-if)#interface serial 0
```

```
Certkiller nat(config-if)#ip address 192.168.17.1 255.255.255.0
```

```
Certkiller nat(config-if)#ip nat outside
```

```
Certkiller nat(config-if)# ^ z
```

```
Certkiller nat#copy running-config startup-config
```

Note, variation:

Name: Certkiller NAT

S0: 192.168.212.1/24

E0: 10.0.243.1/24

Secret Password: Certkiller  
Web Server's External Address  
192.168.212.5/24  
Web Server's Internal Address  
10.0.243.5/24

---

**QUESTION 362:**

One of the virtues of a Link-State routing protocol is its hierarchical design. What advantages does this equate to in network design? (Select two)

- A. It allows link-state protocols to support VLSM.
- B. It allows them to support address summarization.
- C. It reduces the need to flood LSAs to all devices in the internetwork.
- D. Routers are no longer required to keep track of the topology of the entire autonomous system.

Answer: C, D

Explanation:

Hierarchical design can limit the requirement to notify all devices. The use of areas restricts the flooding to the logical boundary of the area rather than to all devices in the OSPF domain. Another advantage of hierarchical network design is that the entire network is divided into multiple smaller sections, so that much of the routing information is kept within a small area. In a non-hierarchical network, each router is required to maintain routing information for all other routers within the entire network.

Incorrect Answers:

- A: Although link state protocols do indeed support VLSM, it is not due to a hierarchical network design. For example, RIP version 2 supports VLSM, but it is not hierarchical.
- B: Although address summarization can indeed be configured on routers configured in a hierarchical design, it is an independent function and is not related to the fact that a tiered network design is being used.

---

**QUESTION 363:**

In regards to the three layered hierarchical network design, which of the following are responsibilities of the distribution layer? (Choose 2)

- A. Reliable transport structure
- B. Route redistribution
- C. Optimized transport structure
- D. Address aggregation
- E. Unauthorized entry access control lists

Answer: B, D

Explanation:

The distribution layer is responsible for what enters and exits a network, including the consolidation of traffic from multiple subnets into a core connection. Route redistribution and address aggregation happen at this layer.

Incorrect Answers:

A, C: This is a function of the core layer.

E. This is a function of the access layer.

---

**QUESTION 364:**

Which of the following are advantages of implementing a hierarchical IP addressing scheme when designing a large network? (Select two)

- A. Smaller routing tables
- B. Efficient address allocation
- C. Translation of private addresses
- D. Support for link-state routing protocols

Answer: A, B


Explanation:

Hierarchical IP addressing uses the addresses more conservatively, by planning for the correct use of subnet masks and route summarization. A hierarchical network design facilitates the use of IP routing summarization, which will reduce the size of the routing tables.

---

**QUESTION 365:**

The Certkiller frame relay network is shown in the diagram below:



Based on the information given in this diagram, which of the following are true??

- A. The network is using a two-layer full-mesh hierarchical design.
- B. The network is using a two-layer hub-and-spoke hierarchical design.
- C. To improve scalability, route summarization at Routers CK1 , CK2 , CK3 , CK4 , CK5 ,

and CK6 should be performed.

D. Routers A and B are the distribution layer routers.

E. Routers CK1 , CK2 , CK3 , CK4 , CK5 and CK6 are the access layer routers.

F. To improve security, packet filtering that uses ACLs at Routers A and B must be implemented.

Answer: C

Explanation:

Routers CK1 , CK2 , CK3 , CK4 , CK5 and CK6 are the distribution layer routers. Route summarization should occur at the distribution layer.

Incorrect Answers:

A: The network shown here is using a 3 level tiered approach: Core, Distribution, and Access.

B: Three layers are used, not two.

D: Routers A and B are the Core routers here.

E: These routers make up the Distribution layer.

---

### **QUESTION 366:**

Certkiller is using Metro Ethernet as their Core backbone network. What does VLAN tunneling in the Metro Ethernet environment do for a customer?

A. Extends their logical network topology across wide geographic networks.

B. Renumbers their LANs.

C. Provides combined wavelength routing.

D. Translates their VLANs at the service provider edge.

Answer: D

Explanation:

Virtual private networks (VPNs) provide enterprise-scale connectivity on a shared infrastructure, often Ethernet-based, with the same security, prioritization, reliability, and manageability requirements of private networks. Tunneling is a feature designed for service providers who carry traffic of multiple customers across their networks and are required to maintain the VLAN and Layer 2 protocol configurations of each customer without impacting the traffic of other customers.

Business customers of service providers often have specific requirements for VLAN IDs and the number of VLANs to be supported. The VLAN ranges required by different customers in the same service-provider network might overlap, and traffic of customers through the infrastructure might be mixed. Assigning a unique range of VLAN IDs to each customer would restrict customer configurations and could easily exceed the IEEE802.1Q specification VLAN limit of 4096.

Using VPN tunneling features, service providers can use a single VLAN to support customers who have multiple VLANs via translation at the service provider edge.

**QUESTION 367:**

You are trying to determine the best routing protocol to use for the large Certkiller network. Which routing protocols should you avoid when deploying a large network? (Select two)

- A. IGRP
- B. OSPF
- C. EIGRP
- D. RIP v.2
- E. RIP v.1

Answer: D, E

Explanation:

Both RIP version 1 and version 2 should only be used in relatively small networks. A large network would be congested by RIP broadcasts. These frequent broadcasts contain the entire routing table, and could saturate links.

---

**QUESTION 368:**

What are some of the features that are important for designing large, scalable networks? (Select three)

- A. A tiered network design model.
- B. Sufficient memory on the router.
- C. Multiple EIGRP autonomous systems.
- D. Good address space allocation scheme.

Answer: A, B, D

Explanation:

A: A tiered network design model such as Core, Distribution, and Access is also needed for large networks.

B: Sufficient capacity of the routers, in particular the memory, is required for large networks.

D: Good allocation of address space is required- Each region should have a unique address space so route summarization is possible.

Incorrect Answers:

C: Configuring multiple EIGRP autonomous systems can add extra overhead and it provides no additional benefits.

---

**QUESTION 369:**

Two of the following routing protocols require a tiered, hierarchical topology. Which ones are they? (Select two)

- A. IS-IS
- B. IGRP
- C. OSPF
- D. EIGRP
- E. RIP v2

Answer: A, C

Explanation:

IS-IS and OSPF both require a hierarchical topology. The use of multiple areas and a core backbone area means that they are hierarchical by nature.

---

**QUESTION 370:**

What has become the leading Interior Routing Protocol (IGP) in use by Internet service providers?

- A. IS-IS
- B. OSPF
- C. EIGRP
- D. RIPv2
- E. BGP4

Answer: A

Explanation:

In recent years, the IS-IS routing protocol has become increasingly popular, with widespread use among Service Providers. It is a link state protocol, which enables very fast convergence with large scalability. It is also a very flexible protocol and has been extended to incorporate leading edge features such as MPLS Traffic Engineering. It is also chosen because it is an IETF standard based protocol, rather than a proprietary protocol such as IGRP or EIGRP.

Reference:

[http://www.cisco.com/en/US/tech/CK365/CK381/technologies\\_white\\_paper09186a00800a3e6f.shtml](http://www.cisco.com/en/US/tech/CK365/CK381/technologies_white_paper09186a00800a3e6f.shtml)

---

**QUESTION 371:**

Two of the following routing protocols support load balancing over unequal cost links. Which ones are they? (Select two)

- A. IGRP
- B. OSPF
- C. EIGRP
- D. RIP v2
- E. RIP v1

Answer: A, C

Explanation:

IGRP and EIGRP support unequal cost path load balancing, which is known as variance. OSPF, RIP v1 and RIP v2 do not support this.

Reference: <http://www.cisco.com/warp/public/103/19.html>

---

**QUESTION 372:**

You are required to choose a routing protocol that supports variable length subnet masks (VLSM) for your network. Which of the following would meet this requirement? (Choose three)

- A. IS-IS
- B. IGRP
- C. OSPF
- D. EIGRP
- E. RIP v1

Answer: A, C, D

Explanation:

IS-IS, EIGRP, OSPF, RIP V2 and static routes all support VLSM.

Incorrect Answers:

B: IGRP is the predecessor to EIGRP and does not support VLSM.

E: RIP version 1 is a distance vector routing protocol that does not support VLSM.

---

**QUESTION 373:**

Cisco routers automatically perform route summarization for various protocols. Which three of the following routing protocols are they? (Select three)

- A. IS-IS
- B. IGRP
- C. OSPF
- D. EIGRP
- E. RIP v.1

Answer: B, D, E

Explanation:

Sending route summaries - Routing information advertised out an interface is automatically summarized at major (classful) network address boundaries by RIP, IGRP, and EIGRP. Specifically, this autonomous summarization occurs for those routes whose classful network address differs from the major network address of the interface to which

the advertisement is being sent.

Reference: Building Scalable Cisco Networks (Cisco Press) page 79

---

**QUESTION 374:**

In some situations, static routing is preferred over dynamic routing. In which situations would it be better to use a static routing protocol? (Select two)

- A. Medium to large networks with redundant paths.
- B. Networks with a single entry point.
- C. Low maintenance routing is required.
- D. Highly adaptable networks.
- E. High degree of control in path selection is required.

Answer: B, E

Explanation:

Networks with a single entry point only have one possible path to take for all networks, so it generally makes the most sense to simply add a single static default route out of that link. With static routes, there is not routing protocol overhead. Static routes may also be preferred when a high degree of control is required, since the entries are static and can not be overridden by a dynamic protocol.

Incorrect Answers:

- A: Larger networks with multiple links are best suited for dynamic routing protocols, since there would be too much administrative overhead with the use of static routes.
  - C: Static routes require the most administrative maintenance, since everything is done manually.
  - D: Statically routed networks are not adaptable, since static routes never change, regardless of any changes to the network topology. Only dynamic protocols can adapt to changing network conditions.
- 

**QUESTION 375:**

One routing protocol considers different path attributes based on preconfigured network policies. Which of the protocols below is it?

- A. EIGRP
- B. OSPF
- C. RIPv2
- D. BGP4
- E. IS-IS
- F. IGRP

Answer: D

Explanation:

BGP uses numerous attributes. They are: origin, AS\_path, next hop, MED, local preference, atomic aggregate, aggregator, community, originator ID, cluster list, and weight. Some of these values can be adjusted to give the administrator flexibility in creating a network routing policy. BGP differs in other routing protocols in that it is used for external use, so the attributes of paths are considered, rather than just the individual routes.

---

**QUESTION 376:**

You are the network administrator at Certkiller . You are in the planning stages of upgrading the Certkiller network, and need to decide which protocol to use. You need to choose a protocol that will meet the following goals:

1. Supports classless routing
2. Supports VLSM
3. Does not rely on TCP/IP to exchange routing information
4. Performs auto-summarization by default
5. Allows for manual route summarization

Which routing protocol should you implement on the Certkiller network?

- A. BGP
- B. EIGRP
- C. OSPF
- D. IS-IS
- E. IGRP
- F. RIPv2

Answer: B

Explanation:

EIGRP is a classless protocol that supports VLSM and performs auto-summarization of subnets by default. The subnet routes are summarized into a single network number aggregate. In addition, IP-EIGRP will allow aggregation on any bit boundary in an IP address and can be configured on a network interface for additional granularity. EIGRP also supports automatic summarization of network addresses at major network borders. There is an automatic redistribution mechanism used so IGRP routes are imported into EIGRP and vice versa. Since the metrics for both protocols are directly translatable, they are easily comparable as if they were routes that originated in their own AS.

---

**QUESTION 377:**

Which routing protocol: (1) is intended to support large routing domains consisting of combinations of many media types, (2) may be used as an IGP, and (3) supports multiple routed protocols in an integrated manner rather than "Ships in the Night" method?

- A. EIGRP

- B. OSPF
- C. IS-IS
- D. BGP
- E. None of the above

Answer: A

Explanation:

EIGRP is a very scalable Interior Routing Protocol (IGP) that supports large networks. In addition to support for IP, EIGRP can be used for routing IPX and Appletalk networks.

Incorrect Answers:

B, C: Although OSPF and IS-IS are link state protocols that are used as IGPs and are highly scalable, they do not provide support for routed networks such as those used by Novell and Apple.

D: BGP only supports IP, and is generally used only for external routing.

---

**QUESTION 378:**

DRAG DROP

Match the correct routing protocol on the left to its characteristics on the right side.

Note that not all answer choices will be used.

|       |       |                                  |            |
|-------|-------|----------------------------------|------------|
| OSPF  | BGP   | Uses NET to identify the router  | Place here |
| RIPv1 | MPLS  | Maintains feasible successor(s)  | Place here |
| IGRP  | EIGRP | Uses path attributes             | Place here |
| ISIS  | RIPv  | Supports special areas like NSSA | Place here |

Answer:

|       |      |                                  |       |
|-------|------|----------------------------------|-------|
|       |      | Uses NET to identify the router  | ISIS  |
| RIPv1 | MPLS | Maintains feasible successor(s)  | EIGRP |
| IGRP  |      | Uses path attributes             | BGP   |
|       | RIPv | Supports special areas like NSSA | OSPF  |

---

**QUESTION 379:**

ODR is being used as the routing protocol for a new Certkiller network. Which network topology is best suited for ODR?

- A. highly redundant
- B. fully meshed
- C. dual self-healing ring
- D. hub and spoke
- E. NBMA

Answer: D

Explanation:

On Demand Routing (ODR) is a mechanism that provides minimum-overhead IP routing for stub sites. The overhead of a general dynamic routing protocol is avoided, without incurring the configuration and management overhead of using static routing.

A stub router is the peripheral router in a hub and spoke network topology. Stub routers commonly have a WAN connection to the hub router and a small number of LAN network segments (stub networks) that are connected directly to the stub router. To provide full connectivity, the hub routers can be statically configured to know that a particular stub network is reachable via a specified access router. However, if there are multiple hub routers, many stub networks, or asynchronous connections between hubs and spokes, the overhead required to statically configure knowledge of the stub networks on the hub routers becomes too great.

ODR simplifies installation of IP stub networks in which the hub routers dynamically maintain routes to the stub networks. This is accomplished without requiring the configuration of an IP routing protocol at the stub routers. With ODR, the stub advertises IP prefixes corresponding to the IP networks that are configured on its directly connected interfaces. Because ODR advertises IP prefixes, rather than IP network numbers, ODR is able to carry Variable Length Subnet Mask (VLSM) information. Once ODR is enabled on a hub router, the router begins installing stub network routes in the IP forwarding table. The hub router can also be configured to redistribute these routes into any configured dynamic IP routing protocols. IP does not need to be configured on the stub router. With ODR, a router is automatically considered to be a stub when no IP routing protocols have been configured on it.

---

**QUESTION 380:**

Your company, Certkiller Inc, has a central office and a remote branch office. Each site has a separate autonomous system (AS) as well as a separate Internet connection. In this network, what is the best method to route Internet traffic.

- A. IGP routing
- B. BGP routing
- C. Configure "ip route prefix mask"
- D. Configure "ip default-gateway ip address"
- E. Configure IS-IS

Answer: B

Explanation:

Border Gateway Protocol (BGP) is used to exchange routing information for the Internet and is the protocol used between Internet service providers (ISP). In order to take full advantage of multiple Internet connections, as well as multiple Internet providers, BGP is the best choice to ensure full redundancy.

Reference: [http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito\\_doc/bgp.htm](http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm)

Incorrect Answers:

A: Customer networks usually employ an Interior Gateway Protocol (IGP) such as RIP or OSPF for the exchange of routing information within their networks.

C, D: Using statically assigned default routes and default gateways is a simplistic approach that works well in many situations, but will not provide for automatic failover and does not provide the routing options that BGP provides.

E: IS-IS is used exclusively as an IGP.

---

**QUESTION 381:**

Which three network issues indicate a need to implement QoS features? (Choose three)

- A. mis-routed packets
- B. excess jitter
- C. delay of critical traffic
- D. packet loss due to congestion
- E. data link layer broadcast storms
- F. ftp connections unsuccessful

Answer: B, C, D

Explanation:

Quality of service (QoS) is the set of techniques designed to manage network resources. QoS refers to the capability of a network to provide better service to selected network traffic over various LAN and WAN technologies. The primary goal of QoS is to provide flow priority, including dedicated bandwidth, controlled jitter and latency (required by some interactive and delay-sensitive traffic), and improved loss characteristics. This is normally done to prioritize traffic and ensure the timely delivery of critical traffic.

Incorrect Answers:

A. This would indicate a routing problem, and normally results in total connectivity loss. QoS is designed to make adjustments for periodic network congestion.

E. This problem would need to be resolved by segmenting the network either through the use of VLANs or by segmenting the LAN with a router.

F. FTP connections are not typically considered mission critical, and due to the nature of FTP connections may be slow during congestion, but still successful. Lost FTP connections indicate a basic network connectivity issue or server problem.

---

**QUESTION 382:**

What is the most valid reason to carry voice traffic on a separate VLAN in a switched network providing both data and IP telephony services?

- A. IP phones require inline power and must be in separate VLAN to receive inline power.
- B. IP telephony applications require prioritization over other traffic as they are more

delay sensitive.

- C. IP phones can only receive IP addresses through DHCP if they are in separate VLAN.
- D. The CDP frames from the IP phone can only be recognized by the switch if the phone is in an auxiliary vlan.

Answer: B

Explanation:

When designing a VOIP solution, it is always wise to carry the voice traffic on a separate VLAN so that the data traffic does not interfere with the voice traffic. Voice traffic is very delay sensitive and steps are needed to ensure that this traffic is delivered in a timely matter.

Incorrect Answers:

- A. Some IP phones can be plugged into AC adapters, so power over Ethernet ports are not required. Besides, a separate VLAN does not need to be configured for any stations to plug into the POE ports.
- C. This statement is false, since the IP phones reach their DHCP servers in the same way that computer stations do.
- D. CDP frames can be carried over any layer 2 transport type and over any VLAN.

---

**QUESTION 383:**

What protocol specified by RFC 1256 will allow an enabled IP host to find a new router when a router becomes unavailable?

- A. IRDP
- B. SNMP
- C. HSRP
- D. VRRP

Answer: A

Explanation:

Some newer IP hosts use ICMP Router Discovery Protocol (IRDP) to find a new router when a route becomes unavailable. A host that runs IRDP listens for hello multicast messages from its configured router and uses an alternate router when it no longer receives those hello messages. The default timer values of IRDP mean that it's not suitable for detection of failure of the first hop. The default advertisement rate is once every 7 to 10 minutes, and the default lifetime is 30 minutes.

ICMP Router Discovery Protocol (IRDP) is defined by RFC 1256 and is an extension to ICMP that allows routers to notify hosts of their presence.

The full RFC can be found in the link below:

Reference: <http://rfc.net/rfc1256.html>

---

**QUESTION 384:**

Which three QoS mechanisms can be configured to improve VoIP quality on a converged network? (Choose three)

- A. The use of a queuing method that will give VoIP traffic strict priority over other traffic.
- B. The use of RTP header compression for the VoIP traffic.
- C. The proper classification and marking of the traffic as close to the source as possible.
- D. The use of 802.1QinQ trunking for VoIP traffic.
- E. The use of WRED for the VoIP traffic.

Answer: A, C, E

Explanation:

Using strict priority, voice traffic will always be given preference over other data traffic types.

As a general rule, when using any QoS mechanisms it is best to classify the traffic as close to the source as possible, so that the voice traffic can be prioritized end to end.

The WRED algorithm provides for congestion avoidance on network interfaces by providing buffer management, and allowing Transmission Control Protocol (TCP) traffic to throttle back before buffers are exhausted. This helps avoid tail drops, and global synchronization issues, thereby maximizing network utilization and TCP-based application performance. The use of WRED can be used to prevent network congestion, while also prioritizing the VOIP traffic types.

---

**QUESTION 385:**

The Certkiller network is in the process of migrating the IP address scheme to use IPv6. Which of the following address types are associated with IPv6? (Select three)

- A. Unicast
- B. Private
- C. Broadcast
- D. Public
- E. Multicast
- F. Anycast

Answer: A, E, F

Explanation:

IP version 6 introduced the concept of anycasts. The three IP address types used for IPv6 are:

1. Unicast- An IPv6 unicast address is an identifier for a single interface, on a single node. A packet that is sent to a unicast address is delivered to the interface identified by that address.

2. Anycast- An anycast address is an address that is assigned to a set of interfaces that typically belong to different nodes. A packet sent to an anycast address is delivered to the closest interface-as defined by the routing protocols in use-identified by the anycast address.

3. Multicast -

An IPv6 multicast address is an IPv6 address that has a prefix of FF00::/8 (11111111).

An IPv6 multicast address is an identifier for a set of interfaces that typically belong to different nodes.

Reference:

[http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products\\_feature\\_guide\\_chapter09186a0080110dd2.h](http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide_chapter09186a0080110dd2.h)

---

**QUESTION 386:**

What number is a valid representation for the 200F:0000:AB00:0000:0000:0000/56 IPv6 prefix?

- A. 200F:0:0:AB/56
- B. 200F:0:AB00::/56
- C. 200F::AB00/56
- D. 200F:AB/56

Answer: B

Explanation:

Due to the method of allocating certain styles of IPv6 addresses, it will be common for addresses to contain long strings of zero bits. In order to make writing addresses containing zero bits easier a special syntax is available to compress the zeros. The use of "::" indicates multiple groups of 16-bits of zeros. The "::" can only appear once in an address. The "::" can also be used to compress the leading and/or trailing zeros in an address.

For example the following addresses:

1080:0:0:0:8:800:200C:417A a unicast address

FF01:0:0:0:0:0:0:43 a multicast address

0:0:0:0:0:0:0:1 the loopback address

0:0:0:0:0:0:0:0 the unspecified addresses

may be represented as:

1080::8:800:200C:417A a unicast address

FF01::43 a multicast address

::1 the loopback address

:: the unspecified addresses

In our example, the trailing zero's can be compressed into ::, giving us 200F:0:AB00:: as the address.

---

**QUESTION 387:**

You have the binary IP address, {11000000.10100100.11000000.00000001 }. What class of IP address is it, and what is a characteristic of that class?

- A. It is a Class B public address.
- B. It is a Class C public address.
- C. It is a Class B private address
- D. It is a Class C private address.
- E. It is a Class D experimental address.

Answer: B

Explanation:

11000000.10100100.11000000.00000001 = 192.164.192.1 = A public Class C address

Incorrect Answers:

- A, C: Class B addresses begin with a "10" in the first octet.
- D. Class C private (RFC 1918) addresses are in the IP range 192.168.0.0-192.168.255.255

Reference: Building Scalable Cisco Networks (Cisco Press) page 66

---

### **QUESTION 388:**

You are the network administrator at Certkiller . Certkiller is assigned a network of 200.10.5.0/24. You want to create 5 subnets from the assigned address. Which subnet mask will give you 5 subnets and the most possible hosts per subnet?

- A. /21
- B. /23
- C. /26
- D. /27
- E. /29

Answer: D

Explanation:

We need 5 subnets from the given address 200.10.5.0/24. It means we will do 3 bits of subnetting which will provide us 8 subnets as  $2^N = 2^3 = 8$  subnets where n=number of subnet bits. So, we will have a total of  $24+3 = 27$  network bits.

---

### **QUESTION 389:**

You are the network engineer at Certkiller . Certkiller has a branch office with five devices that require public IP addresses. Which of the following subnet masks will provides the most efficient allocation of addresses?

- A. 255.255.255.128
- B. 255.255.255.248
- C. 255.255.255.252
- D. 255.255.255.254

Answer: B

Explanation:

Requirement = 5 public IP addresses to connect 5 devices  
3-bits of subnetting in the host portion will provide  $2^3 = 8$  addresses. After we subtract 2 of these addresses, for the network and broadcast IP address, we are still left with 6 useable IP addresses for each network.

---

**QUESTION 390:**

Which three are characteristics of IPv6? Select three.

- A. An IPv6 address is 128 bits long.
- B. An IPv6 header is 20 bits long.
- C. An IPv6 header contains the next header field.
- D. An IPv6 header contains the next protocol field.
- E. IPv6 routers send RA messages.
- F. An IPv6 header contains the header checksum field.

Answer: A, C, E

Explanation:

In IPv6, the 128 bit IP address is separated to two: network prefix, which identifies network, and interface ID, which identifies a node (interface). Interface ID is configured by the node on its own, and prefix is notified by the network (usually router). These two are combined to form an IPv6 address. IPv6 packet headers contain many of the fields found in IPv4 packet headers; some of these fields have been modified from IPv4. The next header field is the next extension header to examine. This was previously the protocol field in IPv4.

RA stands for router advertisements, used by IPv6 routers for the purpose of automatic address configuration.

---

**QUESTION 391:**

Which address type does the IPv6 address FF05:0:0:0:0:0:0:2 specify?

- A. unspecified
- B. aggregable global unicast
- C. link local
- D. site local unicast
- E. multicast

Answer: E

Explanation:

All IPv6 multicast addresses start with the first 8 bits of the address set to 1. Thus all multicast addresses start with the hexadecimal notation FF (1111 1111). The multicast range is as follows:

FF00::/8

FFFF::/8

The second octet, following to the first octet of FF, identifies both the scope and the lifetime of the multicast address. In this way, IPv6 has millions of group multicast addresses to use in current and emerging technologies.

Reference: CCNP Self-Study CCNP BSCI Exam Certification Guide p.112

---

**QUESTION 392:**

Which command must be globally enabled on a Cisco router to support IPv6?

- A. ip routing ipv6
- B. ipv6 unicast-routing
- C. ipv6 routing
- D. ip classless
- E. ipv6 cef

Answer: B

Explanation:

To support the processing and forwarding of IP version 6 packets, the "ipv6 unicast-routing" command must be enabled in global configuration mode. The following is a configuration example:

```
ipv6 unicast-routing
```

```
interface ethernet0/0
```

```
description connection to enterprise
```

```
ipv6 address 2002:C0A8:2101::1/128
```

---

**QUESTION 393:**

When an IPv6 enabled host boots, it sends a router solicitation (RS) message. An IPv6 router responds with a router advertisement (RA). Which two items are contained in the RA? (Select two)

- A. IPv6 address for the host
- B. lifetime of the prefix
- C. prefixes for the link
- D. keepalive timers
- E. request for the local host IP address

F. any route advertisement it has received

Answer: B, C

Explanation:

In IP Version 6, Router Advertisements have the following attributes:

Routers advertise periodically

- Max. time between advertisements can be in the range from 4 and 1800 seconds

- The advertisement has a lifetime (= 0 if not a default router)

Advertisement contains one or more prefixes

- Prefixes have a lifetime

Preferred lifetime

Valid lifetime

Specifies if state full of state less auto configuration is to be used

---

**QUESTION 394:**

Which three IPv6 notations represent the same address? Select three.

A. 2031:0000:130F:0000:0000:09C0:876A:130B

B. 2031::130F:9C0:876A:130B

C. 2031:0:130F::9C0:876A:130B

D. 2031::130F:0::9C0:876A:130B

E. 2031:0:130F:0:0:09C0:876A:130B

F. 2031:0:130F::9C0:876A:130B

Answer: A, C, E

Explanation:

With IP version 6, octets containing all zero's can be simply represented as :, while consecutive zero fields can be represented as ::. Answer choices C and E are simply the shorthand version of the fully written IPv6 address shown in choice

A. For a full

explanation of the IPv6 address representation, see the reference link below:

Reference: <http://www.ietf.org/rfc/rfc1924.txt?number=1924>

---

**QUESTION 395:**

Which three are characteristics of IPv6? Select three.

A. An IPv6 address is 128 bits long.

B. An IPv6 address is 20 bits long.

C. An IPv6 header contains the next header field.

D. An IPv6 header contains the protocol field.

- E. IPv6 routers send RA messages.
- F. An IPv6 header contains the header checksum field.

Answer: A, C, E

Explanation:

All IPv6 addresses are 128 bits long to accommodate a far larger number of stations than what was possible with the 32 bit IPv4 addresses.

The following displays the IPv6 header field in detail:

IPv6 header:

|   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 7 | 8 | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 3 |   |   |
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0 |

|                |               |            |
|----------------|---------------|------------|
| Version        | Traffic Class | Flow Label |
| Payload Length | Next Header   | Hop Limit  |

Source address :::

Destination address :::

Data :::

Version. 4 bits.

IPv6 version number.

Traffic Class. 8 bits.

Internet traffic priority delivery value.

Flow Label. 20 bits.

Used for specifying special router handling from source to destination(s) for a sequence of packets.

Payload Length. 16 bits unsigned.

Specifies the length of the data in the packet. When cleared to zero, the option is a hop-by-hop Jumbo payload.

Next Header. 8 bits.

Specifies the next encapsulated protocol. The values are compatible with those specified for the IPv4 protocol field.

Hop Limit. 8 bits unsigned.

For each router that forwards the packet, the hop limit is decremented by 1. When the hop limit field reaches zero, the packet is discarded. This replaces the TTL field in the IPv4 header that was originally intended to be used as a time based hop limit.

Source address. 16 bytes.

The IPv6 address of the sending node.

Destination address. 16 bytes.

The IPv6 address of the destination node.

Reference: <http://www.networksorcery.com/enp/protocol/ipv6.htm>

---

**QUESTION 396:**

Which three techniques can be used to transition from IPv4 to IPv6? (Select three)

- A. 6to4 tunneling
- B. flow label
- C. dual stack
- D. anycast
- E. NAT
- F. mobile IP

Answer: A, C, E

Explanation:

The four key strategies for deploying IPv6 are as follows:

1. Deploying IPv6 over IPv4 tunnels: These tunnels encapsulate the IPv6 traffic within the IPv4 packets, and are primarily for communication between isolated IPv6 sites or connection to remote IPv6 networks over an IPv4 backbone. The techniques include using manually configured tunnels, generic routing encapsulation (GRE) tunnels, semiautomatic tunnel mechanisms such as tunnel broker services, and fully automatic tunnel mechanisms such as IPv4-compatible and 6to4.
2. Deploying IPv6 over dedicated data links: This technique enables isolated IPv6 domains to communicate by using the same Layer 2 infrastructure as for IPv4, but with IPv6 using separate Frame Relay or ATM PVCs, separate optical links, or dense Wave Division Multiplexing (dWDM).
3. Deploying IPv6 over MPLS backbones: This technique allows isolated IPv6 domains to communicate with each other, but over an MPLS IPv4 backbone. Multiple techniques are available at different points in the network, but each requires little change to the backbone infrastructure or reconfiguration of the core routers because forwarding is based on labels rather than the IP header itself.
4. Deploying IPv6 using dual-stack backbones: This technique allows IPv4 and IPv6 applications to coexist in a dual IP layer routing backbone. All routers in the network need to be upgraded to be dual-stack with IPv4 communication using the IPv4 protocol stack and IPv6 communication using the IPv6 stack.

In addition to the strategies for deploying IPv6 within your IPv4 environment, you also need protocol translation mechanisms (for example, a NAT-PT device to connect IPv6-only web browsers to IPv4-only web servers) or dual-stack servers (for example, an e-mail server that handles IPv4-only and IPv6-only mail clients) to allow communication between applications using IPv4 and applications using IPv6. These mechanisms become increasingly important as IPv6 deployment moves from the testing to the actual usage phase, and more relevant as application developers decide that continuing to support IPv4 is not cost-effective.

Reference:

[http://www.cisco.com/en/US/tech/ CK8 72/technologies\\_white\\_paper09186a00800c9907.shtml#1098273](http://www.cisco.com/en/US/tech/CK8_72/technologies_white_paper09186a00800c9907.shtml#1098273)

---

**QUESTION 397:**

Route aggregation is a way of reducing the number of entries in a routing table. Three of the routing protocols below support automatic route summarization. Which ones are they? (Select three)

- A. IS-IS
- B. EIGRP
- C. OSPF
- D. IGRP
- E. RIP v.1

Answer: B, D, E

Explanation:

B: IGRP automatically summarize routes on the class network boundary.

D: EIGRP support both manual and automatic route summarization.

E: RIPv1 automatically summarize routes on the class network boundary.

Incorrect Answers:

A, C: IS-IS and OSPF only support manual route summarization.

---

**QUESTION 398:**

The Certkiller network consists of the following IP subnets:

| NetworkNumber   |
|-----------------|
| 192.168.31.0/24 |
| 192.168.32.0/24 |
| 192.168.33.0/24 |
| 192.168.34.0/24 |
| 192.168.35.0/24 |
| 192.168.36.0/24 |
| 192.168.37.0/24 |
| 192.168.38.0/24 |

|                 |
|-----------------|
| 192.168.39.0/24 |
| 192.168.40.0/24 |

How many CIDR blocks will be needed to summarize these subnets in the Certkiller network?

- A. Two
- B. Three
- C. Four
- D. Five
- E. None of the above

Answer: B

Explanation:

To properly summarize 10 contiguous networks, they will have to be broken into three separate advertisements: one will consist of 8 subnets (the middle eight) and the other two will consist of the first and last subnets. The summarizations are as follows:

192.168.31.0/24

192.168.32.0/21

192.168.40.0/24

Reference: Building Scalable Cisco Networks (Cisco Press) page 85

---

### QUESTION 399:

You are in the process of manually summarizing some routes within your network. Which of the addresses below would be included in the manual route summarization of 153.25.200.0/21?

- A. 153.25.198.0
- B. 153.25.206.0
- C. 153.25.208.0
- D. 153.25.224.0

Answer: B

Explanation:

The 21 leftmost bits must match. In particular we should study the 3rd octet.

Decimal 1st Octet 2nd Octet 3rd Octet 4th Octet

153.25.200.0 10011001 00011001 11001000 00000000

153.25.198.0 10011001 00011001 11000110 00000000

153.25.206.0 10011001 00011001 11001110 00000000

153.25.208.0 10011001 00011001 11010000 00000000  
153.25.224.0 10011001 00011001 11100000 00000000

---

**QUESTION 400:**

Which of the following statements about route summarization are true? (Select two)

- A. Private addresses cannot be summarized.
- B. Summarization is not compatible with VLSM.
- C. RIP v.1 automatically summarizes routes on network class boundaries.
- D. Classless routing protocols support summarization on any address bit boundary.

Answer: C, D

Explanation:

C: RIP V.1 is only able to summarize routes on network class boundaries.

D: Classless routing protocols support summarization on any address bit boundary.

Incorrect Answers:

A: Private address ranges can indeed be summarized.

B: Summarization is indeed compatible with VLSM.

---

**QUESTION 401:**

Some of the terms below are synonymous with route summarization? Which ones are they? (Select two)

- A. supernetting
- B. route aggregation
- C. address translation
- D. Classful inter-domain routing

Answer: A, B

Explanation:

Route summarization, also called route aggregation or supernetting, can reduce the number of routes that a router must maintain because it is a method of representing a series of network numbers in a single summary address.

---

**QUESTION 402:**

Your network is using the Class B private IP address scheme, as defined in RFC 1918. If you wanted to summarize this entire private address space into one single supernet, which prefix should you use?

- A. /8
- B. /12

- C. /16
- D. /24

Answer: B

Explanation:

In Class B, the private address is 172.16.0.0 to 172.31.255.255

Four bits are required to take care of the range 16-31. So, from the default subnet mask of 16, four more taken out and so, it becomes 12.

---

**QUESTION 403:**

In an EIGRP network using VLSM, where in the network does the process of route summarization occur?

- A. Manually on any router interface.
- B. Only at classless network boundaries.
- C. Only at classful network boundaries.
- D. Dynamically at the supernet boundary.

Answer: A

Explanation:

EIGRP automatically summarizes at the network boundary by default, but it can also be configured manually on any of the routed interfaces.

Incorrect Answers:

B, C: Since the word "only" is used, these two choices are incorrect. EIGRP can be configured to summarize automatically at the network boundary, or manually.

D: EIGRP provides no means for summarizing routes dynamically.

---

**QUESTION 404:**

Which one of the following addresses can be included in the route summarization of 172.17.200.0/21?

- A. 172.17.198.0
- B. 172.17.206.0
- C. 172.17.217.0
- D. 172.17.224.0

Answer: B

Explanation:

Here, we list the network addresses in binary form and see which subnet address match the 21 leftmost bits of the route summarization.

Decimal 1st Octet 2nd Octet 3rd Octet 4th Octet

172.17.200.0 10101100 00010001 11001000 00000000  
172.17.198.0 10101100 00010001 11000110 00000000  
172.17.206.0 10101100 00010001 11001110 00000000 Match!  
172.17.217.0 10101100 00010001 11011001 00000000  
172.17.224.0 10101100 00010001 11100000 00000000

---

**QUESTION 405:**

The Certkiller network contains the following 4 IP subnets:

172.18.129.0/24  
172.18.130.0/24  
172.18.132.0/24  
172.18.133.0/24

You want to summarize these networks into a single supernet. Which one of the following route summarizations is capable of summarizing every one of the networks in the above exhibit?

- A. 172.18.128.0/21
- B. 172.18.128.0/22
- C. 172.18.130.0/22
- D. 172.18.132.0/20

Answer: A

Explanation:

To illustrate, we list the network addresses in binary form and see how many leftmost bits match.

| Decimal      | 1st Octet | 2nd Octet | 3rd Octet | 4th Octet |
|--------------|-----------|-----------|-----------|-----------|
| 172.18.129.0 | 10101100  | 00010010  | 10000001  | 00000000  |
| 172.18.130.0 | 10101100  | 00010010  | 10000010  | 00000000  |
| 172.18.132.0 | 10101100  | 00010010  | 10000100  | 00000000  |
| 172.18.133.0 | 10101100  | 00010010  | 10000101  | 00000000  |
| 172.18.128.0 | 10101100  | 00010010  | 10000000  | 00000000  |

We see that the 21 leftmost bits match and that 172.18.128.0/21 can summarize the four networks.

Incorrect Answers:

- B, C: Only the 21 leftmost bits match, not 22.
  - D: This is an illegal summarization. It is host address, not a network address.
- 

**QUESTION 406:**

The Certkiller network is using the 172.19.100.0/24 IP address range for its entire network and wishes to summarize these routes. How many Class C IP subnets could be summarized in the 172.19.100.0/22 supernet?

- A. 2

- B. 4
- C. 8
- D. 16
- E. 255
- F. None of the above

Answer: B

Explanation:

As the summarized route is /22 the first 22 bits must be the same. A Class C subnet mask has 24 bits so the lower 2 bits in the 3rd octet can be changed. There are 4 possible combinations (00, 01, 10, 11) so 172.19.100.0/22 can summarize 4 class C addresses.

Reference: Source: Cisco Press, CCNP Self-Study Building Scalable Cisco Internetworks (BSCI). Chapter 2 Extending IP Addresses, Calculating VLSMs, page 73

---

### **QUESTION 407:**

**CORRECT TEXT**

You are administrating a network with subnets ranging from 172.16.94.0/24 - 172.16.118.0/24. A summary statement with the network 172.16.96.0 and the mask of 255.255.240.0 have summarized some of your subnets.

Which subnet is the last subnet to be included in the summary network?

Answer: 172.16.111.0

Explanation:

We can verify this answer as follows:

111 (3rd octet) decimal is 01101111. This is within the summarization. However, 172.16.112.0 would not be ok: 112 decimal is 01110000 (see below).

Decimal 1st Octet 2nd Octet 3rd Octet 4th Octet

255.255.240.0 11111111 11111111 11110000 00000000 subnet mask

172.16.96.0 10101100 10101100 01100000 00000000 summarized network

172.16.111.0 10101100 10101100 01101111 00000000 Included in summary element

172.16.112.0 10101100 10101100 01110000 00000000 Not included in summary element

Therefore, 172.16.111.0 is the last IP subnet that can be included in this summary route.

---

### **QUESTION 408:**

What strategy can a network administrator use to minimize the effect of routing table updates on internal routers when a single WAN interface frequently changes its state from up to down?

- A. Use a distance vector routing protocol.
- B. Use private IP addresses.

- C. Use dial-on-demand routing.
- D. Use route summarization.
- E. Use a routing protocol that tolerates route flapping.
- F. All of the above.

Answer: D


Explanation:

When using route summarization, multiple IP networks are advertised to downstream routers as a single route. As long as any of the IP subnets included within the summarized route are up and active, the summary route will be used, so any single link problem will not have to be propagated to the rest of the network.

---

**QUESTION 409:**

The Certkiller EIGRP network is shown below:



```
Certkiller1 # show ip route 2.0.0.0
Routing entry for 2.0.0.0/8, 4 known subnets
  Attached (2 connections)
  Variably subnetted with 2 masks
  Redistributing via eigrp 2000

C    2.1.3.0/24 is directly connected, Serial2
D    2.1.2.0/24 [90/10537472] via 10.1.1.2, 00:23:24, Serial1
D    2.1.0.0/8 is a summary, 00:23:20, Null0
C    2.1.1.0/24 is directly connected, Serial1
```

Router Certkiller 1 was configured in the following manner:

```
Certkiller 1(config)#router eigrp 2000
Certkiller 1(config-router)# network 2.0.0.0
Certkiller 1(config-router)# exit
Certkiller 1(config)#interface serial 2
Certkiller 1(config-if)#ip summary-address eigrp 2000 2.1.0.0 255.255.0.0
```

Based on the above information, how many routes will Certkiller 2 receive from router Certkiller 1?

- A. None
- B. One
- C. Two
- D. Three
- E. Four

Answer: B

Explanation:

Router Certkiller 1 was configured to summarize the network routes when advertising to the neighbor Certkiller 2. Because of this, router Certkiller 2 will only receive the 2.1.0.0/16 EIGRP route from router Certkiller 1.

**QUESTION 410:**

Which of the following statements are true regarding EIGRP manual summarization? (Select two)

- A. Manual summarization is configured on a per interface basis.
- B. Manual summaries can be configured with the classful mask only.
- C. When manual summarization is configured, auto-summarization is automatically disabled by default.
- D. The summary address is assigned an administrative distance of 10 by default.
- E. The summary address is entered into the routing table and is shown to be sourced from the Null0 interface.

Answer: A, E

Explanation:

By default, EIGRP routers summarize routes at the classful boundary. When routes are manually summarized, they are done so while in interface configuration mode. After the manual summarization is done, the logical null 0 interface is created, and the summary network is installed into the routing table via this interface.

---

**QUESTION 411:**

The CK1 and CK2 networks are consolidating into one large network. To accommodate this, the IP network was separated in the middle by a different major network boundary, creating a discontinuous network. With EIGRP being used as the routing protocol, what can be done to provide full connectivity across these non-contiguous networks? (Select two)

- A. Use private addresses on LAN links.
- B. Use static route redistribution.
- C. Use the no auto-summary command.
- D. Use route summarization.
- E. All of the above

Answer: B, C

Explanation:

Redistributing static route information could be used to advertise the routes across the network, since these static routes will not be automatically summarized.

By default, EIGRP summarizes IP subnets at the network boundary. However, in order to support the use of non-contiguous networks, the "no auto-summary" command can be used to disable the feature.

Incorrect Answers:

A: Using private of public addresses on any of the links will not correct this issue.

D: Route summarization could result in duplicate network advertisements. We want to prevent this by disabling route summarization.

---

**QUESTION 412:**

You are troubleshooting an OSPF problem and you need to view the OSPF neighbor information. Which of the following commands should you enter to examine neighbor adjacencies? (Select two)

- A. show ip ospf database
- B. show ip ospf neighbors
- C. show ip ospf protocols
- D. show ip ospf interfaces

Answer: B, D

Explanation:

B: Using the show ip ospf neighbor command, you can observe the neighbor data structure. This command displays OSPF-related neighbor information. The Interface field shows the interface on which the OSPF neighbor has formed adjacency.

Sample:

```
Router CK2 #show ip ospf neighbor
Neighbor ID Pri State Dead Time Address Interface
192.168.45.1 1 FULL/DR 00:00:36 10.0.0.1 Ethernet0
```

D: The show ip ospf interface command displays area ID and adjacency information

Incorrect Answers:

A: The "show ip ospf database" command displays the link-state database.

C: This is an invalid command.

Reference: <http://www.cisco.com/warp/public/104/16.html>

---

**QUESTION 413:**

You are a network troubleshooter and you've been contracted to examine an OSPF network. You feel that the CPU utilization of your routers is too high and want to figure out how many SPF calculations have occurred. Which one of the commands below could you use?

- A. show ip ospf
- B. show ip route
- C. show ip ospf interface
- D. show ip ospf protocols

Answer: A

Explanation:

The show ip ospf command displays summary information regarding the global OSPF

configuration. The output includes the number of times the Shortest Path First (SPF) algorithm has been run.

Sample output:

```
router CertK #show ip ospf
OSPF is running, process id: 1, router id: 10.1.2.136
Number of areas: 1, normal: 1, stub: 0
Area: 1.2.3.4
Number of interfaces in this area is 1
Type of authentication none
SPF algorithm has run 3 times
SPF interval 5 seconds
```

Incorrect Answers:

- B: The show ip route command displays IP routing table entries.
- C: The show ip interface command displays information about one or more interfaces.
- D: There is no such command.

---

**QUESTION 414:**

Which of the following IOS commands will allow you to view when a topological database purges out of date routes?

- A. show ip ospf
- B. show ip route
- C. show ip ospf interface
- D. show ip ospf protocols
- E. show ip ospf neighbor
- F. None of the above

Answer: C

Explanation:

The OSPF route dead timer is used to purge the outdated routes. The dead timer information can be obtained via the "show ip ospf interface" command.

Example:

```
Router CK1 # show ip ospf interface ethernet 0
Ethernet0 is up, line protocol is up
Internet Address 10.10.10.1/24, Area 0
Process ID 1, Router ID 192.168.45.1, Network Type BROADCAST, Cost: 10
Transmit Delay is 1 sec, State BDR, Priority 1
Designated Router (ID) 172.16.10.1, Interface address 10.10.10.2
Backup Designated router (ID) 192.168.45.1, Interface address
10.10.10.1
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:06
Index 1/1, flood queue length 0
Next 0x0(0)/0x0(0)
```

Last flood scan length is 2, maximum is 2

Last flood scan time is 0 msec, maximum is 4 msec

Neighbor Count is 1, Adjacent neighbor count is 1

Adjacent with neighbor 172.16.10.1 (Designated Router)

Suppress hello for 0 neighbor(s)

Timer Intervals

These are the values of the OSPF timers:

1. Hello-Interval time in seconds that a router sends an OSPF hello packet. On broadcast and point-to-point links, the default is 10 seconds. On NBMA, the default is 30 seconds.
2. Dead-Time in seconds to wait before declaring a neighbor dead. By default, the dead timer interval is four times the hello timer interval.
3. Wait-Timer interval that causes the interface to exit out of the wait period and select a DR on the network. This timer is always equal to the dead timer interval.
4. Retransmit-Time to wait before retransmitting a database description (DBD) packet when it has not been acknowledged.
5. Hello Due In-An OSPF hello packet is sent on this interface after this time. In this example, a hello is sent three seconds from the time the show ip ospf interface is issued.

Reference:

[http://www.cisco.com/en/US/tech/CK365/technologies\\_tech\\_note09186a0080094056.shtml#timer](http://www.cisco.com/en/US/tech/CK365/technologies_tech_note09186a0080094056.shtml#timer)

---

#### **QUESTION 415:**

You are troubleshooting a convergence problem on an OSPF network. Which of the following factors can affect the OSPF convergence time? (Select three)

- A. Hold-down timers.
- B. The size of the network.
- C. The number of routing areas.
- D. Maximum allowed hop count.
- E. The route calculation algorithm.

Answer: B, C, E

Explanation:

B: The size of the network affects the size of the topology table on the routers. A larger topology table will correspond to a higher convergence time.

C: Areas are introduced to put a boundary on the total number of link-state updates. The number of areas affects the convergence time.

E: The route calculation algorithm used for OSPF affects the convergence time.

Incorrect Answers:

A: OSPF does not use hold-down timers. Instead topology changes are flooded immediately. Distance vector routing protocols use hold-down timers.

D: OSPF does not use hop count. Distance vector routing protocols use hop counts.

---

#### **QUESTION 416:**

Which IOS command could you use to see which networks are being routed by an OSPF process?

- A. show ospf
- B. show ip route
- C. show ip protocols
- D. show ip ospf database
- E. show ip ospf protocols

Answer: C

Explanation:

The show ip protocols command display current routing protocols. It displays the parameters and current state of the active routing protocol process. The output includes a list of the networks routing for individual ospf processes.

Sample output:

```
Rt Certkiller # show ip protocols
Routing Protocol is "ospf 200"
Sending updates every 0 seconds
Invalid after 0 seconds, hold down 0, flushed after 0
Outgoing update filter list for all interfaces is not set
Incoming update filter list for all interfaces is not set
Redistributing: ospf 200
Routing for Networks:
172.6.31.5/32
Routing Information Sources:
Gateway Distance Last Update
Distance: (default is 110)
```

Incorrect Answers:


- A: The show ospf command displays summary information regarding the global OSPF configuration.
- B: The show ip route command displays the IP routing table.
- D: The show ip ospf database command displays the contents of the topological database maintained by the router. The command also shows the router ID and the OSPF process ID. However, the output does not include the networks routing for individual ospf processes.
- E: This is an invalid command.

---

#### **QUESTION 417:**

The Certkiller OSPF network is displayed in the following diagram.

642-801



In the Certkiller network, if the link to the designated router were to go down, what will Router CK1 and Router CK3 do?

- A. They send a flash update with poison reverse.
- B. They perform the designated router election process.
- C. They re-broadcast their routing tables to all other neighboring routers.
- D. They send a query to neighboring routers for other routers to the failed link.
- E. None of the above

Answer: E

Explanation:

A Designated Router (DR) router is elected by all other routers on the same LAN to represent all the routers. Each network has one DR. Normally, in this case the BDR would take over as the DR. In this case, the router that was already acting as the BDR (either CK1 or CK3 ) would take over as the DR. A new BDR election would then take place.

Incorrect Answers

- A: IGRP and EIGRP use poison reverse. Poison reverse is not used by OSPF.
- B: The router that was already acting as the BDR would take over as the DR.
- C, D: This is not a necessary step in the OSPF process.

---

**QUESTION 418:**

When issuing the "show ip ospf neighbor" command on router CK1 , a neighbor is reported as a DROTHER. Which of the following is true regarding this OSPF neighbor?

- A. The neighbor is down and currently not forwarding traffic.
- B. It is not in full state and is still exchange LSAs.
- C. It is the backup DR:
- D. It is neither the DR nor BDR.
- E. It is acting as both the DR and the BDR

Answer: D

Explanation:

When viewing the OSPF neighbor relationships, a neighbor can be one of three different

values:

DR - This neighbor is the DR

BDR - This neighbor is the backup DR

DROTHER - This neighbor is neither the DR nor the BDR.

---

**QUESTION 419:**

You're troubleshooting a router that's running Integrated IS-IS and you suspect something abnormal with the neighboring routers. Which IOS command would you enter to display the adjacencies?

- A. show isis route
- B. show clns route
- C. show isis database
- D. show clns neighbors
- E. None of the above

Answer: D

Explanation:

The show clns neighbors command displays ES and IS neighbors. The output includes adjacency information.

Example:

```
router CertK # show clns neighbors
System Id SNPA Interface State Holdtime Type Protocol
0000.0000.0007 aa00.0400.6408 Ethernet0 Init 277 IS ES-IS
0000.0C00.0C35 0000.0c00.0c36 Ethernet1 Up 91 L1 IS-IS
0800.2B16.24EA aa00.0400.2d05 Ethernet0 Up 29 L1L2 IS-IS
0800.2B14.060E aa00.0400.9205 Ethernet0 Up 1698 ES ES-IS
0000.0C00.3E51 *HDLC* Serial1 Up 28 L2 IS-IS
0000.0C00.62E6 0000.0c00.62e7 Ethernet1 Up 22 L1 IS-IS
0A00.0400.2D05 aa00.0400.2d05 Ethernet0 Init 24 IS ES-IS
```

Incorrect Answers

A: The show isis routes command is used to display the IS-IS Level 1 forwarding table for IS-IS learned routes.

B: The show clns route command is used to display all of the destinations to which this router knows how to route packets.

C: The show isis database command is used to display the Intermediate System-to-Intermediate System (IS-IS) link state database.

---

**QUESTION 420:**

The IS-IS routing table for Certkiller 1 is displayed in the diagram below:

**Certkiller1#show ip route isis**

```
10..0.0.0/8 is variably subnetted, 7 subnets, 3 masks
i L2 10.200.200.14/32 [115/30] via 10.1.0.2, Serial1/0
i L1 10.200.200.13/32 [115/20] via 10.1.1.3, Ethernet0/0
i L1 10.1.3.0/24 [115/20] via 10.1.1.3, Ethernet0/0
i L2 10.1.2.0/23 [115/20] via 10.1.0.2, Serial1/0
i su 10.1.0.0/23 [115/10] via 0.0.0.0, Null0
```

Regarding the command output on Certkiller 1 in the exhibit, which statement is true?

- A. The 10.1.0.0/23 route is an IS-IS external route
- B. The Certkiller 1 IS-IS router is an ASBR.
- C. The 10.1.0.0/23 route is a suppressed route.
- D. The 10.1.0.0/23 route is a summary route.
- E. The Certkiller 1 IS-IS router is an ASBR that belongs to multiple IS-IS areas.
- F. The Certkiller 1 IS-IS router is performing route aggregation and is suppressing the more specific 10.1.0.0/23 prefix.

Answer: D

Explanation:

The "su" in the routing table denotes an IS-IS summary null route. This route is automatically created within the router when IS-IS route summarization is configured.

---

**QUESTION 421:**

The "show isis route" command was issued on router Certkiller 3 as shown below:

Certkiller3#show isis rout

```
IS-IS Level-1 Routing Table      version
System Id      Next-Hop      Interface  SNPA          Metric  State
Certkiller3    Certkiller1   Et0/0      aabb.cc00.3300  10     Up    L2-IS
Certkiller1    --
```

Default route out of area - (via 1 L2-attached IS)

```
System Id      Next-Hop      Interface  SNPA          Metric  State
Certkiller1    Certkiller1   Et0/0      aabb.cc00.3300  10     Up
```

Based on the information above, which statement is true?

- A. Certkiller 1 is the exit point out of the area for Certkiller 3
- B. Certkiller 1 is a level-1 only IS-IS router.
- C. Certkiller 1 has been configured with a non-default IS-IS metric.
- D. Certkiller 3 is a level-2 only IS-IS router.
- E. Certkiller 3 has been configured with a non-default IS-IS metric.
- F. Certkiller 3 routing table should contain L2 entries.
- G. None of the above

Answer: A

Explanation:

In the IS-IS routing table of router Certkiller 3, the default router is displayed as Certkiller 1, which is a Level 2 router. Level 2 routers are used for routing inter-area traffic, so Certkiller 3 will use this router as the exit point for all traffic destined to another area.

Incorrect Answers:

B: Certkiller 1 is a Level 2 router.

C, E: The default IS-IS metric for all links is 10, which is the value shown here.

D, F: Certkiller 3 is a level 1 router.

---

### QUESTION 422:

The routing table of Certkiller 1 is displayed below:

```
Certkiller1#show ip route
```

```
Gateway of fast resort is not set
```

```
140.140.0.0 is subnetted (mask is 255.255.255.0), 3 subnets
```

```
C 140.140.64.0 255.255.255.0 is possibly down, routing via 0.0.0.0, Ethernet 0
```

```
i L2 140.140.67.0 [115/20] via 140.140.64.240, 0:00:12, Ethernet 0
```

```
i L2 140.140.66.0 [115/20] via 140.140.64.240, 0:00:12, Ethernet 0
```

Based on this information, how was network 140.140.66.0 discovered?

- A. It is an IGRP route that has been redistributed in from an external protocol.
- B. It is an IS-IS route via inter-area routing.
- C. It is an IS-IS route via intra-area routing.
- D. It is an IBGP route.
- E. It is a static route.

Answer: B

Explanation:

Level-2 routers are used for routing traffic between areas (inter-area). Based on the routing table of Certkiller 1, the 140.140.66.0/24 network is known via a Level 2 IS-IS router.

---

### QUESTION 423:

You have a multi-area OSPF network and you're concerned because one of the sites is having connectivity problem to resources in a different area. Which IOS privileged mode command would you enter to confirm that your network: A) has a path to its ABR, B) has a path to its ASBR, and C) the SPF calculation is functional?

- A. show ip protocols
- B. show running-config

- C. show ip ospf neighbor
- D. show ip ospf border-routers

Answer: D

Explanation:

The show ip ospf border-routers command displays the internal OSPF routing table entries to an area border router (ABR) and autonomous system boundary router (ASBR). The SPF No in the output is the internal number of SPF calculation that installs this route.

Example:

```
Router Certkiller # show ip ospf border-routers
OSPF Process 109 internal Routing Table
Destination Next Hop Cost Type Rte Type Area SPF No
160.89.97.53 144.144.1.53 10 ABR INTRA 0.0.0.3 3
160.89.103.51 160.89.96.51 10 ABR INTRA 0.0.0.3 3
160.89.103.52 160.89.96.51 20 ASBR INTER 0.0.0.3 3
160.89.103.52 144.144.1.53 22 ASBR INTER 0.0.0.3 3
```

Incorrect Answers:

- A: The show ip protocols command only displays routing protocol parameters and current timer values.
- B: The show running-config command displays the currently used configuration mode. The required information will not be displayed.
- C: The show ip ospf neighbor command displays OSPF-neighbor information on a per-interface basis. It does not include ABR, ASBR or SPF information.

---

**QUESTION 424:**

You have logged onto the console of a router running OSPF, and for your analysis log you need information on: the filters, default metric, maximum paths, and the number of areas configured. Which of the IOS commands below should you use?

- A. show ip protocol
- B. show ip route
- C. show ip ospf interface
- D. show ip ospf
- E. Show ip routing

Answer: A

Explanation:

The show ip protocols command, displays parameters about timers, filters, metrics, network, and other information for the entire router.

Reference: Building Scalable Cisco Networks (Cisco Press) page 133

---

**QUESTION 425:**

Which IOS command tells you how many times the OSPF shortest path first algorithm was executed?

- A. show ip protocol
- B. show ip ospf interface
- C. show ip ospf
- D. show ip ospf database
- E. None of the above

Answer: C

Explanation:

The following table describes the output of the "show ip ospf" command and their meanings:

| Table52 show ip ospf Field Descriptions |                                      |
|---|--------------------------------------|
| Field                                   | Description                          |
| Routing process "ospf 201" with ID      | Process ID and OSPF router ID.       |
| 10.0.0.1                                |                                      |
| Supports...                             | Number of types of service supported |
|   | (Type 0 only).                       |
| SPF schedule delay                      | Delay time of SPF calculations.      |
| Minimum LSA interval                    | Minimum interval between link-state  |
|   | advertisements.                      |
| LSA group pacing timer                  | Configured LSA group pacing timer    |
|   | (in seconds).                        |
| Interface flood pacing timer            | Configured LSA flood pacing timer    |
|   | (in milliseconds).                   |
| Retransmission pacing timer             | Configured LSA retransmission        |
|   | pacing                               |

**642-801**

|  |   |
|--|---|
|  | timer (in milliseconds).  |
| Number of...                                   | Number and type of link-state advertisements that have been received.   |
| Number of external LSA                         | Number of external link-state advertisements.                           |
| Number of opaque AS LSA                        | Number of opaque link-state advertisements.                             |
| Number of DCbitless external and opaque AS LSA | Number of demand circuit external and opaque link-state advertisements. |
| Number of DoNotAge external and opaque AS LSA  | Number of do not age external and opaque link-state advertisements.     |
| Number of areas in this router is              | Number of areas configured for the router.                              |
| External flood list length                     | External flood list length.   |

Reference:

[http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/iprrp\\_r/ip2\\_s3g.htm#wp1036469](http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/iprrp_r/ip2_s3g.htm#wp1036469)

---

**QUESTION 426:**

An OSPF link can be in multiple states at any given moment (ie. Exstart, exchange, full). Which two IOS commands let you view the state of the link? (Select two)

- A. show ip ospf
- B. show ip protocols
- C. show ip ospf neighbor
- D. show ip ospf interface

Answer: C, D

Explanation:

The link state exstart is an OSPF link state (see note below). We need retrieve OSPF link state information.

C: The output of the show ip ospf neighbor command is used To display OSPF-neighbor information on a per-interface basis. It includes link state information.

D: The show ip ospf interface command is used to display OSPF-related interface information for a particular interface. This includes the link state of the specified interface.

Note: exstart state: After two OSPF neighboring routers establish bi-directional communication and complete DR/BDR election (on multi-access networks), the routers transition to the exstart state.

Incorrect Answers:

A: The show ip ospf command is used to display general information about OSPF routing processes. However, it does not include any link state information.

B: The command "show ip protocols" displays the parameters and current state of the active routing protocol process. It does not show any link state information.

---

**QUESTION 427:**

While performing a routine inspection of your OSPF network you come to the realization that the Designated Router (DR) may be overloaded, and you're considering changing the DR. What are three ways you could manipulate the election of the DR? (Select three)

- A. Use of the priority command.
- B. Use of the router-id command.
- C. Assignment of the loopback address.
- D. Assigning a lower IP address
- E. By adding additional memory to the desired router.

Answer: A, B, C

Explanation:

A: The OSPF router priority is only used on multi-access networks such as LAN's. This establishes whether the router is eligible to become the Designated Router (DR) for the LAN. A priority of zero means that the router is not eligible to become DR.

B: The router ID is used to break ties during the DR and BDR election processes if the priority values are equal.

C: The loopback address is used in DR election. The highest IP address on the router is the router ID. If a loopback address is configured, then it is the router ID.


Incorrect Answers:

D: The highest IP address is used as the router ID, not the lowest.

E: This would not make it a more likely DR candidate, did the DR election process does not take into consideration the amount of memory.

**QUESTION 428:**

You are the network engineer at Certkiller . The Certkiller network topology is shown in the following graphic:



OSPF area 1 has been configured to receive only the intra-area routes and a default route from the area 1 ABR. During routine maintenance, you issue the "show ip route" command on area 1 internal routers and notice that O and O IA routes appear, but no O E2 routes appear.

What is the probable cause of this problem?

- A. At the ABR for area 1, the area 1 stub command is missing the no-summary option.
- B. All the routers in area 1 are missing the area 1 stub no summary option.
- C. At the internal routers in area 1, the area 1 stub command is missing.
- D. At the ABR for area 1, the area 1 stub command is missing.
- E. The virtual link that transits area 1 is configured incorrectly.
- F. Area 1 should be configured as a not-so-stubby-area with the area 1 nssa command.


Answer: A

Explanation:

An autonomous system boundary router (ASBR) advertises external destinations throughout the OSPF autonomous system. In many cases, external link states make up a large percentage of the link states in the databases of every router. A stub area is an area in which you don't allow advertisements of external routes, thus reducing the size of the database even more. Instead, a default summary route (0.0.0.0) is inserted into the stub area in order to reach these external routes. If you have no external routes in your network, then you have no need to define stub areas.

**QUESTION 429:**

The complex Certkiller OSPF network topology is displayed in the following diagram:



The routing table of Certkiller 3 contains all the proper OSPF routes except the 10.4.4.0/24 OSPF route. Which show command on Certkiller 3 will be most useful to determine the cause of this problem and why?

- A. Perform `show is ospf virtual-link` to examine if the virtual link between Certkiller 3 and Certkiller 4 is configured and in the up state.
- B. Perform `show run` to examine if area 3 is configured as a stubby area.
- C. Perform `show run` to examine if area 3 is configured as a nssa area.
- D. Perform `show ip ospf neighbor` to examine which router is elected as the DR on the e2 interface.
- E. Perform `show ip ospf neighbor` to examine the OSPF database entries.

Answer: A

Explanation:

All OSPF areas need to be directly connected to the backbone (area 0) in order to ensure proper OSPF operation. If an area is not able to be directly connected with area 0, a virtual link needs to be configured in order to logically connect that area to the backbone network. In this case, a virtual link needs to be configured between routers Certkiller 3 and Certkiller 4. This will logically extend the OSPF area 3 into area 4, connecting it into the backbone. In this example, if the route to 10.4.4.0/24 is missing, a misconfigured or broken virtual link is the most likely cause of the problem.

**QUESTION 430:**

Which command would display OSPF parameters such as filters, default metric, maximum paths, and number of areas configured on a router?

- A. `show ip protocol`
- B. `show ip route`
- C. `show ip ospf interface`

- D. show ip ospf
- E. show ip interface
- F. None of the above

Answer: A

Explanation:

The "show ip protocol" command displays values about routing timers and network information associated with the entire router

. This includes, the AS number associated with the routing process, number of areas configured on the router, the metric, and the maximum paths.

---

**QUESTION 431:**

What can cause an OSPF virtual link not to be established?

- A. There are DR/BDR election errors on the OSPF virtual link.
- B. The ABRs are using the wrong OSPF router ID in the virtual link configuration.
- C. The virtual link transit area is not Area 0.
- D. The virtual link is set up over an NBMA network.
- E. There is a mismatch of the OSPF network type on the virtual link configuration.

Answer: B

Explanation:

Every OSPF area must be contiguous to the Backbone Area 0. Non-backbone areas can be thought of "leaves" to the OSPF tree where the backbone is the trunk.

If it is not physically possible to have an area connected to the backbone, you must configure a Virtual Link in the two routers which are providing the tunnel to the backbone. These two routers do not need to be physically connected, but they must share a common area called the "transit area".

Each virtual link section is identified by the Router ID (NOT the interface IP address) of the virtual link neighbor (NOT the Router ID of THIS router). The Router ID of the neighbor is the largest IP interface address associated with that router, and the IP address used for the virtual link must specify the router ID of the peer. You can request the router ID of the virtual neighbor by issuing the command "show ospf rtrid".

Reference:

[http://www.cisco.com/en/US/products/hw/routers/ps4081/products\\_configuration\\_example09186a0080094260.s](http://www.cisco.com/en/US/products/hw/routers/ps4081/products_configuration_example09186a0080094260.s)

---

**QUESTION 432:**

While troubleshooting a BGP problem on the Certkiller network you notice that the 10.10.10.0/24 prefix is not being injected into the local BGP table of CK1 . The relevant configuration of router CK1 is shown below:

```
router bgp 65001
```

```
network 10.0.0.0
neighbor 172.16.1.1 remote-as 65002
no auto-summary
Routing table information:
show ip route | include 10
O 10.10.10.0/24 [110/11] via 192.168.1.1, 2d00h,
Ethernet0/0
Why doesn't the local BGP table have the prefix?
```

- A. This route is not a BGP learned route.
- B. The network command is wrong.
- C. The 172.16.1.1 neighbor is down.
- D. The prefix 10.10.10.0/24 is not a connected route.

Answer: B

Explanation:

The correct syntax should be "network 10.10.10.0 mask 255.255.255.0" under the BGP routing process. Without the correct subnet mask specified, the route will not get injected into the BGP routing table, even if it is learned via an IGP. In this case, the route is known via OSPF.

---

**QUESTION 433:**

What would happen if a full mesh of BGP sessions were configured within a single large autonomous system? (Select two)

- A. Many UDP sessions will be created.
- B. More memory and CPU will be consumed.
- C. This configuration is not permitted by default.
- D. A significant amount of bandwidth for BGP updates and retransmissions can be used.
- E. Permanent Virtual Circuits (PVCs) must be created to link the fully meshed BGP sessions.

Answer: B, D

Explanation:

The number of TCP connections that are required for a full mesh grows exponentially. The formula for the total number of BGP sessions is  $N(N-1)/2$ . For example, for 10 IBGP routers to become fully meshed 45 peering sessions will need to be configured.

B: More memory and CPU resources are required on the routers to support the full mesh BGP.

D: There will be many TCP sessions and a significant amount of bandwidth will be required for BGP traffic.

Incorrect Answers:

A: Many TCP, not UDP, sessions will be created.

- C: It is permitted.
- E: PVCs are not required.

---

**QUESTION 434:**

An ISP is running a large IBGP network with 25 routers. The full mesh topology that is currently in place is inefficient on bandwidth and BGP traffic. What can the administrator configure to reduce the number of BGP neighbor relationships within the AS?

- A. route reflector
- B. route maps
- C. route redistribution
- D. peer groups
- E. aggregate addresses

Answer: A

Explanation:

Route reflectors can be used to overcome this scalability issues with configuring a full IBGP mesh within a network. With route reflectors, the clients need to only peer with one single router (the route reflector). In this way, fewer peering sessions need to be configured and maintained within the IBGP network.

---

**QUESTION 435:**

To establish a full mesh IBGP between ten routers, how many IBGP sessions are required?

- A. 10
- B. 45
- C. 50
- D. 90
- E. 100

Answer: B

Explanation:

For full mesh  $X*(X-1)/2$  connections are needed. So,  $10(10-1)/2 = (10*9)/2=45$ .

---

**QUESTION 436:**

While verifying the OSPF operation on your network, you issue the following command:

```
show ip ospf interface
```

What information parameters will you get after you execute this command? (Select

two)

- A. The router ID
- B. The summary link counts
- C. The neighbor adjacencies
- D. The link-state update interval

Answer: A, C

Explanation:

A: The router ID is displayed.

C: Neighbor adjacencies are included in the output of this command.

D: OSPF use the hello interval as update interval. The hello interval is displayed by this command.

---

**QUESTION 437:**

You issue the "show ip route" command on router CK1 and receive the following output:

```
S 62.99.153.0/24 [1/0] via 209.177.64.130
172.209.12.0/32 is subnetted, 1 subnets
D EX 172.209.12.1
[170/2590720] via 209.179.2.114, 06:47:28,
Serial0/0/0.1239
62.113.17.0/24 is variably subnetted, 2 subnets, 2 masks
D 62.113.17.0/29 [90/30208] via 62.113.20.10, 07:35:24,
ATMO/1/0.130
S 62.113.17.0/24 [1/0] via 62.113.1.25
D EX 99.3.215.0/24
[170/27316] via 209.180.96.45, 09:52:10,
FastEthernet11/0/0
[170/27316] via 209.180.96.44, 09:52:10,
FastEthernet11/0/0
D 25.248.17.0/24
[90/1512111] via 209.179.66.25, 10:33:13,
Serial0/0/0.1400001
[90/1512111] via 209.179.66.41, 10:33:13,
Serial0/0/0.1402001
62.113.1.0/24 is variably subnetted, 12 subnets, 2 masks
D 62.113.1.227/32
[90/24823552] via 209.180.96.45, 07:35:24,
FastEthernet1/0/0
[90/24823552] via 209.180.96.44, 07:35:24,
FastEthernet1/0/0
S* 0.0.0.0/0 [1/0] via 209.180.96.14
```

One of the routing entries in router CK1 was learned via a static route. Which one is

it?

- A. 99.3.215.0
- B. 62.99.153.0
- C. 172.209.12.1
- D. 62.113.1.227

Answer: B

Explanation:

The S denotes a static route (see below).

S 62.99.153.0/24 [1/0] via 209.177.64.130

Incorrect Answers:

A, C: The D EX denotes a route learned through EIGRP external.

D: The D denotes a route learned through EIGRP.

---

**QUESTION 438:**

While verifying BGP operation on the Certkiller router, you issue the "show ip bgp" command as shown below:

```
routerCK>show ip bgp
```

```
BGP table version is 1046033, local router ID is 198.32.162.100
```

```
Status codes: s suppressed, d damped, h history, * valid, > best,
```

```
i - internal
```

```
Origin codes: i - IGP, e EGP, ? -incomplete
```

```
Network Next Hop Metric LocPrf Weight Path
```

```
* > 143.16.0.0 128.214.63.2 0 400 0 200 1
```

```
* 143.16.0.0 192.208.10.5 0 300 0 300 1
```

```
* 143.16.0.0 143.16.63.5 0 100 0 200 1
```

```
* 143.16.0.0 203.250.13.41 0 100 0 500 1
```

From the information above, which path will the network 143.16.0.0 prefer to take to exit the AS?

- A. 128.214.63.2
- B. 192.208.10.5
- C. 128.213.63.5
- D. 203.250.13.41
- E. All of the above will be used in a round robin fashion.

Answer: A

Explanation:

Local preference (LocPref) is a well-known discretionary attribute that provides an indication to routers in the AS about which path is preferred to exit the AS. A path with a higher local preference is more preferred. In this scenario the following entry has the highest local preference value of 400.

Network Next Hop Metric LocPrf Weight Path

```
* > 128.213.0.0 128.214.63.2 0 400 0 200 1
```

The preferred exit path of the AS is therefore 128.214.63.2, as noted by the ">" which refers to the best path for this destination.

---

**QUESTION 439:**

Which IOS command would you use to find out which networks are routed by a particular OSPF process?

- A. show ospf
- B. show ip route
- C. show ip protocols
- D. show ip ospf database
- E. None of the above

Answer: C

Explanation:

The show ip protocols command display current routing protocols. It displays the parameters and current state of the active routing protocol process. The output includes a list of the networks routing for individual ospf processes.

Sample output:

```
Rt Certkiller # show ip protocols
Routing Protocol is "ospf 200"
Sending updates every 0 seconds
Invalid after 0 seconds, hold down 0, flushed after 0
Outgoing update filter list for all interfaces is not set
Incoming update filter list for all interfaces is not set
Redistributing: ospf 200
Routing for Networks:
172.6.31.5/32
Routing Information Sources:
Gateway Distance Last Update
Distance: (default is 110)
```

Incorrect Answers:

- A: The show ospf command displays summary information regarding the global OSPF configuration.
  - B: The show ip route command displays the IP routing table.
  - D: The show ip ospf database command displays the contents of the topological database maintained by the router. The command also shows the router ID and the OSPF process ID. However, the output does not include the networks routing for individual ospf processes.
- 

**QUESTION 440:**

**642-801**

The routing table of router CK is shown below:

```
CK#show ip routeCodes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B -  
BGP D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area E1 - OSPF  
external type 1, E2 - OSPF external type 2, E - EGP i - IS-IS, L1 - IS-IS level-1, L2 -  
IS-IS level-2, * - candidate default<Some output deleted>Gateway of last resort is  
30.64.0.2 to network 0.0.0.0 30.0.0.0/8 is variably subnetted, 9 subnets, 2 masksO IA  
30.2.0.0/16 [110/74] via 30.64.0.2, 00:09:13, Ethernet0C 30.1.3.0/24 is directly  
connected, Serial0O IA 30.3.0.0/16 [110/148] via 30.64.0.2, 00:05:22, Ethernet0C  
30.1.2.0/24 is directly connected, Serial1
```

Based on the output above, which of the following statements are true regarding the routing table? (Select two)

- A. The area is a stub area.
- B. The area is totally stubby.
- C. Network 30 is using VLSM
- D. The routing table is for an ABR

Answer: A, C

Explanation:

Incorrect Answers:

A: This is a stub area.

C: The exhibits states that the 30.0.0.0/8 subnet is variably subnetted and the networks used are /16 and /24, so VLSM is used

Incorrect Answers:

B: The area is stub, not totally stubby.

D: This is not an ABR (Area border router) since there are no external routes, only internal and directly connected.

---

**QUESTION 441:**

While troubleshooting an OSPF routing issue you issue the "show ip ospf interface" command. Which of the following parameters are provided from this? (Select two)

- A. Router ID
- B. Summary link counts
- C. Neighbor adjacencies
- D. Link-state update interval

Answer: A, C

Explanation:

The show ip ospf interface command verifies that interfaces have been configured in the intended areas. If no loopback address is specified, the interface with the highest address is the taken router ID. It also gives the timer intervals, including the hello interval, and

shows the neighbor adjacencies.

Reference: Building Scalable Cisco Networks (Cisco Press) page 134.

---

**QUESTION 442:**

Which command should you use to obtain information on OSPF link state advertisements? In particular: which LSA's have been sent, which LSA's have been received, and the time when the last LSA was received?

- A. show ip ospf database
- B. show ip ospf neighbors
- C. show ip ospf protocols
- D. show ip ospf interfaces
- E. None of the above

Answer: A

Explanation:

The show ip ospf database command is used to display lists of information related to the OSPF topological, the link state, database for a specific router.

Sample output:

```
R_ Certkiller #show ip ospf database
OSPF Router with ID (192.168.0.12) (Process ID 1)
Router Link States (Area 0)
Link ID ADV Router Age Seq# Checksum Link count
192.168.0.10 192.1680.10 817 0x80000003 0xFF56 1
192.168.0.11 192.1680.11 817 0x80000003 0xFD55 1
192.168.0.12 192.168.0.12 816 0x80000003 0xFB54 1
192.168.0.13 192.168.0.13 816 0x80000003 0xF953 1
192.168.0.14 192.168.0.14 817 0x80000003 0xD990 1
Net Link States (Area 0)
Link ID ADV Router Age Seq# Checksum
192.168.0.14 192.168.0.14 812 0x80000002 0x4AC8
```

Incorrect Answers:

B: The show ip ospf neighbor is used to display OSPF-neighbor information on a per-interface basis. It does show the required information though.

Sample output:

```
Neighbor ID Pri State Dead Time Address Interface
192.168.0.13 1 2WAY/DROTHER 00:00:31 192.168.0.13 Ethernet0
192.168.0.14 1 FULL/BDR 00:00:38 192.168.0.14 Ethernet0
```

C: There is no such command.

D: The show ip ospf interface command is used to display OSPF-related interface information. It displays the circuit name and state, IP address, network mask, broadcast address, redundancy, Internet Control Message Protocol (ICMP) settings, and RIP settings. However, it does not display LSAs.

Sample output:

```
RouterCK# show ip interfaces
IP Interface Summary:
Circuit Name: VLAN2 State: active
IP Address: 172.16.1.200 Network Mask: 255.255.0.0
Broadcast Address: 172.16.255.255 Redundancy: disabled
ICMP Redirect: enabled ICMP Unreachable: enabled
RIP: enabled
```

---

**QUESTION 443:**

You have a console connection on a router running RIP. Which IOS command would you enter if you wanted to view the RIP routing transactions that are occurring?

- A. show ip rip database
- B. show ip route
- C. show ip protocols rip
- D. debug ip rip
- E. debug ip routing
- F. None of the above

Answer: D

Explanation:

Use the debug ip rip EXEC command to display information on RIP routing transactions in real time. The no form of this command disables debugging output.

Reference:

[http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products\\_command\\_reference\\_chapter09186a008007](http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products_command_reference_chapter09186a008007)

---

**QUESTION 444:**

While troubleshooting a problem on the Certkiller network, you issue the "show ip protocol" as shown below:

```
CK1 #show ip protocol
Outgoing update filter list for all interface is
Incoming update filter list for all interface is
Default networks flagged in outgoing updates
Default networks accepted from incoming updates
EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
EIGRP maximum hopcount 100
EIGRP maximum metric variance 1
Redistributing: eigrp 101, ospf 101 (internal, external 1, external 2)
Automatic network summarization is in effect
Routing for Networks:
192.168.1.0
```

Routing Information Sources:

Gateway Distance Last Update

192.168.1.34 90 00:05:21

192.168.1.40 90 00:05:21

192.168.1.18 90 00:05:21

Distance: internal 90 external 170

Routing Protocol is "ospf 101"

Sending updates every 0 seconds

Invalid after 0 seconds, hold down 0, flushed after 0

Outgoing update filter list for all interfaces is

Incoming update filter list for all interfaces is

Redistributing: eigrp 101, ospf 101

Routing for Networks:

10.1.1.1/32

Routing Information Sources:

Gateway Distance Last Update

172.16.11.100 110 00:00:16

Based on this information, which of the following statements are true for Router CK1 ? (Choose three)

- A. Redistribution has been configured for EIGRP into OSPF.
- B. Redistribution has been configured for OSPF into EIGRP.
- C. EIGRP auto-summary has been disabled.
- D. The OSPF process-id and the EIGRP autonomous system number can not be identical.
- E. EIGRP has been configured for un-equal cost paths load balancing.
- F. Router CK1 has received routing updates from three EIGRP neighbors and from one OSPF neighbor.

Answer: A, B, F

Explanation:

A: Under the EIGRP configuration the line "Redistributing: eigrp 101, ospf 101" is present.

B: Under the OSPF configuration the line "Redistributing: eigrp 101, ospf 101" is present.

F: This is correct because of the three neighbor updates listed under the Gateway Distance Last Update section of the EIGRP configuration

Incorrect Answers:

C: The output explicitly states that automatic network summarization is in effect.

D: This is not true. In this router, they are both configured for AS 101


E: is incorrect because under the EIGRP configuration the EIGRP maximum metric variance is 1, which is the default, specifying equal cost load balancing.

---

### **QUESTION 445:**

The Certkiller network is displayed below, along with the relevant router configurations for CK1 and CK2 .

## 642-801



PC1 is unable to ping PC2 successfully.

Given the above diagram and configurations, which solution would fix this problem?

- A. Enable IP classless on CK1 and CK2 .
- B. To support discontinuous subnet, enable RIPv2 on CK1 and CK2 then disable auto-summary.
- C. Enable both RIPv1 and RIPv2 on CK1 and CK2 .
- D. Enable no auto-summary in RIP router configuration mode on CK1 and CK2 .

Answer: B

Explanation:

With RIP version 1 configured, both routers will summarize the LAN subnets and each will advertise the 10.0.0.0/8 route to each other. If we enable use a routing protocol that supports VLSM, such as RIPv2, and then disable the auto-summarization feature, then router CK1 will advertise the 10.1.1.0/24 subnet to CK2 , and CK2 will advertise the 10.3.3.0/24 subnet to CK1 .

Incorrect Answers:

- A: This will not solve the problem, since each router has a 10s network on each of their LANs.
- C: This alone will not solve the problem, since we also need to disable the automatic summarization of network routes.
- D: This is not an option for RIP version 1 networks, since it is a classful routing protocol.

---

### **QUESTION 446:**

The IP BGP table of router P1R3 is displayed in the following diagram:

```
P1R3#show ip bgp
BGP table version is 16, local router ID is 10.200.200.13
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
  *>i192.168.11.0    10.200.200.12      0   100   101 64998 65222 65223 i
  * i                10.200.200.11      0   200     0 64998 65222 65223 i
  *>i192.168.12.0    10.200.200.12      0   100   101 64998 65222 65223 i
  * i                10.200.200.11      0   200     0 64998 65222 65223 i
```

Based on the P1R3 output shown above, which statement is true?

- A. The best path to reach the 192.168.11.0 prefix is via 10.200.200.11.
- B. The best path to reach the 192.168.11.0 prefix is via 10.200.200.12.
- C. The best path to reach the 192.168.11.0 prefix is via both 10.200.200.11 and 10.200.200.12; BGP will automatically load balance between the two.
- D. The 192.168.11.0 and 192.168.12.0 prefixes were learned via EBGP from the 10.200.200.11 and 10.200.200.12 EBGP neighbors.
- E. None of the above

Answer: B

Explanation:

The best path to any given destination is noted by the ">" in the IP BGP table. In this case, the best path to 192.168.11.0 is via next hop 10.200.200.12 due to the fact that the weight is higher (101) than the path via the alternative next hop. Weight is a Cisco proprietary method for path determination and the weight value is used above all other values. Within a router, the path with the highest weight will be preferred.

---

**QUESTION 447:**

Which IOS command could you use to verify if a BGP router is or isn't a route reflector?

- A. show bgp neighbor
- B. show running-config
- C. show route-reflector
- D. show route-reflector-client
- E. None of the above

Answer: A

Explanation:

The show bgp neighbor command indicates if a particular neighbor is a route reflector client. Note: Route reflectors modify the BGP split horizon rule by allowing the router configured as the route reflector to propagate routes learned by IBGP to other IBGP peers.

Incorrect Answers:

B: Displaying the current configuration would not be useful in finding information in the route reflector status.

C, D: There are no such commands.

---

**QUESTION 448:**

Which IOS command would you enter if you wanted to view a list of IBGP and EBGP neighbor relationships that are configured?

- A. show ip bgp
- B. show ip bgp paths
- C. show ip bgp peers
- D. show ip bgp summary
- E. show ip bgp protocols

Answer: D

Explanation:

The show ip bgp summary command displays the status of all BGP connections. Neighbors with corresponding AS values will be listed; both interior and external.

Incorrect Answers:

A: The show ip bgp command displays routes in the BGP routing table, not the neighbors.

B: The show ip bgp paths command is used to display all the BGP paths in the database. However, it does not list the neighbors.

C, E: There is no such command.

Reference:

[http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1\\_r/1rprt1/1rbgp.htm](http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_r/1rprt1/1rbgp.htm)

---

**QUESTION 449:**

The 192.168.0.0/16 network is not being propagated throughout the network via BGP as expected. Observe the BGP configuration commands from the advertising router shown below.

```
Router bgp 65111
neighbor 172.16.1.1 remote-as 65111
neighbor 172.16.2.1 remote-as 65112
network 192.168.0.0
network 10.0.0.0
!
```

```
ip route 192.168.0.0 255.255.0.0 null0
```

What is the reason the 192.168.0.0/16 route is not being advertised?

- A. The network 192.168.0.0 statement is missing mask 255.255.0.0
- B. The network 192.168.0.0 statement is missing mask 0.0.255.255
- C. The network 10.0.0.0 statement is missing mask 255.0.0.0

- D. The network 10.0.0.0 statement is missing mask 0.255.255.255
- E. The auto-summary configuration is missing

Answer: A

Explanation:

The network 192.168.0.0 statement is missing mask 255.255.0.0. Without the mask command used in a network statement, the route may not get properly injected into the BGP routing process.

---

#### **QUESTION 450:**

The Certkiller network is redistributing routing information from OSPF and EIGRP. Which three steps are most helpful in verifying proper route redistribution? (Select three)

- A. On the routers not performing the route redistribution, use the show ip route command to see if the redistributed routes show up.
- B. On the ASBR router performing the route redistribution, use the show ip protocol command to verify the redistribution configurations
- C. On the ASBR router performing the route redistribution, use the show ip route command to verify that the proper routes from each routing protocol are there.
- D. On the routers not performing the route redistribution, use the show ip protocols command to verify the routing information sources.
- E. On the routers not performing the route redistribution, use the debug ip routing command to verify the routing updates from the ASBR.

Answer: A, B, C

Explanation:

In order to verify proper route redistribution, use the "show ip route" command on all routers within the network, as well as the ASBR, to verify that the routes are properly being advertised to all routers. In addition, issuing the "show ip protocol" can be used on the router performing the redistribution to verify that routes are being redistributed into each other.

Incorrect Answers:

- D: Issuing this command on a non-redistribution router will not tell us where and how the routes are originating from. This command will only be useful on the redistributing routers.
  - E: This command can not be used to verify the redistributed routes.
- 

#### **QUESTION 451:**

Exhibit:

Certkiller A# show ip route isis  
10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks

```
IL2 10.200.200.14/32 [115/30] via 10.1.0.2, Serial 1/0
IL1 10.200.200.13/32 [115/20] via 10.1.1.3, Ethernet 0/0
IL1 10.1.3.0/24 [115/20] via 10.1.1.3, Ethernet 0/0
IL2 10.1.2.0/23 [115/20] via 10.1.0.2, Serial 1/0
I su 10.1.0.0/23 [115/10] via 0.0.0.0, Null0
```

Based on the show ip route isis output on Router Certkiller A, which statement is true?

- A. The "I su" 10.1.0.0/23 route is an IS-IS external route.
- B. The Certkiller A IS-IS router is an ASBR.
- C. The "I su" 10.1.0.0/23 route is a suppressed route.
- D. The "I su" 10.1.0.0/23 route is a summary route.
- E. The R1 IS-IS router is an ABR that belongs to multiple IS-IS areas.
- F. The R1 IS-IS route is performing route aggregation and is suppressing the more 10.1.0.0/23 prefix.

Answer: D

Explanation:

When viewing the "show IP route ISIS" output, the "su" entry represents a summarized route. This is the direct result of the "summary-address" router configuration command. When creating the summary route. The Cisco IOS automatically creates the summary route and point it to interface Null 0.

---

### **QUESTION 452:**

Exhibit:

```
Certkiller A# show ip eigrp topology
IP-EIGRP Topology Table for process 200
Codes:P - Passive, A - Active, U - Update, Q - Query, R -
Reply, r - Reply status
P 192.168.1.64/28, 1 successors, FD is 281600
Via Connected, Ethernet0
P 192.168.1.32/28, 1 successors, FD is 40512000
Via Connected, Serial1
P 192.168.1.48/28, 1 successors, FD is 40537600
Via 192.168.1.66 (40537600/40512000), Ethernet0
Via 192.168.1.77 (41024000/40512000), Serial0
Via 192.168.1.33 (41024000/40512000), Serial1
P 192.168.1.16/28, 1 successors, FD is 40512000
Via Connected, Serial0
```

Based on the above show ip eigrp topology output, which three statements are true?  
(Choose three.)

- A. Certkiller A is in AS 200
- B. Certkiller A will balance between three paths to reach the 192.168.1.48/28 prefix, because all three paths have the same AD of 40512000.

- C. The best path for Certkiller A to reach the 192.168.1.48/28 prefix is via 192.168.1.66.
- D. 40512000 is the AD to reach the 192.168.1.48/28 prefix.
- E. All of the routes are in the passive mode because these routes are in the hold-down state.
- F. All the routes are in the passive mode, because Certkiller A is in the query process for those routes.

Answer: A, C, D

Explanation:

It can be determined that AS 200 is used, from the fact that the IS-IS process ID is labeled as 200. The best path to reach the network 192.168.1.48/28 is the first one displayed in the routing table. This can be further demonstrated by the fact that the metric is less than the alternative route, via serial 0. Finally, the AD can be found by viewing the second number within the parentheses, which in this case is 40512000.

---

#### **QUESTION 453:**

The following output was seen on router Certkiller 1 and Certkiller 3:  
The Certkiller 1 and Certkiller 3 routers are OSPF neighbors over the Ethernet 0/0 connection. Based on the show ip ospf neighbor output from the Certkiller 1 and Certkiller 3 routers, which statement is true?

- A. Certkiller 1 is the DR because it has a higher OSPF router priority.
- B. Certkiller 1 is the DR because it has a lower OSPF router ID.
- C. Certkiller 3 is the DR because it has a higher OSPF router priority.
- D. Certkiller 3 is the DR because it has a lower OSPF router ID.
- E. Both Certkiller 1 and Certkiller 3 are using the default OSPF router priority.

Answer: A


Explanation:

The output shown above displays information about the neighbors, so from Certkiller 3 we see that the priority of Certkiller 1 is 2, and that Certkiller 1 is the DR. From Certkiller 1 we can see that its neighbor ( Certkiller 3) is the BDR and has an OSPF priority of 1.

---

#### **QUESTION 454:**

The Certkiller WAN is depicted below:



Certkiller 3 is redistributing the EIGRP routes into OSPF. What will the EIGRP routes appear in the routing table of Certkiller 1?

- A. O
- B. O IA
- C. O E2
- D. D
- E. D EX

Answer: C

Explanation:

O E1 or O E2. The routes in this LSA are external to the autonomous system. They can be configured to have one of two values. E1 will include the internal cost to the ASBR added to the external cost reported by the ASBR. E2 does not compute the internal cost - it just reports the external cost to the remote destination.

---

**QUESTION 455:**

Which three types of OSPF route entries can be found in the routing table of an internal OSPF router within an OSPF not so stubby area? (Select three)

- A. O
- B. O\*IA
- C. O\*OA
- D. O E1
- E. O E2
- F. O N1

Answer: A, B, F

Explanation:

The various route types used by OSPF are as follows:

O - OSPF

IA - OSPF inter area

N1 - OSPF NSSA external type 1

N2 - OSPF NSSA external type 2


E1 - OSPF external type 1

E2 - OSPF external type 2

An OSPF NSSA will receive inter-area, external type 2, OSPF routes.

**QUESTION 456:**

Four Certkiller routers are connected together as shown in the diagram below:



Router Certkiller 2 is configured as shown below:

```
hostname Certkiller2
!
interface serial0/0/24
 ip policy route-map TEST
!
access-list 1 permit 192.168.10.0 0.0.0.255
access-list 101 permit tcp host 192.168.10.5 eq telnet any
!
route-map TEST permit 5
 match ip address 1
 set ip next-hop 192.168.16.254
!
route-map TEST permit 15
 match ip address 101
 set ip next-hop 192.168.17.254
```

You work as a network technician at Certkiller .com. Study the two exhibits. Policy-based routing is configured on Certkiller 2 to direct all traffic coming from sources with address prefixes 192.168.10.0/24 to the next hop 192.168.16.254. The exception should be made for packets originating from the Telnet port on host 192.168.10.5 to be forwarded to the next hop 192.168.17.254. The debug ip packet command executed on Certkiller 3 shows that the policy routing does not work correctly.

What could be the problem?

- A. The telnet keyword in the access-list 101 is associated with the wrong port.
- B. The standard access list 1 should be replaced with an extended access-list.
- C. The ip local route-map should be configured on Certkiller 1 Serial0 interface.
- D. The route map statements are in the wrong order.

Answer: D

Explanation:

Policy routing map statements work similarly to access lists in that the different rules are consulted in order until a match is found. Once a match is found, the traffic is policy

routed immediately and the following policy rule statement is not consulted. In order to meet the requirements of this example, the first and second statements should be switched.

---

**QUESTION 457:**

Part of the configuration file and debug output for one of the Certkiller routers is displayed in the following exhibit:

```
interface Ethernet0/0
 ip address 100.100.100.1 255.255.255.0
 ip policy route-map blah
!
interface Serial1/0
 ip address 10.10.10.1 255.255.255.0
!
interface Serial2/0
 ip address 20.20.20.1 255.255.255.0
!
ip classless
no ip http server
!
access-list 100 permit ip host 100.100.100.3 host 200.200.200.4
!
route-map blah permit 10
 match ip address 100
 set ip default next-hop 10.10.10.2

routing debugging is on

Policy routing debugging is on
*Dec 4 12:50:57.360: IP: s=100.100.100.3 (Ethernet0/0), d=200.200.200.4, len 100, policy match
*Dec 4 12:50:57.360: IP: route map blah, item 10, permit
*Dec 4 12:50:57.360: IP: s=100.100.100.3 (Ethernet0/0), d=200.200.200.4 (Serial2/0), len 100,
policy rejected -- normal forwarding
*Dec 4 12:50:57.431: IP: s=100.100.100.3 (Ethernet0/0), d=200.200.200.4, len 100, policy match
*Dec 4 12:50:57.431: IP: route map blah, item 10, permit
*Dec 4 12:50:57.431: IP: s=100.100.100.3 (Ethernet0/0), d=200.200.200.4 (Serial2/0), len 100,
policy rejected -- normal forwarding
*Dec 4 12:50:57.491: IP: s=100.100.100.3 (Ethernet0/0), d=200.200.200.4, len 100, policy match
*Dec 4 12:50:57.491: IP: route map blah, item 10, permit
*Dec 4 12:50:57.491: IP: s=100.100.100.3 (Ethernet0/0), d=200.200.200.4 (Serial2/0), len 100,
policy rejected -- normal forwarding
```

Study the exhibit carefully.

Which statement is true regarding policy-based routing?

- A. IP packets matching the access list are being forwarded to 10.10.10.2 because of the set ip default next-hop command.
- B. IP packets not matching the access list are being forwarded to 10.10.10.2 because of the set ip default next-hop command.
- C. IP packets matching the access list are not being forwarded to 10.10.10.2 because a path to 200.200.200.4 exists in the routing table.
- D. IP packets matching the access list are not being forwarded to 10.10.10.2 because a path to 200.200.200.4 does not exist in the routing table.
- E. IP packets matching the access list are being forwarded to 10.10.10.2 because OSPF has not been configured on the Serial1/0 interface.

Answer: D


Explanation:

In the example above, the packets sourced from 100.100.100.3 destined to 200.200.200.4 are correctly matching the policy map as specified in access-list 100, but a route to 200.200.200.4 does not exist in the routing table, so the policy is rejected and normal forwarding will occur.

---

**QUESTION 458:**

Three Certkiller routers are connected together as shown in the diagram below:



Exhibit

Certkiller1#show ip route

...

- C 10.1.3.0/24 is directly connected, Serial2
- O 10.1.2.0/24 [90/10537472] via 10.1.1.2, 00:23:24, Serial1
- C 10.1.1.0/24 is directly connected, Serial 1
- S 192.1668.20.0/24 is directly connected, ethernet0

You work as a network technician at Certkiller .com. Study the exhibits carefully. If the following command, "default information-originate", is added to the OSPF configuration on Certkiller 1, what will the neighboring routers see in their route table?

- A. S\* 0.0.0.0/0 [1/0] via 6.0.0.3, 00:28:00
- B. O E1 0.0.0.0 [1/0] via 6.0.0.3, 00:28:00
- C. O E2 0.0.0.0 [1/0] via 6.0.0.3, 00:28:00
- D. No default routes will be propagated to neighboring routers.

Answer: D

Explanation:

Default routes injected into a normal area can be originated by any OSPF router. The OSPF router does not, by default, generate a default route into the OSPF domain. In order for OSPF to generate a default route, you must use the default-information originate command. However, this will only advertise a default route to the OSPF neighbors if the router itself also contains a default route. Based on the routing table shown in this example, a default route does not already exist.

There are two ways to advertise a default route into a normal area. The first is to advertise 0.0.0.0 into the OSPF domain, provided the advertising router already has a default route. The second is to advertise 0.0.0.0 regardless of whether the advertising router already has a default route. The second method can be accomplished by adding the keyword always to the default-information originate command. In this example, the "always" keyword would be required in order to propagate this default route.

---

**QUESTION 459:**

EIGRP uses five generic packet types (hello, updates, queries, replies, acknowledgements). If you wished to view the statistics for these packets, which IOS command should you use?

- A. debug eigrp packets
- B. show ip eigrp traffic
- C. show ip eigrp topology
- D. show ip eigrp neighbors

Answer: B

Explanation:

The show ip eigrp traffic command displays the number of Enhanced IGRP (EIGRP) packets sent and received.

Example:

The following is sample output from the show ip eigrp traffic command:

```
Router# show ip eigrp traffic
IP-EIGRP Traffic Statistics for process 77
Hellos sent/received: 218/205
Updates sent/received: 7/23
Queries sent/received: 2/0
Replies sent/received: 0/2
Acks sent/received: 21/14
```

Reference

:[http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products\\_command\\_reference\\_chapter09186a00800](http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products_command_reference_chapter09186a00800)

---

**QUESTION 460:**

While troubleshooting a routing problem on the Certkiller EIGRP network you discover that one of the routers is failing to establish adjacencies with its neighbor.

What is a likely cause of this problem between neighbors? (Select two)

- A. The K-values do not match.
- B. The hold times do not match.
- C. The hello times do not match.
- D. The AS numbers do not match.

Answer: A, D

Explanation:

Peer relationships and adjacencies between routers will not be formed between EIGRP routers if the neighbor resides in a different autonomous system or if the metric-calculation mechanism (K values) is misaligned for that link.

Incorrect Answers:

B, C: It is possible for two routers to become EIGRP neighbors even though the hello and hold timers do not match.

---

**QUESTION 461:**

While troubleshooting an EIGRP routing issue, you are seeing a high number of SIA (stuck in active) routes. Which of the following are causes of a route becoming SIA? (Select two)

- A. Some query or reply packets are lost between the routers.
- B. The neighboring router stops receiving ACK packets from this router.
- C. The neighboring router starts receiving route updates from this router.
- D. A failure causes traffic on a link between two neighboring routers to flow in only one direction (unidirectional link).

Answer: A, D

Explanation:

The acknowledgement does not reach the destination or they are too delayed. This is normally due to too many routing topology changes, or a router with insufficient memory.

Note: In some circumstances, it takes a very long time for a query to be answered. So long, in fact, that the router that issued the query gives up and clears its connection to the router that isn't answering, effectively restarting the neighbor session. This is known as a stuck in active (SIA) route. The most basic SIA routes occur when it simply takes too long for a query to reach the other end of the network and for a reply to travel back.

Incorrect Answers:

B: Ack packets don't reply to Query, only Reply do.


C: Does not apply to SI

A. This is the normal operation of EIGRP.

Reference:<http://www.cisco.com/warp/public/103/eigrp3.html>

**QUESTION 462:**

The Certkiller network is displayed in the diagram below:



If the FDDI interface in the EIGRP network above were to fail or shut down, which router(s) will become Router A's next-hop to Network 7?

- A. B
- B. D only
- C. H only
- D. D and H

Answer: C

Explanation:

Router H will be the successor, and that route will be placed in the Routing table. Router A detects the link failure between Router B and network 7. It checks the topology table for a successor. It finds that H is the successor since the advertised distance for H (30) is less than the feasible distance for B (31). However, there is no next best route - no feasible successor. The candidate route through D has an advertised distance (220) that is higher than the feasible distance of the successor route (40).

**QUESTION 463:**

The Certkiller network consists of a hub and spoke topology with a main router supporting about 20 regional offices. A point-to-point Frame Relay WAN connects the regional offices to the main office, and EIGRP is deployed as the routing protocol. The committed information rate (CIR) for each of the Frame Relay PVC's is different, and the bandwidth command IS NOT configured on any of the interfaces or subinterfaces. You want to ensure that EIGRP routes everything properly. How should you configure the network?

- A. Convert each Frame Relay PVC to point-to-multipoint connection
- B. Manually configure the bandwidth of the major interface to the lowest CIR x 24

- C. Manually configure the bandwidth of the major interface to the highest CIR x 24
- D. Manually configure the bandwidth of each of these PVCs to equal to their respective CIR.

Answer: D

Explanation:

Although this will require some administrative effort, the only way to ensure that EIGRP properly considers the actual bandwidth to use in the routing decision, each link should be set to the CIR.


Incorrect Answers:

- A: This is not a Cisco recommended solution.
- B: This will force all PVC to run at a low speed.
- C: This would give too high a bandwidth.

---

**QUESTION 464:**

The Certkiller EIGRP network is displayed in the diagram below:



All routers are running EIGRP. Based on this, what will router CK3 do if the link between CK3 and CK1 were to go down? (Select two)

- A. It elects a new designated router.
- B. It sends a flash update with poison reverse.
- C. It checks its topology table for an alternate route.
- D. It re-broadcasts its routing table to all other neighboring routers.
- E. It sends a query to neighboring routers for other routes to the failed link.

Answer: C, E

Explanation:

The steps of convergence in an EIGRP network are as follows:

1. Router CK3 detects the link failure between CK1 and CK3 . It checks the topology table for a feasible successor, but it doesn't find a qualifying alternate route and enters in an active convergence state. (C)
2. CK3 sends a Query out all interfaces for other routes to the failed link (E). The neighboring routers acknowledge the query.
3. The reply from CK4 indicates no other route to the failed link.
4. CK2 's reply contains a route to the failed link, although it has a higher feasible distance.

5. Router CK3 accepts the new path and metric information, places it in the topology table, and creates an entry for the routing table.
6. CK3 sends an update about the new route out all interfaces.

---

**QUESTION 465:**

The EIGRP topology table for router Certkiller 1 is displayed below:

```
Certkiller1#show ip eigrp topology
IP-EIGRP Topology Table for process 200

Codes: P-Passive, A-Active, U-Update, Q-Query, R-Reply,
       r - Reply status

P 192.168.1.64/28, 1 successors, FD is 281600
   via Connected, Ethernet0
P 192.168.1.32/28, 1 successors, FD is 40512000
   via Connected, Serial1
P 192.168.1.48/28, 1 successors, FD is 40537600
   via 192.168.1.66 (40537600/40512000), Ethernet0
   via 192.168.1.17 (41024000/40512000), Serial0
   via 192.168.1.33 (41024000/40512000), Serial1
P 192.168.1.16/28, 1 successors, FD is 40512000
   via Connected, Serial0
```

Regarding the command output on Certkiller 1 in the exhibit, which statements are true? (Select three)

- A. Certkiller 1 is in AS 200
- B. Certkiller 1 will load balance between three paths to reach the 192.168.1.48/28 prefix, because all three paths have the same AD of 40512000.
- C. The best path for Certkiller 1 to reach the 192.168.1.48/28 prefix is via 192.168.1.66.
- D. 40512000 is the advertised metric via 192.168.1.66 to reach the 192.168.1.48/28 prefix.
- E. All the routes are in the passive mode because these routes are in the hold-down state.
- F. All the routes are in the passive mode, because Certkiller 1 is in the query process for those routes.

Answer: A, C, D

Explanation:

The Certkiller 1 router resides in AS 200, as displayed by the "IP EIGRP topology for process 200" output.

Regarding the numbers specified in the parenthesis, the first number is the EIGRP metric that represents the cost to the destination. The second number is the EIGRP metric that this peer advertised.

Based on this, the best path to the 192.168.1.48/28 destination is via 192.168.1.66, because the metric is less than the alternatives.

---

**QUESTION 466:**

The following show command was issued on router CK1 :

```
CK1#show ip eigrp topology
IP-EIGRP Topology Table for AS(24)/ID(1.1.1.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 10.2.1.0/24, 1 successors, FD is 28160
   via Connected, FastEthernet0/0
P 10.1.1.0/24, 1 successors, FD is 20512000
   via 10.1.1.2 (25632000/5120256), Serial0/0
P 17.19.0.0/16, 1 successors, FD is Inaccessible, Q
   via 10.1.1.2 (40992000/40480000), Serial0/0
P 172.16.0.0/16, 1 successors, FD is 23072000
   via 10.1.1.2 (23072000/2560256), Serial0/0
P 172.17.0.0/16, 1 successors, FD is 25632000
   via 10.1.1.2 (25632000/5120256), Serial0/0
P 172.18.0.0/16, 1 successors, FD is 32032000
   via 10.1.1.2 (32032000/11520256), Serial0/0
A 172.19.0.0/16, 1 successors, FD is Inaccessible, Q
  2 replies, active never, query-origin: Successor Origin
  via 10.1.1.3 (391248640/390736640), r, Serial0/0
  Remaining replies:
  via 10.1.1.2, r, Serial0/0
CK1#
```

Based on the exhibited command output, which two statements are true? (Choose two)

- A. The EIGRP network is stable.
- B. The router at 10.1.1.3 has not replied to the CK1 query packet.
- C. The route to 172.19.0.0/16 is undergoing recomputation.
- D. The route to 172.19.0.0/16 is stuck-in-active.
- E. CK1 has sent a query packet to 10.1.1.2.

Answer: B, C


Explanation:

The output shown above is a snapshot from an EIGRP router re-computing the route to 172.19.0.0/16. As can be seen above, a query was originated by CK1 to the neighbor at 10.1.1.3, but the query has not yet been replied to by this router.

---

**QUESTION 467:**

Router CK1 is configured for EIGRP and OSPF as shown below:



During the redistribution process configured on CK1, some of the EIGRP routes, such as 10.1.1.0/24 and 10.2.2.0/24, are not being redistributed into the OSPF routing domain.

Which two items could be a solution to this problem? (Choose two)

- A. Change the metric-type to 2 in the redistribute command.
- B. Configure the redistribute command under router eigrp 1 instead.
- C. Change the EIGRP AS number from 100 to 1 in the redistribute command.
- D. Add the subnets option to the redistribute command.
- E. Change the metric to an EIGRP compatible metric value (bandwidth, delay, reliability, load MTUs) in the redistribute command.

Answer: A. B

Explanation:

A: Change the metric-type to 2 in the redistribute command.

because metric type 2 is the default metric for OSPF and metric type 1 as given in the command is used to redistribute BGP. Which should be the second answer no clue, that's why maybe you could help.

B: In Given Configuration, Only Network 172.16.0.0 0.0.255.255 of OSPF is going to redistribute into the EIGRP routing table. That is called one way redistribute. For Complete Communication, EIGRP also should redistribute the EIGRP routing table for OSPF.

Example:

```


# router eigrp 1
# network 10.1.1.0 0.0.0.255
# redistribute ospf 100 metric 1000 100 255 1 1500

```

---

### QUESTION 468:

Exhibit:



Based on the information in the exhibit, which statement is true?

- A. RTC will be able to access the 10.0.0.0 network.
- B. RTC will not have the 10.0.0.0 network in its routing table.
- C. RTC will not have the 192.168.10.0 network in its routing table.
- D. RTB will not have the 10.0.0.0 network in its routing table.
- E. RTB and RTC will not have the 10.0.0.0 network in their routing tables.

Answer: B

Explanation:

Use the distribute-list command to pick and choose which routing updates a router will send or receive. By referencing an access list, the distribute-list creates a route filter. This is a set of rules that precisely controls what routes a router will send or receive in a routing update. This command is available for all IP routing protocols and can be applied to either inbound or outbound routing updates. When applied to inbound updates, the syntax for configuring a route filter is as follows:  
 Router(config-router)#distribute-list access-list-number in  
 [interface-name]

---

### QUESTION 469:

Which two provide inter-area routing services? (Choose two.)

- A. L1 IS
- B. L1 ES
- C. L2 IS
- D. L2 ES
- E. L1/L2 IS

Answer: C, E

Explanation:

1. Level 1 (L1) ISs, closely equivalent to OSPF internal non-backbone routers, are


responsible for routing to ESs inside an area. L1 ISs enable communication between ESs in an area.

2. Level 2 (L2) ISs, closely equivalent to backbone routers in OSPF, route between areas only.

3. Level 1 and Level 2 (L1L2) Intermediate ISs, closely equivalent to area border routers (ABRs) in OSPF, route between areas and the backbone. They participate in the L1 intra-area routing and the L2 inter-area routing.

### QUESTION 470:

Exhibit:



Refer to the exhibit. Assume that the Ethernet LANs on Certkiller A, Certkiller B, and Certkiller C have the appropriate network commands configured under the referenced IGP. Passive interfaces have been configured as indicated. Which three statements are true? (Choose three.)

- A. If RIP v1 is configured on all routers, the Certkiller C will receive updates about the Certkiller A LAN.
- B. If RIP v2 is configured on all routers, then the Certkiller B LAN will only be advertised to Certkiller D and Certkiller E.
- C. If IGRP is configured on all routers, the Certkiller A will receive updates about the LANs on Certkiller B, Certkiller D, and Certkiller E.
- D. If EIGRP is configured on all routers, the Certkiller B will establish neighbor relationship with Certkiller A and Certkiller C.
- E. If EIGRP is configured on all routers, then Certkiller B will not be able to establish neighbor relationships with Certkiller C, Certkiller D, and Certkiller E.
- F. If OSPF is configured on all routers, the Certkiller B will establish neighbor relationships with Certkiller D and Certkiller E.

Answer: A, C, E

Explanation:

A passive interface receives updates, but does not send them. The passive-interface

command can be used with all IP interior gateway protocols. It can be used with RIP, IGRP, EIGRP, OSPF, and IS-IS. The syntax of the command is as follows:

Router(config-router)#passive-interface type number

Passive-interface default : puts all interfaces in passive mode.

Ethernet interface of Router CK is configured passive interface, S0/0,

S0/1 interface is non-passive interface. So, Routers will get the routing updates from s0/0 and s0/1 interface of CKB.

---

**QUESTION 471:**

Exhibit:

```
R1#show ip route
Gateway of last resort is 10.1.1.2 to network 0.0.0.0

C    1.0.0.0/8 is directly connected, Loopback0
    172.17.0.0/24 is subnetted, 1 subnets
D    172.17.1.0 [90/25632000] via 10.1.1.2, 00:05:20, Serial0/0
    172.16.0.0/24 is subnetted, 1 subnets
D    172.16.1.0 [90/20640000] via 10.1.1.3, 00:05:21, Serial0/0
     [90/20640000] via 10.1.1.3, 00:05:21, Serial0/0
D    172.19.0.0/16 [90/20640000] via 10.1.1.3, 00:05:21, Serial0/0
D    172.22.0.0/16 [90/20640000] via 10.1.1.3, 00:05:21, Serial0/0
D EX 172.25.0.0/16 [170/32032000] via 10.1.1.2, 00:00:10, Serial0/0
    10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D    10.2.0.0/16 is a summary, 00:06:18, Null0
C    10.2.1.0/24 is directly connected, FastEthernet0/0
C    10.1.1.0/24 is directly connected, Serial0/0
D*EX 0.0.0.0/0 [170/20514560] via 10.1.1.2, 00:00:11, Serial0/0
R1#
```

Based on the exhibited output, which three statements are true? (Choose three.)

- A. R1 is configured with the ip summary-address command.
- B. The route 10.2.0.0/16 was redistributed into EIGRP.
- C. R1 is configured with the variance command.
- D. A default route has been redistributed into the EIGRP autonomous system.
- E. The router at 10.1.1.2 is configured with the ip default-network 0.0.0.0 command.
- F. R1 is sourcing an external EIGRP route from Null0.

Answer: A, C, D

Explanation:

EIGRP also enables manual configuration of a prefix to use as a summary address. Manual summary routes are configured on a per-interface basis. The interface that will propagate the route summary must first be selected and then defined with the ip summary-address eigrp command, which has the following syntax:

Router(config-if)#ip summary-address eigrp autonomous-system-number

ip-address mask administrative-distance

D 2.1.0.0/16 is a summary, 00:00:22, Null0

Notice that the summary route is sourced from Null0, and not an actual interface

Although the redistribution command is available for all IP routing protocols, it

behaves differently depending on the actual IP routing protocols involved. However, the underlying principles are the same. Therefore, the examples in this section can be used as a starting point for any redistribution scheme.


Example of Redistribute

```
# router eigrp 1
# network 10.1.1.0 0.0.0.255
# redistribute ospf 100 metric 1000 100 255 1 1500
```

---

**QUESTION 472:**

The Certkiller OSPF network is displayed below:



| Neighbor ID | Pri | State          | Dead Time | Address     | Interface |
|-------------|-----|----------------|-----------|-------------|-----------|
| 192.168.1.2 | 1   | output omitted | 00:00:34  | 192.168.1.2 | Ethernet0 |
| 192.168.1.4 | 1   | output omitted | 00:00:32  | 192.168.1.4 | Ethernet0 |
| 192.168.1.3 | 1   | output omitted | 00:00:39  | 192.168.1.3 | Ethernet0 |

Refer to the exhibit. All routers have simultaneously been reloaded and the DR election has concluded as expected. What state is Certkiller 2 in?

- A. 2WAY/DROTHER
- B. 2WAY/BDR
- C. 2WAY/DR
- D. FULL/DROTHER
- E. FULL/BDR
- F. FULL/DR

Answer: A

Explanation

1. Designated router (DR)

- For every multiaccess IP network, one router will be elected the DR. The DR has two main functions. The first function is to become adjacent to all other routers on the network, and second, is to act as a spokesperson for the network. As spokesperson, the DR will send network LSAs for all other IP networks to every other router. Because the DR becomes adjacent to all other routers on the IP network, it is the focal point for collecting routing information (LSAs).

2. Backup designated router (BDR) - The DR could represent a single point

of failure, so a second router is elected as the BDR to provide fault tolerance. Therefore, the BDR must also become adjacent to all routers on the network and must serve as a second focal point for LSAs. However, unlike the DR, the BDR is not responsible for updating the other routers or sending network LSAs. Instead, the BDR keeps a timer on the update activity of the DR to ensure that it is operational. If the BDR does not detect activity from the DR before the timer expires, the BDR takes over the role of DR and a new BDR is elected.

3. Those router, which unable to become the DR or BDR will become DROTHER.

---

**QUESTION 473:**

Which statement is true concerning 6to4 tunneling?

- A. IPv4 traffic is encapsulated with an IPv6 header.
- B. The edge routers can use any locally configured IPv6 address.
- C. Hosts and routers inside a 6to4 site will need a special code.
- D. An edge router must use IPv6 address of 2002::/16 in its prefix.

Answer: D


Explanation:

6to4 Tunneling allows IPv6 to be run over an automatically configured tunnel. It requires that the routers connecting the IPv6 remote sites through the IPv4 cloud need to be running dual stacks. The edge routers responsible for running the tunnel will use the prefix 2002::/16 and append the IPv4 interface address to create an address. The interface is the IPv4 address converted to hexadecimal and added to the routing prefix 2002::/16.

---

**QUESTION 474:**

Exhibit:



Refer to the exhibit. If Certkiller 1 is configured for 6to4 tunneling, what will the prefix of its IPv6 network be?

- A. 1723:1100:1::/48
- B. FFFF:AC1F:6401::/16
- C. AC1F:6401::/32
- D. 2002:AC1F:6401::/48

E. 3FFE:AC1F:6401::/32

Answer: D

Explanation:

6to4 Tunneling allows IPv6 to be run over an automatically configured tunnel. It requires that the routers connecting the IPv6 remote sites through the IPv4 cloud need to be running dual stacks. The edge routers responsible for running the tunnel will use the prefix 2002::/16 and append the IPv4 interface address to create an address. The interface is the IPv4 address converted to hexadecimal and added to the routing prefix 2002::/16.

Example:

```
Router(config)# ipv6 general-prefix my-prefix 6to4 ethernet0
```

When defining a general prefix based on a 6to4 interface, specify the 6to4 keyword and the interface-type interface-number argument.

When defining a general prefix based on an interface used for 6to4 tunneling, the general prefix will be of the form 2001:a.b.c.d::/48, where "a.b.c.d" is the IPv4 address of the interface referenced.

---

**QUESTION 475:**

Exhibit:



A network administrator recently redistributed RIP routes into an OSPF domain. However, the administrator wants to configure the network so that instead of 32 external type-5 LSAs flooding into the OSPF network, there is only one. What must the administrator do to accomplish this?

- A. Configure summarization on R1 with area 1 range 172.16.32.0 255.255.224.0
- B. Configure summarization on R1 with summary-address 172.16.32.0 255.255.224.0
- C. Configure area 1 as a stub area with area 1 stub
- D. Configure area 1 as a NSSA area with area 1 stub nssa

Answer: B

Explanation:

Sending the mask in the updates allows the network to support classless interdomain routing (CIDR). Summarization for the Internet, and variable-length subnet mask (VLSM) summarization for the autonomous system. Summarization requires manual configuration at either the routing process or the interface.

**QUESTION 476:**

Which two statements are true about RIP? (Choose two.)

- A. By default, RIP interfaces will only accept and send version 1 updates.
- B. By default, RIP interfaces will accept version 1 and 2 updates, but will only send version 1 updates.
- C. When configured with the version 1 command, RIP interfaces will only accept and send version 1 updates.
- D. When configured with the version 1 command, RIP interfaces will accept version 1 and 2 updates, but will only send version 1 updates.
- E. When configured with the version 2 command, RIP interfaces will only accept version 2 updates, but will send version 1 and version 2 updates.
- F. When configured with the version 2 command, RIP interfaces will accept version 1 and 2 updates, but will only send version 2 updates.

Answer: B, C

Explanation:

By default, a RIP process configured on a Cisco router sends only RIP v1 messages but listens to both RIP v1 and RIP v2. This default is changed with the version command. In this mode, the router sends and receives only RIP v2 messages.

Example:

```
Router(config)#router rip
```

```
Router(config-router)#version 2
```

```
Router(config-router)#network network-number
```

The router can also be configured to send and receive only RIP v1 messages:

```
Router(config)#router rip
```


```
Router(config-router)#version 1
```

```
Router(config-router)#network network-number
```

---

**QUESTION 477:**

Exhibit:



```

CK1# show ip route ospf
C    192.168.10.0/24 is directly connected, Serial1
C    10.10.10.0/24 is directly connected, Serial0
O*IA 0.0.0.0/0 [110/65] via 10.10.10.3, 00:04:11, Serial0

```

Refer to the exhibit and the output on CK1 . What is the OSPF area type configured for area 1?

- A. stub area
- B. totally stubby area
- C. NSSA
- D. NSSA totally stubby area

Answer: B

Explanation:

The following are several area types that are possible:

1. Standard area- A standard area can accept link updates and route summaries.
2. Backbone area (transit area) - When interconnecting multiple areas, the backbone area is the central entity to which all other areas connect. The backbone area is always Area 0. All other areas must connect to this area to exchange route information. The OSPF backbone has all the properties of a standard OSPF area.
3. Stub area  
-A stub area is an area that does not accept information about routes external to the autonomous system, the OSPF internetwork, such as routes from non-OSPF sources. If routers need to reach networks outside the autonomous system, they use a default route. A default route is noted as 0.0.0.0/0.
4. Totally stubby area- A totally stubby area is an area that does not accept external autonomous system (AS) routes and summary routes from other areas internal to the autonomous system. Instead, if the router needs to send a packet to a network external to the area, it sends it using a 0.0.0.0/0 default route. Totally stubby areas are a Cisco proprietary feature.
5. Not-so-stubby area (NSSA)- An NSSA is an area that is similar to a stub area but allows for importing external routes as Type 7 LSAs and translation of specific Type 7 LSA routes into Type 5 LSAs.

When we use the show ip route command we will get different route code:

O : OSPF networks from within the same area as the router

O IA : OSPF interarea. This is sent out by the ABRs and is created from the

summary link LSA (type 3 and type 4).

O N1 : OSPF NSSA external Type 1

O N2 : OSPF NSSA external 2

O E1 : OSPF external type 1

O E2 OSPF external type 2.

---

**QUESTION 478:**

Which command displays statistics on EIGRP hello, update, replies, and acknowledgements?

- A. debug eigrp packets
- B. show ip eigrp traffic
- C. show ip eigrp topology
- D. show ip eigrp neighbors

Answer: B

Explanation:

show ip eigrp traffic displays the number of eigrp packets send and receive.

---

**QUESTION 479:**

Which statement is true concerning EIGRP neighbors forming adjacencies?

- A. Hello and hold time values are configured locally and do not have to match.
- B. Different autonomous systems can be configured on each router.
- C. K-values are configured locally and do not have to match.
- D. Configuration of the variance command must match between peers.

Answer: A

Explanation:

EIGRP routers establish adjacencies with neighbor routers by using small hello packets. Hellos are sent by default every five seconds. An EIGRP router assumes that as long as it is receiving hello packets from known neighbors, those neighbors (and their routes) remain viable. By forming adjacencies, EIGRP routers do the following:

1. Dynamically learn of new routes that join their network
2. Identify routers that become either unreachable or inoperable
3. Rediscover routers that had previously been unreachable

Hello and hold time values are configured locally so to form adjacencies no need the same value but required the same K value.

---

**QUESTION 480:**

Which two are characteristics of the IS-IS protocol but not OSPF? (Choose two.)

- A. provides for network scalability by allowing the network to be separated into areas
- B. provides routing support for multiple network layer protocols
- C. three layers of hierarchical routing
- D. utilizes SPF algorithm
- E. forms adjacencies with all neighbors
- F. supports demand circuit routing

Answer: B, E

Explanation:

Similarities of IS-IS and OSPF

1. Classless
2. Link-state database and Dijkstra's algorithm
3. Hello Packets to form and maintain adjacencies
4. Use areas to form hierarchical topologies
5. Support address summarization between areas
6. Elect Designates Router on Multi-access network
7. Link-state representation, aging and metrics
8. Update, decision and flooding processes
9. Convergence capabilities
10. Deployed on ISP backbones


OSPF and IS-IS

1. IS-IS does not elect a Backup Designated Router
2. IS-IS repeats the election process whenever a new router become active
3. Every time the DR changes, a new set of LSPs must be flooded
4. IS-IS routers form adjacencies with all neighbors, not just the DR.
5. Each IS-IS router belongs to exactly one Level 2 area
6. IS-IS allows a more flexible approach to extending the backbone
7. One link-state packets per IS-IS router in one area, including redistributed prefixes, compared to many OSPF LSA.
8. OSPF supports 50 routers per area.
9. IS-IS supports 1000 routers per area.

---

**QUESTION 481:**

Exhibit:



Refer to the exhibit. What is the correct configuration to enable router Certkiller 4 to exchange RIP routing updates with router Certkiller 1 but not with router Certkiller 3?

- A. Certkiller 4(config)# interface fa0/0  
Certkiller 4(config-if)# neighbor 192.168.10.3  
Certkiller 4(config-if)# passive-interface fa0/0
- B. Certkiller 4(config)# router rip  
Certkiller 4(config-router)# neighbor 192.168.10.3  
Certkiller 4(config-router)# passive-interface fa0/0
- C. Certkiller 4(config)# interface fa0/0  
Certkiller 4(config-if)# neighbor 192.168.10.3  
Certkiller 4(config-if)# passive-interface 192.168.10.34
- D. Certkiller 4(config)# router rip  
Certkiller 4(config-router)# neighbor 192.168.10.34 no broadcast  
Certkiller 4(config-router)# passive-interface fa0/0

Answer: B

Explanation:

A passive interface receives updates, but does not send them. The passive-interface command can be used with all IP interior gateway protocols. It can be used with RIP, IGRP, EIGRP, OSPF, and IS-IS. The syntax of the command is as follows:

Router(config-router)#passive-interface type number

So Answer B is correct.

---

**QUESTION 482:**

Which statement best describes autoconfiguration for IPv6?

- A. A router connected to the local link sends a broadcast periodically to advertise the

subnet and the host appends its locally configured IPv4 address to create a unique IPv6 address.


- B. A router connected to the local link sends its 64 bit prefix address and the host appends its MAC address to create its unique IPv6 address.
- C. A host will contact its local router to receive its unique IPv6 address.
- D. The local DHCP server will detect a new host has been added to the LAN segment and automatically send a unique IPv6 address to the new host.

Answer: B

Explanation: The local or directly connected router sends out the prefix of the local link and the router's default route. This is sent to all the nodes on the wire, allowing them to autoconfigure their own IPv6 addresses. The local router provides the 48-bits global prefix and the SLA or subnet information to each end system. The end system simply adds its own Layer 2 address, which is unique because it is the burned-in address (MAC address) taken from the interface card. This MAC address, when appended to the 48-bits global address and the 16-bit subnet address, makes up the unique 128-bits IPv6 end system address.

**QUESTION 483:**

Exhibit:



```
Certkiller1# show ip route ospf
C    30.30.30.0/24 is directly connected, Serial1
C    10.10.10.0/24 is directly connected, Serial0
O    20.0.0.0/24 is subnetted, 1 subnets
O      20.20.20.0 [110/1562] via 30.30.30.3, 00:00:08, Serial1
      [110/1562] via 10.10.10.2, 00:00:08, Serial0
O    40.0.0.0/24 is subnetted, 1 subnets
O*IA  40.40.40.0 [110/782] via 30.30.30.3, 00:00:08, Serial1
O*IA  0.0.0.0/0 [110/782] via 30.30.30.3, 00:00:09, Serial1
```

Refer to the exhibit and the output within it. What type of OSPF area is area 1?

- A. stub area
- B. totally stubby area
- C. NSSA
- D. NSSA totally stubby area

Answer: A

Explanation:

The following are several area types that are possible:

1. Standard area- A standard area can accept link updates and route summaries.
2. Backbone area (transit area) - When interconnecting multiple areas, the backbone area is the central entity to which all other areas connect. The backbone area is always Area 0. All other areas must connect to this area to exchange route information. The OSPF backbone has all the properties of a standard OSPF area.
3. Stub area -A stub area is an area that does not accept information about routes external to the autonomous system, the OSPF internetwork, such as routes from non-OSPF sources. If routers need to reach networks outside the autonomous system, they use a default route. A default route is noted as 0.0.0.0/0.
4. Totally stubby area- A totally stubby area is an area that does not accept external autonomous system (AS) routes and summary routes from other areas internal to the autonomous system. Instead, if the router needs to send a packet to a network external to the area, it sends it using a 0.0.0.0/0 default route. Totally stubby areas are a Cisco proprietary feature.
5. Not-so-stubby area (NSSA)- An NSSA is an area that is similar to a stub area but allows for importing external routes as Type 7 LSAs and translation of specific Type 7 LSA routes into Type 5 LSAs.

When we use the show ip route command we will get different route code:

- O : OSPF networks from within the same area as the router
- O IA : OSPF interarea. This is sent out by the ABRs and is created from the summary link LSA (type 3 and type 4).
- O N1 : OSPF NSSA external Type 1
- O N2 : OSPF NSSA external 2
- O E1 : OSPF external type 1
- O E2 OSPF external type 2.

---

**QUESTION 484:**

Which two statements are true concerning IPv6? (Choose two.)

- A. Mobil IP is built into IPv6 by default.
- B. The leading zeros in an address format are mandatory
- C. Like IPv4, IPv6 broadcasts are sent to all nodes on a LAN segment.
- D. IPSec is mandatory and built into IPv6.

Answer: A, D

Explanation: IPv6 offers the following benefits and features:


1. Larger address space
2. Unicast and Multicast addressing
3. Address aggregation
4. Autoconfiguration

5. Renumbering
6. A simple and efficient header
7. Security
8. Mobility
9. Options for transitioning from IPv4 to IPv6
10. Routing Protocols

---

**QUESTION 485:**

Exhibit:



Refer to the network topology diagram shown in the exhibit. The routing table of Certkiller 3 contains all the proper OSPF routes except the 10.4.4.0/24 OSPF route. Which show command on Certkiller 3 will be the most useful to determine the cause of this problem?

- A. Perform `show ip ospf virtual-link` to examine if the virtual link between Certkiller 3 and Certkiller 4 is configured and in the up state.
- B. Perform `show run` to examine if area 3 is configured as a stubby area.
- C. Perform `show run` to examine if area 3 is configured as a nssa area.
- D. Perform `show ip ospf neighbor` to examine which router is elected as the DR on the e2 interface.
- E. Perform `show ip ospf database` to examine the OSPF database entries.

Answer: A

Explanation: A Virtual link is a virtual connection to a remote area that does not have any connections to the backbone (area 0). Typically, this is because the network has become segmented. Although OSPF treats this link as a direct, single-hop connection to the backbone area, it is a virtual connection that tunnels through the network.

1. `show ip ospf border-routers` - Displays the internal OSPF routing table entries to an ABR.
2. `show ip ospf virtual-links` - Displays parameters about the current state of OSPF virtual links.
3. `show ip ospfprocess-id` - Displays information about each area to which the router

is connected, and indicates whether the router is an ABR, ASBR, or both. The process ID is a user defined identification parameter. It is locally assigned and can be any positive integer number. The number used here is the number assigned administratively when enabling the OSPF routing process.

---

**QUESTION 486:**

Which two are characteristics of the IS-IS protocol but not OSPF? (Choose two.)

- A. provides for network scalability by allowing the network to be separated into areas
- B. provides routing support for multiple network layer protocols
- C. three layers of hierarchical routing
- D. utilizes SPF algorithm
- E. forms adjacencies with all neighbors
- F. supports demand circuit routing

Answer: B, E

Explanation:

Similarities of IS-IS and OSPF

1. Classless
2. Link-state database and Dijkstra's algorithm
3. Hello Packets to form and maintain adjacencies
4. Use areas to form hierarchical topologies
5. Support address summarization between areas
6. Elect Designates Router on Multi-access network
7. Link-state representation, aging and metrics
8. Update, decision and flooding processes
9. Convergence capabilities
10. Deployed on ISP backbones


OSPF and IS-IS

1. IS-IS does not elect a Backup Designated Router
2. IS-IS repeats the election process whenever a new router become active
3. Every time the DR changes, a new set of LSPs must be flooded
4. IS-IS routers form adjacencies with all neighbors, not just the DR.
5. Each IS-IS router belongs to exactly one Level 2 area
6. IS-IS allows a more flexible approach to extending the backbone
7. One link-state packets per IS-IS router in one area, including redistributed prefixes, compared to many OSPF LSA.
8. OSPF supports 50 routers per area.
9. IS-IS supports 1000 routers per area.

---

**QUESTION 487:**

Exhibit:



Refer to the exhibit. BGP sessions are established between all autonomous systems. AS 100 receives updates about network 200.200.200.0/24 from routers Certkiller 2, Certkiller 3, Certkiller 4, and Certkiller 5 with different MED values. Which router will be chosen as a next hop to reach the network 200.200.200.0/24 if all other attributes are the same?

- A. Certkiller 2
- B. Certkiller 3
- C. Certkiller 4
- D. Certkiller 5

Answer: A

Explanation:

The Multiple-exit-discriminator (MED) attribute is an optional nontransitive attribute, type code 4. MED informs external neighbors about the preferred path into an AS that has multiple entry points. A lower MED is preferred over a higher MED.

Unlike Local Preference, the MED attribute is exchanged between autonomous systems, but a MED attribute that comes into an AS does not leave the AS. When an update enters the AS with a certain MED value, that value is used for decision making within the AS. When BGP forwards the routing update to another AS, the MED is reset to zero. This is true unless the outgoing MED is set to a specific value.

By default, BGP compares only the MED attributes of routes coming from neighbors in the same external AS, such as AS 200 in the example. This means that RTA will compare the MED attribute coming from Certkiller 2 (100) only to the MED attribute coming from Certkiller 3 (200). Even though the update coming from Certkiller 2 has the lowest MED value, Certkiller 1 will choose Certkiller 2 as the best path for reaching network 200.200.200.0/24.

---

### QUESTION 488:

Exhibit:

```
CertKiller1#sh ip ospf_neighbor
```

| Neighbor ID   | Pri | State    | Dead Time | Address  | Interface   |
|---------------|-----|----------|-----------|----------|-------------|
| 10.200.200.13 | 1   | FULL/BDR | 00:00:33  | 10.1.1.3 | Ethernet0/0 |

```
CertKiller3#sh ip ospf_neighbor
```

| Neighbor ID | Pri | State   | Dead Time | Address  | Interface   |
|-------------|-----|---------|-----------|----------|-------------|
| 172.31.1.1  | 2   | FULL/DR | 00:00:31  | 10.1.1.1 | Ethernet0/0 |

The Certkiller 1 and Certkiller 3 routers are OSPF neighbors over the Ethernet 0/0 connection. Based on the show ip ospf neighbor output from the Certkiller 1 and Certkiller 3 routers, which statement is true?

- A. Certkiller 1 is the DR because it has a higher OSPF router priority.
- B. Certkiller 1 is the DR because it has a lower OSPF router ID.
- C. Certkiller 3 is the DR because it has a higher OSPF router priority.
- D. Certkiller 3 is the DR because it has a lower OSPF router ID.
- E. Both Certkiller 1 and Certkiller 3 are using the default OSPF router priority.

Answer: A


Explanation: Router Certkiller 1 and Certkiller 3 are connected to each other so these are neighbors, While selecting the DR from OSPF neighbor, it checks the priority value first. That router having highest priority value wins the race. So in comparing two routers, Certkiller 1 has highest priority so it should be the DR.

---

**QUESTION 489:**

Exhibit.

## 642-801



Refer to the exhibit. Router Certkiller 1 is not able to ping the loopback Lo0 on Router Certkiller 2. What could the problem be?

- A. The loopback addresses have different subnet masks.
- B. Router Certkiller 1 is using RIP version 1 while Router Certkiller 2 is using RIP version 2.
- C. Router Certkiller 1 is using plain text authentication while Router Certkiller 2 is using encrypted authentication.
- D. MD5 authentication is not possible with RIP.
- E. The key strings do not match.
- F. The key chains do not match.

Answer: C

Explanation:

The implementation of RIP v2 message authentication by Cisco includes the choice of a simple password or MD5 authentication, and the option of defining multiple keys, or passwords on a key chain. The router may then be configured to use different keys at different times. Plain text authentication is the default setting in every RIP v2 packet.

Default Authentication is plain text, in Certkiller 1 authentication type is not defined but in Certkiller 2 MD5 authentication is defined so authentication is mismatched.

---

### **QUESTION 490:**

What are two routing protocols defined by the OSI protocol suite at the network layer? (Choose two.)

- A. End System-to-End System
- B. Routing Information Protocol
- C. Interior Gateway Routing Protocol
- D. End System-to-Intermediate System
- E. Intermediate System-to-Intermediate System

Answer: D,E

Explanation:

ISO has developed standards for two types of network protocols used in routing. These two protocols are ES-IS and IS-IS. The ES-IS discovery protocols are used for routing between end systems and intermediate systems. The IS-IS routing protocols are used for hierarchical routing between intermediate systems.


ES-IS is an analogous to Address Resolution Protocol (ARP) in IP. Although not technically a routing protocol, ES-IS is included here because it is commonly used with routing protocols to provide end-to-end data movement through an internetwork. Routing between end systems and intermediate systems is sometimes referred to as Level 0 routing.

With IS-IS routing, OSI distinguishes between Level 1, Level 2, and Level 3 routing to simplify router design and operation. Level 1 ISs communicate with other Level 1 ISs in the same area. Level 2 ISs route between Level 1 areas and form an intradomain routing backbone. Level 3 routing is done between separate domains. Hierarchical routing simplifies backbone design, because Level 1 ISs only need to know how to get to the nearest Level 2 IS.

---

**QUESTION 491:**

Exhibit:



```

CertKiller4#show running-configuration
|
router bgp 200
no synchronization
neighbor 20.20.43.3 remote-as 200
neighbor 10.10.45.5 remote-as 100
neighbor 10.10.45.5 route-map LP1 in
|
access list 10 permit 100.100.100.0.0.255.255.255
|
route-map LP1 permit 10
match ip address 101
set local-preference 300
|

```

Refer to exhibit. BGP is configured on all routers, synchronization is turned off, and none of the default attributes have been changed except the local preference attribute on Certkiller 4. Which path will be preferred by Certkiller 2 to reach the network 100.100.100.0/24?

- A. Certkiller 2 - Certkiller 3 - Certkiller 4 - Certkiller 5 because it has a lower admin distance
- B. Certkiller 2 - Certkiller 3 - Certkiller 4 - Certkiller 5 because it has a higher local preference
- C. Certkiller 2 - Certkiller 1 because it has the shortest AS-path
- D. Certkiller 2 - Certkiller 1 because it has a lower local preference

Answer: B

Explanation:

A BGP speaker can receive updates from multiple autonomous systems that describe different paths to the same destination. It must then choose the single best path for reaching that destination. Once chosen, BGP propagates the best path to its neighbors. The decision is based on the value of attributes, such as Next Hop or Local Preference, that the update contains and other configurable BGP factors. The following sections provide an overview of these key attributes that BGP uses in the decision making process:

1. Next Hop
2. AS\_Path
3. Atomic Aggregate
4. Aggregator
5. Local Preference

6. Weight

7. Multiple Exit Discriminator (MED)


8. Origin

As per configuration, all routers has default value of Next Hop, AS\_Path, Automatic Affregate, Aggregator but Local preference is higher then default so it select the Certkiller 2- Certkiller 3- Certkiller 4- Certkiller 5 path.

---

**QUESTION 492:**

Exhibit:



Router Certkiller 5 is configured with the EIGRP variance 2 command. What path will Router Certkiller 5 take to reach Router Certkiller 1?

- A. only Certkiller 5- Certkiller 4- Certkiller 1
- B. only Certkiller 5- Certkiller 2- Certkiller 1
- C. only Certkiller 5- Certkiller 3- Certkiller 1
- D. both Certkiller 5- Certkiller 2- Certkiller 1 and Certkiller 5- Certkiller 3- Certkiller 1
- E. Certkiller 5. both Certkiller 5- Certkiller 2- Certkiller 1 and Certkiller 5- Certkiller 4- Certkiller 1
- F. all available paths.

Answer: D

Explanation: Every routing protocol supports equal cost path load balancing. In addition, Interior Gateway Routing Protocol (IGRP) and EIGRP also support unequal cost path load balancing. Use the variance n command in order to instruct the router to include routes with a metric of less than n times the minimum metric route for that destination. The variable n can take a value between 1 and 128. The default is 1, which means equal cost load balancing.

When variance 2 is configured, router can use 2 different path for the destination. That path will select, which have low metric.

Example of Variance configuration:

```
router eigrp 1
```

network x.x.x.x  
variance 2

**QUESTION 493:**

**DRAG DROP**

Match the appropriate IS-IS routing terms to its correct definition.

Routing terms, select from these

- Level 0 routing
- Level 1/2 routing
- Level 1 routing
- Level 2 routing
- Level 3 routing

Definitions

- Between ES's and IS's on the same subnet. OSI routing begins at this level.
- Routing between separate domains. Similar to BGP
- Between IS's on the same area. Also called area-routing.
- Called inter-area routing

Routing terms, place here

- Place here
- Place here
- Place here
- Place here

Answer:

Routing terms, select from these

- Level 2 routing

Definitions

- Between ES's and IS's on the same subnet. OSI routing begins at this level.
- Routing between separate domains. Similar to BGP
- Between IS's on the same area. Also called area-routing.
- Called inter-area routing

Routing terms, place here

- Level 0 routing
- Level 3 routing
- Level 1/2 routing
- Level 1 routing

Explanation:

- i. Level 0 Routing between ESs and ISs on the same subnet
- ii. Level 1 routing between ISs within same area
- iii. Level 2 routing between different areas within the same domain
- iv. Level 3 routing between separate domains.

**QUESTION 494:**

Which three types of OSPF route entries can be found in the routing table of an internal OSPF router within an OSPF stubby area? (Choose three.)

- A. O
- B. O IA
- C. O\* IA
- D. O E1
- E. O E2
- F. O N1

Answer: A,B,C

Explanation:

The following are several area types that are possible:

1. Standard area- A standard area can accept link updates and route summaries.
2. Backbone area (transit area) - When interconnecting multiple areas, the backbone area is the central entity to which all other areas connect. The backbone area is always Area 0. All other areas must connect to this area to exchange route information. The OSPF backbone has all the properties of a standard OSPF area.
3. Stub area -A stub area is an area that does not accept information about routes external to the autonomous system, the OSPF internetwork, such as routes from non-OSPF sources. If routers need to reach networks outside the autonomous system, they use a default route. A default route is noted as 0.0.0.0/0.
4. Totally stubby area- A totally stubby area is an area that does not accept external autonomous system (AS) routes and summary routes from other areas internal to the autonomous system. Instead, if the router needs to send a packet to a network external to the area, it sends it using a 0.0.0.0/0 default route. Totally stubby areas are a Cisco proprietary feature.
5. Not-so-stubby area (NSSA)- An NSSA is an area that is similar to a stub area but allows for importing external routes as Type 7 LSAs and translation of specific Type 7 LSA routes into Type 5 LSAs.

When we use the show ip route command we will get different route code:

- O : OSPF networks from within the same area as the router
- O IA : OSPF interarea. This is sent out by the ABRs and is created from the summary link LSA (type 3 and type 4).
- O N1 : OSPF NSSA external Type 1
- O N2 : OSPF NSSA external 2
- O E1 : OSPF external type 1
- O E2 OSPF external type 2.

---

**QUESTION 495:**

Which three are characteristics of IPv6? (Choose three.)

- A. An IPv6 address is 128 bits long.
- B. An IPv6 header is 20 bits long.
- C. An IPv6 header contains the next header field.
- D. An IPv6 header contains the protocol field.
- E. IPv6 routers send RA messages.
- F. An IPv6 header contains the header checksum field.

Answer: A,C,E

Explanation:

Unlike the usual dotted-decimal format of the IPv4 address, IPv6 is represented by hexadecimal numbers. A hexadecimal number is equivalent to four bits, also known as a nibble because it is half a byte, and is numbered 0-9 and A-F. A represents 10 and F represents 15, and they are not case-sensitive. The IPv6 address is a 32-digit hexadecimal numeric value, in eight four-digit clusters, known as fields, separated by colons (:), representing the 128-bit address. Here is an example of a valid IPv6 address: 1041:0000:130B:0000:0000:09C0:586C:1305.

There are some techniques used to shorten the IPv6 address. One of these techniques is to omit leading 0s in the address field, so 0000 can be compressed to just 0 and 09C0 can be compressed to 9C0. You can omit leading 0s but not trailing 0s. The previous IPv6 address example could be shortened to 1041:0:130B:0:0:9C0:586C:1305. Another technique is to use double colons (::) to represent a contiguous block of 0s. Again, the previous IPv6 address can be further shortened to 1041:0:130B::9C0:586C:1305.

---

**QUESTION 496:**

Exhibit:

```
CertKiller1#show ip eigrp topology all-links
IP-EIGRP Topology Table for AS(65001)/ID(172.16.10.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       _ - reply status   s - sin Status

P 172.16.3.128/25, 2 successors, FD is 30720, serno 9
   via 172.16.10.2 (30720/28160), FastEthernet0/1
   via 172.16.3.2 (30720/28160), FastEthernet0/3
P 192.168.0.0/24, 1 successors, FD is 156160, serno 16
   via 172.16.3.2 (156160/120256), FastEthernet0/3
   via 172.16.10.2 (158720/156160), FastEthernet0/1
P 172.16.10.0/24, 1 successors, FD is 28160, serno 1
   via Connected, FastEthernet0/1
P 172.16.0.0/30, 1 successors, FD is 20514560, serno 15
   via 172.16.1.1 (20514560/20512000), FastEthernet0/2
   via 172.16.10.2 (20517120/20514560), FastEthernet0/1
P 172.16.1.0/24, 1 successors, FD is 28160, serno 2
   via Connected, FastEthernet0/2
P 172.16.2.0/24, 1 successors, FD is 30720, serno 8
   via 172.16.10.2 (30720/28160), FastEthernet0/1
   via 172.16.3.2 (33280/30720), FastEthernet0/3
P 172.16.3.0/25, 1 successors, FD is 28160, serno 3
   via Connected, FastEthernet0/3

CertKiller1#
```

Which three networks is the router at 172.16.10.2 directly connected to? (Choose three)

- A. 172.16.0.0/30
- B. 172.16.1.0/24
- C. 172.16.2.0/24
- D. 172.16.3.0/25
- E. 172.16.3.128/25
- F. 192.168.0.0/24

Answer: C, E, F

Explanation: show ip eigrp topology all-links command displays all routes, not just feasible successors, in the EIGRP topology. According to exhibit Answer C,E,F are correct.

---

**QUESTION 497:**

Which IS-IS NET represents a locally administered private address?

- A. 39.0040.0010.0c99.1112.00
- B. 45.0004.0000.0d35.4554.00
- C. 47.0010.0000.0a11.3564.00
- D. 49.0001.0000.0c12.3456.00

Answer: D

---

**QUESTION 498:**

Which three types of OSPF route entries can be found in the routing table of an internal OSPF router within an OSPF NSSA area? (Choose three.)

- A. O
- B. O N2
- C. O\* N2
- D. O E1
- E. O E2

Answer: A,B,C

Explanation:

The following are several area types that are possible:

1. Standard area- A standard area can accept link updates and route summaries.
2. Backbone area (transit area) - When interconnecting multiple areas, the backbone area is the central entity to which all other areas connect. The backbone area is always Area 0. All other areas must connect to this area to exchange route information. The OSPF backbone has all the properties of a standard OSPF area.
3. Stub area  
-A stub area is an area that does not accept information about routes external to the autonomous system, the OSPF internetwork, such as routes from non-OSPF sources. If routers need to reach networks outside the autonomous system, they use a default route. A default route is noted as 0.0.0.0/0.
4. Totally stubby area- A totally stubby area is an area that does not accept external autonomous system (AS) routes and summary routes from other areas internal to the autonomous system. Instead, if the router needs to send a packet to a network external to the area, it sends it using a 0.0.0.0/0 default route. Totally stubby areas are a Cisco proprietary feature.
5. Not-so-stubby area (NSSA)- An NSSA is an area that is similar to a stub area but allows for importing external routes as Type 7 LSAs and translation of specific Type 7 LSA routes into Type 5 LSAs.


When we use the show ip route command we will get different route code:

- O : OSPF networks from within the same area as the router
- O IA : OSPF interarea. This is sent out by the ABRs and is created from the summary link LSA (type 3 and type 4).
- O N1 : OSPF NSSA external Type 1
- O N2 : OSPF NSSA external 2
- O E1 : OSPF external type 1
- O E2 OSPF external type 2.

---

**QUESTION 499:**

Exhibit:



Refer to the exhibit. Routers Certkiller 2, Certkiller 3, Certkiller 4, and Certkiller 5 have OSPF enabled. What should be configured on the routers in area 1 to ensure that all default summary routes and redistributed EIGRP routes will be forwarded from Certkiller 6 to area 1, and only a default route for all other OSPF routes will be forwarded from Certkiller 5 to area 1.

- A. Certkiller 5(config-router)# area 1 stubR6(config-router)# area 1 stub
- B. Certkiller 5(config-router)# area 1 stub no-summaryR6(config-router)# area 1 stub
- C. Certkiller 5(config-router)# area 1 nssa Certkiller 6(config-router)# area 1 nssa
- D. Certkiller 5(config-router)# area 1 nssa no-summaryR6(config-router)# area 1 nssa

Answer: D

Explanation:

Not-so-stubby area (NSSA)- An NSSA is an area that is similar to a stub area but allows for importing external routes as Type 7 LSAs and translation of specific Type 7 LSA routes into Type 5 LSAs.

Stub area -A stub area is an area that does not accept information about routes external to the autonomous system, the OSPF internetwork, such as routes from non-OSPF sources. If routers need to reach networks outside the autonomous system, they use a default route. A default route is noted as 0.0.0.0/0. Router 5 Should configured for no-summary of area 1 of nssa and router 6 should forward the summary routes and redistributed EIGRP routes to area 1.

---

### QUESTION 500:

Which two statements are EIGRP characteristics? (Choose two.)

- A. Updates are sent as multicast.
- B. Updates are sent as broadcast.
- C. Metric values are represented in a 32-bit format for granularity.
- D. LSAs are sent to adjacent neighbors .


Answer: A,C

Explanation:

Despite being compatible with IGRP, EIGRP uses a different metric calculation and hop-count limitation. EIGRP scales the IGRP metric by a factor of 256. That is because EIGRP uses a metric that is 32-bits long, and IGRP uses a 24-bit metric. By multiplying or dividing by 256, EIGRP can easily exchange information with IGRP. As well as EIGRP sends the updates to multicast address 224.0.0.10

### QUESTION 501:

Exhibit:



Refer to the exhibit. Which configuration command sequence would prevent the propagation of updates for network 192.168.236.0/24 to the Certkiller 2 router?

- A. Certkiller 2(config)# access-list 1 deny 192.168.236.0  
Certkiller 2(config)# access-list 1 permit any  
Certkiller 2(config)# interface serial 0/0  
Certkiller 2(config-if)# ip access-group 1 in
- B. Certkiller 1(config)# access-list 1 deny 192.168.236.0  
Certkiller 1(config)# access-list 1 permit any  
Certkiller 1(config)# interface serial 0/0  
Certkiller 1(config-if)# ip access-group 1 out
- C. Certkiller 2(config)# access-list 1 deny 192.168.236.0  
Certkiller 2(config)# access-list 1 permit any  
Certkiller 2(config)# router rip  
Certkiller 2(config-router)# distribute-list 1 out serial 0/0
- D. Certkiller 1(config)# access-list 1 deny 192.168.236.0  
Certkiller 1(config)# access-list 1 permit any  
Certkiller 1(config)# router rip  
Certkiller 1(config-router)# distribute-list 1 out serial 0/0
- E. Certkiller 2(config)# router rip

```
Certkiller 2(config-router)# passive-interface s0/0  
F. Certkiller 1(config)# router rip  
Certkiller 1(config-router)# passive-interface s0/0
```

Answer: D

Explanation:

Use the distribute-list command to pick and choose which routing updates a router will send or receive. By referencing an access list, the distribute-list creates a route filter. This is a set of rules that precisely controls what routes a router will send or receive in a routing update. This command is available for all IP routing protocols and can be applied to either inbound or outbound routing updates. When applied to inbound updates, the syntax for configuring a route filter is as follows:

```
Router(config-router)#distribute-list access-list-number in  
[interface-name]
```

When applied to outbound updates, the syntax can be more complicated as shown in the following:

```
Router(config-router)#distribute-list access-list-number out  
[interface-name | routing-process | as-number]
```

---

### QUESTION 502:

Exhibit:

```
CertKiller3 # show ip route  
...  
C   10.1.3.0/24 is directly connected, Serial2  
D   10.1.2.0/24 [90/10537472] via 10.1.1.2, 00:23:24, Serial1  
D   10.0.0.0/8 is a summary, 00:22:20, Null0  
C   10.1.1.0/24 is directly connected, Serial1  
S   192.168.20.0/24 is directly connected, Ethernet0
```

Refer to the exhibit. If the configuration shown below is added to Certkiller 3, which three route entries will EIGRP advertise to neighboring routers?(Choose three.)  
router eigrp 10 network 10.0.0.0 eigrp stub

- A. 10.1.3.0/24
- B. 10.1.2.0/24
- C. 10.0.0.0/8
- D. 10.1.1.0/24
- E. 192.168.20.0/24

Answer: A,C,D


Explanation: Remote router can configure as an EIGRP stub router. A stub router is typically used on small capacity routers in a hub-and-spoke WAN environment.

The stub router in EIGRP is similar to the concepts of on Demand Routing. As in ODR, the router in EIGRP network has no other neighbors and accesses the network through a distribute layer router. It is not necessary, therefore, for this remote router to have a complete routing table that may overwhelm its limited resources. The remote router needs only a default route to the distribute router that can serve all its needs. When router Certkiller 3 will configure as a stub router, there will be 10.1.3.0/24, 10.0.0.0/8, 10.1.1.0/24 entries.

---

**QUESTION 503:**

Exhibit:



Refer to the exhibit. Routers RTA and RTB are running BGP but the session is active. What command needs to be added to establish the BGP session?

- A. ip route 10.10.10.1 255.255.255.255 s0/0 ip route 10.10.10.1 255.255.255.255 s0/1
- B. network 10.10.10.0
- C. neighbor 10.10.10.1 next-hop-self
- D. no synchronization

Answer: A

Explanation: while configuring ebgp-multihop, you need to set the static route of neighbor of remote as.

---

**QUESTION 504:**

Which two statements are true about IBGP neighbor relationships? (Choose two.)

- A. An EGP or static routing is required between IBGP neighbors.
- B. A full-mesh IBGP requires that neighbor relationships be established between all BGP enabled routers in the autonomous system.
- C. IBGP neighbors must be in different autonomous systems.
- D. The BGP split-horizon rule specifies that routes learned via EBGP are never propagated to other IBGP peers.
- E. The BGP split horizon rule specifies that routes learned via IBGP are never propagated to other IBGP peers.

Answer: B,E


Explanation:

IBGP - Occurs between routers in the same AS, and is used to coordinate and synchronize routing policy within the AS. Neighbors may be located anywhere in the AS, even several hops away from one another. An IBGP session typically occurs between routers in the same AS in an ISP. Full-mesh IBGP requires that neighbor relationships be established between all BGP enabled router in the AS.

---

**QUESTION 505:**

Exhibit:



Refer to the exhibit. Which two statements are true? (Choose two.)

- A. Certkiller 2 will receive only the summary address of 2.0.0.0/8.
- B. Certkiller 2 will receive two routes, the manual summary address of 2.1.0.0/16 and the automatic, summarized classful address of 2.0.0.0/8.
- C. The summary route will be sourced from the Certkiller 1 s0/0 interface in the Certkiller 1 routing table.
- D. The summary route will be sourced from the Certkiller 1 Null0 interface in the Certkiller 1 routing table.
- E. The summary route will be assigned the administrative distance of 90 in the Certkiller 1 routing table.
- F. The summary route will be assigned the administrative distance of 5 in the Certkiller 2 routing table.

Answer: A,D

Explanation:

EIGRP also enables manual configuration of a prefix to use as a summary address. Manual summary routes are configured on a per-interface basis. The interface that will propagate the route summary must first be selected and then defined with the ip summary-address eigrp command, which has the following syntax:


```
Router(config-if)#ip summary-address eigrp autonomous-system-number  
ip-address mask administrative-distance
```

D 2.1.0.0/16 is a summary, 00:00:22, Null0

Notice that the summary route is sourced from Null0, and not an actual interface

**QUESTION 506:**

Exhibit:



Refer to the exhibit. A network administrator is troubleshooting IS-IS routing. Currently, Certkiller 1 is not receiving Link State Packets (LSPs) from Certkiller 2. What is the problem?

- A. Certkiller 2 must be configured with the is-type level-1 command under router isis.
- B. Certkiller 1 must be configured with the isis circuit-type level-1 under interface s0.
- C. Certkiller 1 and Certkiller 2 are configured in different areas.
- D. Certkiller 1 and Certkiller 2 should both be configured for is-type level-2 under router isis.

Answer: C

Explanation:

**Area ID** The area ID is typically a one- or two-octet field that precedes the system ID. The area ID is used to signify the area the router belongs to. The area ID can span up to two octets if need be. For the BSCI, we will concentrate on a one-octet area ID. For the NSAP addressing scheme used with IS-IS, the area ID is everything to the left of the system ID.

**System ID** The system ID is used to identify the router. It is similar to the router ID in OSPF. The system ID can be up to eight octets in length. Cisco supports only a six-octet field for the system ID. Whether or not the vendor you are using supports more than six octets, the number of octets used must be the same throughout the IS-IS routing domain. Normally, the system ID will be set to one of the MAC addresses on the router. This is not a rule, and the system ID can be represented in a number of ways. The system ID that is used for a router must be unique throughout the IS-IS routing domain. There must be only one system ID set per router, regardless of the number of routing processes. The system ID must be common

across all routing processes on the same router.

SEL The NSAP selector byte (SEL) is a one-octet field that represents the service being offered at the network level of the device. For our implementation of IS-IS, the SEL will always be set to 00. 00 represents the router. A good comparison to the SEL is the TCP or UDP port number being associated with an IP address, forming a socket that identifies a unique application instance on the internetwork. Basically, it tells you what service is being referenced for the particular address.


As previously mentioned, the NET can come in and be entered in a variety of NSAP formats, the last eight bytes of each serving the same function. Let's go ahead and take a look at the three most common formats

In Between two routers, Area ID is different, one router has 192 area ID and another router has 1921 Area ID so different in Area ID.

---

**QUESTION 507:**

Exhibit:



Refer to the exhibit. Which statement is true?

- A. Certkiller 1 will redistribute the RIP routes into the NSSA as type 7 LSAs. Certkiller 2 will translate the type 7 LSAs into type 5 LSAs and flood them throughout the OSPF backbone.
- B. Certkiller 1 will redistribute the RIP routes into the NSSA as type 7 LSAs. Certkiller 2 will flood the type 7 LSAs throughout the backbone.
- C. Certkiller 1 will redistribute the RIP routes into the NSSA as type 5 LSAs. Certkiller 2 will flood the type 5 LSAs throughout the backbone.
- D. Certkiller 1 will redistribute the RIP routes into the NSSA as type 5 LSAs. Certkiller 2 will translate the type 5 LSAs into type 7 LSAs and flood them throughout the OSPF backbone.
- E. Certkiller 1 will not redistribute the RIP routes into the NSSA.

Answer: A

Explanation:

Not-so-stubby area (NSSA)- An NSSA is an area that is similar to a stub area but allows for importing external routes as Type 7 LSAs and translation of specific Type 7 LSA routes

into Type 5 LSAs.

Stub area -A stub area is an area that does not accept information about routes external to the autonomous system, the OSPF internetwork, such as routes from non-OSPF sources. If routers need to reach networks outside the autonomous system, they use a default route. A default route is noted as 0.0.0.0/0. Router 5 Should configured for no-summary of area 1 of nssa and router 6 should forward the summary routes and redistributed EIGRP routes to area 1. So, in Certkiller 1 router redistribute the RIP into the NSSA type 7 and in router Certkiller 2 translate the type 7 into type 5.

---

**QUESTION 508:**

Which two statements about the EIGRP DUAL process are correct? (Choose two.)

- A. An EIGRP route will go active if there are no successors or feasible successors in the EIGRP topology table.
- B. An EIGRP route will go passive if there are no successors in the EIGRP topology table.
- C. DUAL will trigger an EIGRP query process while placing the flapping routes in the holddown state.
- D. A feasible successor in the EIGRP topology table can become the successor only after all the query requests have been replied to.
- E. The stuck in active state is caused when the wait for the query replies have timed out.
- F. EIGRP queries are sent during the loading state in the EIGRP neighbor establishment process.

Answer: A,E

Explanation:

EIGRP routers converge quickly because they rely on a state-of-the-art routing algorithm called the Diffusing Update Algorithm (DUAL). DUAL guarantees loop-free operation at every instant during route computation and allows all routers involved in a topology change to synchronize at the same time.

EIGRP routers keep route and topology information readily available in RAM so they can react quickly to changes. Like OSPF, EIGRP keeps this information in several tables, or databases. The following terms are related to EIGRP and its tables and are used throughout this module:

1. Neighbor table- Each EIGRP router maintains a neighbor table that lists adjacent routers. This table is comparable to the adjacency database used by OSPF. There is a neighbor table for each protocol that EIGRP supports.
2. Topology table- Every EIGRP router maintains a topology table for each configured network protocol. This table includes route entries for all destinations that the router has learned. All learned routes to a destination are maintained in the topology table.
3. Routing table- EIGRP chooses the best routes to a destination from the topology table and places these routes in the routing table. Each EIGRP router maintains a routing table for each network protocol.
4. Successor- A successor is a route selected as the primary route to use to reach a destination. Successors are the entries kept in the routing table. Multiple successors for a

destination can be retained in the routing table.

5. Feasible successor- A feasible successor is a backup route. These routes are selected at the same time the successors are identified, but are kept in the topology table. Multiple feasible successors for a destination can be retained in the topology table

---

**QUESTION 509:**

Exhibit:

```
CertKiller1 # show ip route
...
C 10.1.3.0/24 is directly connected, Serial2
D 10.1.2.0/24 [90/10537472] via 10.1.1.2, 00:23:24, Serial1
D 10.0.0.0/8 is a summary, 00:22:20, Null0
C 10.1.1.0/24 is directly connected, Serial1
S 192.168.20.0/24 is directly connected, Ethernet0
```

Refer to the exhibit. Which command was used to inject the summary route into the route table?

- A. Router(config-if)# ip summary-address eigrp 101 10.0.0.0 255.0.0.0
- B. Router(config)# ip route 10.0.0.0 255.0.0.0 null 0
- C. Router(config-router)# no auto-summary
- D. Router(config-router)# default-information originate

Answer: A

Explanation:


EIGRP also enables manual configuration of a prefix to use as a summary address. Manual summary routes are configured on a per-interface basis. The interface that will propagate the route summary must first be selected and then defined with the ip summary-address eigrp command, which has the following syntax:

```
Router(config-if)#ip summary-address eigrp autonomous-system-number
ip-address mask administrative-distance
```

---

**QUESTION 510:**

Exhibit:



Refer to the exhibit. By default, when Certkiller 2 passes BGP advertisements from Certkiller 1 about network 192.168.2.0 to Certkiller 3, what address will be listed as the next-hop address?

- A. 10.1.1.1
- B. 10.1.1.2
- C. 192.168.1.49
- D. 192.168.1.50

Answer: D

Explanation:

For a BGP router to establish a neighbor relationship with another BGP router, issue the following configuration command:


Router(config-router)#neighbor ip-address remote-as AS-number

This command serves to identify a peer router with which the local router will establish a session.

The ip-address argument is the IP address of the neighbor interface. The AS-number argument determines whether the neighbor router is an EBGP or an IBGP neighbor.

## QUESTION 511:

Exhibit:



Refer to the exhibit and the partial configuration on router Certkiller 10. EIGRP routes are redistributed into the OSPF domain as external type 2 routes. With the OSPF link cost that is shown in the exhibit, what would be the preferred path for

Certkiller 1 to reach Network Certkiller D?

- A. Certkiller 1 will prefer Link A as a fastest internal path to the ASBR.
- B. Certkiller 1 will prefer Link B as a shortest internal path to the ASBR.
- C. Certkiller 1 will prefer Link C as a path with the lowest internal cost to the ASBR.
- D. Certkiller 1 will load balance across the three links because of the equal external cost to Network Certkiller D.

Answer: C

Explanation:

The default-metricCommand


An alternative to this redistribution configuration is to use the default-metriccommand instead of including the same seed metric with each redistribute statement.

Whenever the redistributecommand is used and the metric is not specified, the router will use the default metric value as the seed metric. The default metric value can be administratively configured for each routing protocol, as shown in the second show running-configoutput. So, Certkiller 1 router takes the path of link C.

---

### QUESTION 512:

Exhibit:



Refer to the exhibit. Which two statements are true? (Choose two.)

- A. Certkiller 2 will receive only the summary address of 2.0.0.0/8.
- B. Certkiller 2 will receive two routes, the manual summary address of 2.1.0.0/16 and the automatic, summarized classful address of 2.0.0.0/8.
- C. The summary route will be sourced from the Certkiller 1 s0/0 interface in the Certkiller 1 routing table.
- D. The summary route will be sourced from the Certkiller 1 Null0 interface in the Certkiller 1 routing table.
- E. The summary route will be assigned the administrative distance of 90 in the Certkiller 1 routing table.
- F. The summary route will be assigned the administrative distance of 5 in the Certkiller 2 routing table.

Answer: A,D

Explanation:

EIGRP also enables manual configuration of a prefix to use as a summary address. Manual summary routes are configured on a per-interface basis. The interface that will propagate the route summary must first be selected and then defined with the ip summary-address eigrp command, which has the following syntax:

```
Router(config-if)#ip summary-address eigrp autonomous-system-number  
ip-address mask administrative-distance
```

D 2.1.0.0/16 is a summary, 00:00:22, Null0

Notice that the summary route is sourced from Null0, and not an actual interface

---

**QUESTION 513:**

Which set of commands will configure summary addressing for RIPv2 on a Cisco router?

- A. Certkiller 1 (config)# router rip  
Certkiller 1 (config-router)# ip summary-address rip  
Certkiller 1 (config-router)# auto-summary
- B. Certkiller 1 (config)# router rip  
Certkiller 1 (config-router)# ip summary-address rip  
Certkiller 1 (config-router)# no auto-summary
- C. Certkiller 1 (config)# router rip  
Certkiller 1 (config-router)# auto-summary  
Certkiller 1 (config)# interface s0/0 Certkiller 1 (config-if)# ip summary-address rip
- D. Certkiller 1 (config)# router rip  
Certkiller 1 (config-router)# no auto-summary  
Certkiller 1 (config)# interface s0/0  
Certkiller 1 (config-if)# ip summary-address rip

Answer: D

Explanation:


RIP also enables manual configuration of a prefix to use as a summary address. Manual summary routes are configured on a per-interface basis. The interface that will propagate the route summary must first be selected and then defined with the ip summary-address rip command, which has the following syntax:

```
Router(config-if)#ip summary-address rip ip netmask  
no auto-summary: disables the auto summary feature.
```

---

**QUESTION 514:**

Exhibit:



In reference to the network topology diagram shown in the exhibit, which statement is true?

- A. The network is using a two-layer full-mesh hierarchical design.
- B. The network is using a two-layer hub-and-spoke hierarchical design.
- C. To improve scalability, route summarization at Routers R1,R2,R3,R4,R5,R6 must be performed.
- D. To improve security, packet filtering that uses ACLs at Routers A and B must be implemented.
- E. Routers A and B are the distribution layer routers.
- F. Routers R1,R2,R3,R4,R5 and R6 are the access layer routers.

Answer: C

Explanation: Summarization is the consolidation of multiple routes into one single advertisement. This is normally done at the boundaries of Area Border Routers (ABRs). Although summarization could be configured between any two areas, it is better to summarize in the direction of the backbone.

Example:

To configure an ABR to summarize routes for a specific area before injecting them into a different area, use the following syntax:


```
Router(config-router)#area area-id range address mask
```

To configure an ASBR to summarize external routes before injecting them into the OSPF domain, use the following syntax:

```
Router(config-router)#summary-address address mask
```

### QUESTION 515:

Exhibit:



Refer to the exhibit. You have been assigned the network address 64.186.80.0 /24 from your service provider. Which three statements would most efficiently address the IP addressing requirements outlined in the exhibit? (Choose three.)

- A. The network address 64.186.80.64 /26 would be assigned to the network attached to the Fa0 interface of the Certkiller Branch router.
- B. The network address 64.186.80.128 /27 would be assigned to the network attached to the Fa0 interface of the Certkiller Branch router.
- C. The network address 64.186.80.192 /28 would be assigned to the network attached to the E0 interface of the Certkiller SOHO router.
- D. The network address 64.186.80.200 /28 would be assigned to the network attached to the E0 interface of the Certkiller SOHO router.
- E. The network address 64.186.80.0 /25 would be assigned to the network attached to the Fa0/0 interface of Certkiller HQ router.
- F. The network address 64.186.80.64 /25 would be assigned to the network attached to the Fa0/0 interface of Certkiller HQ router.

Answer: B,C,E

Explanation: Answer B, C and E are correct, In branch Network required IP addresses are 20, SOHO 10 and 100 in Head office. When you subnet the 64.186.80.0/28 you will get 30 usable addresses and can use for further extend also. While you subnet 64.186.80.0/25, you will get 128 usable hosts.

---

**QUESTION 516:**

Exhibit:

```
CertKiller1 # show ip route
...
C 10.1.3.0/24 is directly connected, Serial2
D 10.1.2.0/24 [90/10537472] via 10.1.1.2, 00:23:24, Serial1
D 10.0.0.0/8 is a summary, 00:23:20, Null0
C 10.1.1.0/24 is directly connected, Serial1
S 192.168.20.0/24 is directly connected, Ethernet0
```

Refer to the exhibit. If the configuration shown below is added to Certkiller 1, which three route entries will EIGRP advertise to neighboring routers? (Choose three.)  
router eigrp 10  
network 10.0.0.0  
eigrp stub

- A. 10.1.3.0/24
- B. 10.1.2.0/24
- C. 10.0.0.0/8
- D. 10.1.1.0/24
- E. 192.168.20.0/24


Answer: A,C,D

Explanation: Remote router can configure as an EIGRP stub router. A stub router is typically used on small capacity routers in a hub-and-spoke WAN environment. The stub router in EIGRP is similar to the concepts of on Demand Routing. As in ODR, the router in EIGRP network has no other neighbors and accesses the network through a distribute layer router. It is not necessary, therefore, for this remote router to have a complete routing table that may overwhelm its limited resources. The remote router needs only a default route to the distribute router that can serve all its needs. When router Certkiller 3 will configure as a stub router, there will be 10.1.3.0/24, 10.0.0.0/8, 10.1.1.0/24 entries.

---

**QUESTION 517:**

Exhibit:



Refer to the exhibit. The administrator of this internetwork needs to provide subnets with a /30 mask for the three WAN links connecting Certkiller 1 to Certkiller 2, Certkiller 3, and Certkiller 4. The address ranges allocated for these connections are within the 172.16.14.0/27 address range. Which three subnets can be used for the WAN links?

- A. 172.16.14.36/30
- 172.16.14.40/30
- 172.16.14.44/30
- B. 172.16.14.68/30
- 172.16.14.72/30
- 172.16.14.76/30
- C. 172.16.14.100/30


- 172.16.14.104/30
- 172.16.14.108/30
- D. 172.16.14.132/30
- 172.16.14.136/30
- 172.16.14.140/30

Answer: D

Explanation: Answer D is correct because 172.16.14.32/27 address is used in Router Certkiller 2, 172.16.14.64/27 is used on Router Certkiller 3 and 172.16.14.96/27 is used in Router Certkiller 4. 172.16.14.32 is the first subnet, 172.16.14.64/27 is second subnet, 172.16.14.96 is the third subnet. When you use the 27 bits for network address, you will get the 30 usable host address. Upto third subnet of 172.16.14.32/27, address reaches 172.16.14.128. Now you have free address after 130 for /30 mask. While you use the 30 bits network address you will get the 2 usable host address. So according to previous subnet, /30 bits subnet starts from 172.16.14.132/30, 172.16.14.136/30, 172.16.14.140/30 for Three router's WAN links.

**QUESTION 518:**

Exhibit:



Refer to the exhibit. A network administrator is troubleshooting IS-IS routing. Currently, Certkiller 1 is not receiving Link State Packets (LSPs) from Certkiller 2. What is the problem?

- A. Certkiller 2 must be configured with the is-type level-1 command under router isis.
- B. Certkiller 1 must be configured with the isis circuit-type level-1 under interface s0.
- C. Certkiller 1 and Certkiller 2 are configured in different areas.
- D. Certkiller 1 and Certkiller 2 should both be configured for is-type level-2 under router isis.

Answer: C

Explanation:

**Area ID** The area ID is typically a one- or two-octet field that precedes the system ID. The area ID is used to signify the area the router belongs to. The area ID can span up to two octets if need be. For the BSCI, we will concentrate on a one-octet area ID. For the NSAP addressing scheme used with IS-IS, the area ID is everything to the left of the system ID.

**System ID** The system ID is used to identify the router. It is similar to the router ID in OSPF. The system ID can be up to eight octets in length. Cisco supports only a six-octet field for the system ID. Whether or not the vendor you are using supports more than six octets, the number of octets used must be the same throughout the IS-IS routing domain. Normally, the system ID will be set to one of the MAC addresses on the router. This is not a rule, and the system ID can be represented in a number of ways. The system ID that is used for a router must be unique throughout the IS-IS routing domain. There must be only one system ID set per router, regardless of the number of routing processes. The system ID must be common across all routing processes on the same router.

**SEL** The NSAP selector byte (SEL) is a one-octet field that represents the service being offered at the network level of the device. For our implementation of IS-IS, the SEL will always be set to 00. 00 represents the router. A good comparison to the SEL is the TCP or UDP port number being associated with an IP address, forming a socket that identifies a unique application instance on the internetwork. Basically, it tells you what service is being referenced for the particular address.


As previously mentioned, the NET can come in and be entered in a variety of NSAP formats, the last eight bytes of each serving the same function. Let's go ahead and take a look at the three most common formats

In Between two routers, Area ID is different, one router has 1921 area id value while another router has 192 value in area id.

---

**QUESTION 519:**

Exhibit:



All routers in the diagram are configured with EIGRP. If Certkiller 2 and Certkiller 3 fail, which action will Certkiller 1 take with respect to the HQ network?

- A. Certkiller 1 will automatically route packets via RTD to the HQ network.
- B. Certkiller 1 will place the route via RTD into the hold down state.
- C. Certkiller 1 will go into the active state for all routes.

D. Certkiller 1 will go into the active state for the route to HQ network.


Answer: D

Explanation: EIGRP routing protocol selects the path to reach the destination having lowest cost among multiple path. EIFRP supports equal and unequal path selection methods. While Certkiller 3 and CertK in2 fail packets EIGRP selects the path to reach HQ network through Certkiller 4 having highest cost then Certkiller 3 or Certkiller 2.

---

**QUESTION 520:**

Exhibit:



Refer to the exhibit. During the process of configuring a virtual link to connect area 2 with the backbone area, the network administrator received this console message on Certkiller 3: \*Mar 1 00:25:01.084: %OSPF-4-ERRRCV: Received invalid packet: mismatch area ID, from backbone area must be virtual link but not found from 20.20.20.1, Serial 0 How should the virtual link be configured on the OSPF routers to establish full connectivity between the areas?

- A. Certkiller 1(config-router)# area 1 virtual-link 30.30.30.3  
Certkiller 3(config-router)# area 1 virtual-link 20.20.20.1
- B. Certkiller 1(config-router)# area 1 virtual-link 20.20.20.2  
Certkiller 3(config-router)# area 1 virtual-link 30.30.30.2
- C. Certkiller 1(config-router)# area 0 virtual-link 1.1.1.1  
Certkiller 3(config-router)# area 2 virtual-link 3.3.3.3
- D. Certkiller 1(config-router)# area 1 virtual-link 3.3.3.3  
Certkiller 3(config-router)# area 1 virtual-link 1.1.1.1
- E. Certkiller 1(config-router)# area 1 virtual-link 2.2.2.2  
Certkiller 3(config-router)# area 1 virtual-link 2.2.2.2

Answer: D

Explanation: When designing a multi-area OSPF network, all areas should be connected to the backbone area. However, there may be instances when an area will

need to cross another area to reach the backbone area, as shown in below figure. Because, in this example, Area 20 does not have a direct link to Area 0, we need to create a virtual link.

The syntax for creating a virtual link across an area is

area area-id virtual-link router-id


where area-id is the number of the transit area, in this example Area 10, and router-id is the IP address of the highest loopback interface configured on a router or can be manually set.

To manually set the router-id, you need to enter the command router-id id-in-IP-address-format in router configuration mode. If you do not manually set the router-id and a loopback interface has not been configured on the router, then the router-id is the highest IP address configured on the router.

---

**QUESTION 521:**

Exhibit:



Which type of Link State Announcement will Certkiller 5 use to announce changes in network Certkiller B to Area 1?

- A. Type 1 LSA
- B. Type 2 LSA
- C. Type 3 LSA
- D. Type 4 LSA
- E. Type 5 LSA

Answer: E


Explanation:

Type 5 LSA Referred to as AS external link advertisements, Type 5 LSAs are sent by autonomous system boundary routers. These ASBRs use Type 5 LSAs to advertise routes that are external to the OSPF autonomous system or a default route external to the OSPF autonomous system that is reachable through them.

---

**QUESTION 522:**

Exhibit:



Refer to the exhibit. Given the exhibited router output, which command sequence can be added to Certkiller 1 to generate a default route into the OSPF domain?

- A. default-router
- B. ip default-network
- C. default-information originate always
- D. ip default-gateway

Answer: C

Explanation: By default, in normal areas OSPF routers don't generate default routes into their routing domains, even if one exists To have an OSPF router generate a default route, use the default-information originate [always] [metric metric value] [metric -type type value ] [route-map map-name] command in the OSPF router configuration mode.

### QUESTION 523:

A Level-2 IS can establish an adjacency with which routers? (Choose three.)

- A. any Level-1 IS in any area
- B. any Level-2 IS in any area
- C. any Level-1 IS in the same area
- D. any Level-1/Level-2 IS in any area
- E. any Level-1/Level-2 IS in the same area

Answer: B,D,E

Explanation:

L1 routers are also referred to as station routers because they enable stations to communicate with each other and the rest of the network. A contiguous group of L1 routers defines an area. The L1 routers maintain the L1 link-state PDU database (LSPD), which defines the picture of the area itself and the exit points to neighboring areas. L2 routers are also referred to as area routers because they interconnect the L1 areas. L2 routers store a separate LSPD, which contains only the inter-area topology information.

L1L2 routers act as if they were two IS-IS routers. Physically, an L1L2 router connects to L1 routers inside its area and to L2 routers in the backbone. Notice that the boundary between areas in IS-IS exists on a link between routers and not on an ABR itself, as in OSPF. Logically, the L1L2 router acts, for the purposes of IS-IS routing, as if it were two logical routers. An L1L2 router operates an L1 routing process, with its own L1 topology table and adjacency table, to handle its association with other L1 routers and ESs. An L1L2 router also operates an L2 routing process, with a separate L2 topology table and a separate L2 adjacency table, to handle its association with its neighbor backbone routers. L1L2 routers support an L1 function to communicate with the other L1 routers in their area and maintain the L1 LSP information in an L1 LSPD. They inform other L1 routers that they are an exit point for the area. They also support an L2 function to communicate with the rest of the backbone and maintain an L2 topology database separately from their L1 LSPD. IS-IS does not share the concept of an Area 0 with OSPF. An IS-IS domain appears as a set of distinct areas interconnected by a chain of L2 routers, weaving their way through and between the Level 1 areas.

An L1 area is a collection of L1 and L1L2 routers. The backbone area is a collection of L2 and L1L2 routers and has to be contiguous. IS-IS does not have a backbone area like OSPF Area 0. The IS-IS backbone is a contiguous collection of L2 capable routers, each of which can be in a different area.

---

**QUESTION 524:**

Exhibit:

```
CertKiller1#show isis route
IS-IS Level 1 Rolling Total Session
System Id Next-Hop Interface SNPA Metric State
CertKiller1 CertKiller1 Et0/0 aabb.cc00.3300 10 Up L2-IS
CertKiller3 ---
Default route out of area - (via 1 L2-attached IS)
System Id Next-Hop Interface SNPA Metric State
CertKiller1 CertKiller1 Et0/0 aabb.cc00.3300 10 Up
```

In reference to the Certkiller 3 show isis route output, which statement is true?

- A. Certkiller 1 is the exit point out of the area for Certkiller 3.
- B. Certkiller 1 is a level-1 only IS-IS router.
- C. Certkiller 1 has been configured with a nondefault IS-IS metric.
- D. Certkiller 3 is a level-2 only IS-IS router.
- E. Certkiller 3 has been configured with a nondefault IS-IS metric.
- F. Certkiller 3 routing table should contain i L2 entries.

Answer: A

Explanation:

The show clns route and show isis route commands are the next commands to be looked at. There is a common source of confusion for those learning Integrated IS-IS. That confusion comes from the frequent use of commands directly referencing CLNS, but used for the purpose of verifying and troubleshooting IP routing.


The command show clns route shows the local NET entry. This command also shows the L2 routes to its own area and the neighbor areas. Notice that L2 regards the route to the area of R2 as being through itself. This further emphasizes that the L1 and L2 processes

operate separately. The command show isis routes shows the IS-IS neighbors. While using the show isis route command you will get the next hop router name.

---

**QUESTION 525:**

Exhibit:



```
Certkiller 1(config)# router eigrp 2000
Certkiller 1(config-router)# network 2.0.0.0
Certkiller 1(config-router)# exit
Certkiller 1(config)# interface serial 2
Certkiller 1(config-if)# ip summary-address eigrp 2000 2.1.0.0 255.255.0.0
```

Examine the exhibit.

How many routes will Certkiller 2 have with this configuration assuming there is no additional connectivity to Certkiller 2 at this time?

- A. 0
- B. 1
- C. 2
- D. 3
- E. 4

Answer: C

Explanation:

EIGRP also enables manual configuration of a prefix to use as a summary address. Manual summary routes are configured on a per-interface basis. The interface that will propagate the route summary must first be selected and then defined with the ip summary-address eigrp command, which has the following syntax:

```
Router(config-if)#ip summary-address eigrp autonomous-system-number ip-address mask administrative-distance
```


```
D 2.1.0.0/16 is a summary, 00:00:22, Null0
```

Notice that the summary route is sourced from Null0, and not an actual interface, by examining the configuration Certkiller 2 router publish two different IP addresses in it's routing table.

---

**QUESTION 526:**

Exhibit:



Refer to the exhibit. You have been assigned the network address 64.186.80.0 /24 from your service provider. Which three statements would most efficiently address the IP addressing requirements outlined in the exhibit? (Choose three.)

- A. The network address 64.186.80.64 /26 would be assigned to the network attached to the Fa0 interface of the Certkiller 1 router.
- B. The network address 64.186.80.128 /27 would be assigned to the network attached to the Fa0 interface of the Certkiller 1 router.
- C. The network address 64.186.80.192 /28 would be assigned to the network attached to the E0 interface of the Certkiller 3 router.
- D. The network address 64.186.80.200 /28 would be assigned to the network attached to the E0 interface of the Certkiller 3 router.
- E. The network address 64.186.80.0 /25 would be assigned to the network attached to the Fa0/0 interface of Certkiller 2 router.
- F. The network address 64.186.80.64 /25 would be assigned to the network attached to the Fa0/0 interface of Certkiller 2 router.

Answer: B, C, E

Explanation: We require the 20 IP Addresses for the network connected to fa0 interface of Certkiller 1 router, while we use 27 bits network address, we will get 30 usable IP address. So answer B is correct.


10 usable IP addresses are required for Certkiller 3 connected through E0 interface. While we use the 28 bits network address, we will get 14 usable IP addresses, four IP address can use for further extend.

100 IP addresses are required for the network connected to fa0/0 port of Certkiller 2 router. When 25 bits use for network, we will get 126 usable host address so 26 host address can use for further extend,

---

### QUESTION 527:

Exhibit:



Refer to the exhibit. Based upon the output of show ip bgp, which statement is true?

- A. The default route 0.0.0.0/0 is preferred through the Certkiller 2 ISP-B link.
- B. The default route 0.0.0.0/0 is preferred through the Certkiller 1 ISP-A link.
- C. The route to network 30.30.30.0/24 is preferred through the Certkiller 2 ISP-B link.
- D. The route to network 40.40.40.0/24 is preferred through the Certkiller 1 ISP-A link.

Answer: B

Explanation:

A BGP speaker can receive updates from multiple autonomous systems that describe different paths to the same destination. It must then choose the single best path for reaching that destination. Once chosen, BGP propagates the best path to its neighbors. The decision is based on the value of attributes, such as Next Hop or Local Preference, that the update contains and other configurable BGP factors.

Next Hop

AS\_Path

Atomic Aggregate

Aggregator

Local Preference

Weight

Multiple Exit Discriminator (MED)

Origin

Local Preference is a well-known discretionary attribute, type code 5. The Local Preference attribute is a degree of preference given to a route for comparison with other routes for the same destination. Higher Local Preference values are preferred. Local Preference, as indicated by the name, is local to the AS and is exchanged between IBGP peers only. Local Preference is not advertised to EBGP peers. So ISP-A has higher Local Preference Value so default router goes through ISP-A link.

## QUESTION 528:

Which three are benefits of IS-IS over OSPF? (Choose three.)

- A. supports more routers in an area
- B. does not require Hello packets to establish neighbor relationships
- C. produces fewer link state advertisements for a given network
- D. supports route tags
- E. supports network layer protocols other than IP
- F. requires fewer neighbor relationships in a broadcast multiaccess network

Answer: A,C,E

Explanation:

Similarities of IS-IS and OSPF

- \* Classless
- \* Link-state database and Dijkstra's algorithm
- \* Hello Packets to form and maintain adjacencies
- \* Use areas to form hierarchical topologies
- \* Support address summarization between areas
- \* Elect Designated Router on Multi-access network
- \* Link-state representation, aging and metrics
- \* Update, decision and flooding processes
- \* Convergence capabilities
- \* Deployed on ISP backbones

OSPF and IS-IS

- \* IS-IS does not elect a Backup Designated Router
- \* IS-IS repeats the election process whenever a new router becomes active
- \* Every time the DR changes, a new set of LSPs must be flooded
- \* IS-IS routers form adjacencies with all neighbors, not just the DR.
- \* Each IS-IS router belongs to exactly one Level 2 area
- \* IS-IS allows a more flexible approach to extending the backbone
- \* One link-state packet per IS-IS router in one area, including redistributed prefixes, compared to many OSPF LSAs.
- \* OSPF supports 50 routers per area.
- \* IS-IS supports 1000 routers per area.

---

**QUESTION 529:**

Exhibit:

```
<output omitted>
|
ipv6 unicast-routing
|
interface Ethernet0
 ip address 192.168.99.1 255.255.255.0
|
interface Tunnel1202
 no ip address
 no ip redirects
 ipv6 address 2002:COA8:6301::1/128
 tunnel source ethernet0
 tunnel mode ipv6ip 6to4
|
 ipv6 route 2002::/16 Tunnel2002
 ipv6 route ::/0 2002:COA8:2101::1
|
<output omitted>
```

Refer to the exhibit. Given the partial configuration in the exhibit, which IPv6 statement is true?

- A. The configuration is an example of an encrypted IPv6 VPN tunnel.
- B. The configuration is an example of a one to one IPv6 tunnel.
- C. The configuration is an example of a 6to4 tunnel.
- D. The configuration is an example of a 4to6 tunnel.

Answer: C

Explanation:

6to4 Tunneling allows IPv6 to be run over an automatically configured tunnel. It requires that the routers connecting the IPv6 remote sites through the IPv4 cloud need to be running dual stacks. The edge routers responsible for running the tunnel will use the prefix 2002::/16 and append the IPv4 interface address to create an address. The interface is the IPv4 address converted to hexadecimal and added to the routing prefix 2002::/16.

---

**QUESTION 530:**

To create an NSSA totally stubby area in Area 1, what commands should be configured on the NSSA ABR?

- A. router ospf 1 area 1 nssa
- B. router ospf 1 area 1 nssa no-summary
- C. router ospf 1 area 1 nssa no-redistribution
- D. router ospf 1 area 1 nssa default-information originate
- E. router ospf 1 area 1 nssa default-information originate metric-type 2

Answer: B

Explanation:

Not-so-stubby area (NSSA)- An NSSA is an area that is similar to a stub area but allows for importing external routes as Type 7 LSAs and translation of specific Type 7 LSA routes into Type 5 LSAs.

To configure an OSPF area as a NSSA, configure all OSPF router interfaces that belong to the area using the following command syntax:

```
Router(config-router)#area area-id nssa [no-summary]
```

Typically, use the optional keyword no-summary when configuring NSSA on an ABR.

This prevents Type 3 and Type 4 summary routes from flooding the NSSA


area and minimizes the routing tables within the area. In effect, the

no-summary keyword makes the NSSA totally stubby.

---

**QUESTION 531:**

Exhibit:



Refer to the exhibit. By default, when Certkiller 2 passes BGP advertisements from Certkiller 1 about network 192.168.2.0 to Certkiller 3, what address will be listed as the next-hop address?

- A. 10.1.1.1
- B. 10.1.1.2
- C. 192.168.1.49
- D. 192.168.1.50

Answer: C

Explanation:

To begin configuring a BGP process, issue the following familiar command:

```
router(config)#router bgp AS-number
```

Notice that the Cisco IOS permits only one BGP process to run at a time. Therefore, a router cannot belong to more than one BGP AS.

The network command is used with IGPs, such as RIP, to determine the interfaces on which to send and receive updates. The command also indicates which directly connected networks to advertise. However, when configuring BGP, the network command does not affect what interfaces BGP runs on. Therefore, configuring just a network statement will not establish a BGP neighbor relationship. This is a major difference between BGP and IGPs. The network statement follows this syntax:

```
Router(config-router)#network network-number [mask network-mask]
```

In BGP, the network command tells the BGP process what locally learned networks to advertise. The networks can be connected routes, static routes, or routes learned by way of a dynamic routing protocol, such as RIP. These networks must also exist in the routing table of the local router or they will not be sent out in updates. The mask keyword can be used with the network command to specify individual subnets. Routes learned by the BGP process are propagated by default but are often filtered by a routing policy.

For a BGP router to establish a neighbor relationship with another BGP router, issue the following configuration command:

```
Router(config-router)#neighbor ip-address remote-as AS-number
```

This command serves to identify a peer router with which the local router will establish a session. The ip-address argument is the IP address of the neighbor interface. The AS-number argument determines whether the neighbor router is an EBGP or an IBGP neighbor.

So, while configuring the BGP in Certkiller 1

```
#router bgp 65000  
#network 192.168.2.0  
#neighbor 192.168.1.49 remote-as 64520
```

---

**QUESTION 532:**

When an IPv6 enabled host boots, it sends a router solicitation (RS) message. An IPv6 router responds with a router advertisement (RA). Which two items are contained in the RA? (Choose two.)

- A. IPv6 address for the host
- B. lifetime of the prefix
- C. prefixes for the link
- D. keepalive timers
- E. request for the local host IP address
- F. any route advertisements it has received

Answer: B,C

---

**QUESTION 533:**

In order for two routers to become EIGRP neighbors, which two values must match? (Choose two.)

- A. hello time
- B. hold time
- C. autonomous system
- D. K values
- E. delay

Answer: C,D

Explanation: To establish the EIGRP neighbors, autonomous system should be same and K values also should be same. But hello time and hold time may differ.

---

**QUESTION 534:**

What does IS-IS use to establish and maintain neighbor relationships between IS's?

- A. IIH
- B. LSP
- C. CLNS
- D. CLNP
- E. ISH

Answer: A

Explanation: While you use the debug isis adj-packets command, you can find out what isis uses for establish and maintain neighbors.

```
router#debug isis adj-packets
ISIS-Adj: Rec L1 IIH from 0000.0c00.40af (Ethernet0), cir type 3, cir idBBBB.BBBB.BBBB.01
ISIS-Adj: Rec L2 IIH from 0000.0c00.40af (Ethernet0), cir type 3, cir idBBBB.BBBB.BBBB.01
ISIS-Adj: Rec L1 IIH from 0000.0c00.0c36 (Ethernet1), cir type 3, cir idCCCC.CCCC.CCCC.03
ISIS-Adj: Area mismatch, level 1 IIH on Ethernet1
ISIS-Adj: Sending L1 IIH on Ethernet1
ISIS-Adj: Sending L2 IIH on Ethernet1
ISIS-Adj: Rec L2 IIH from 0000.0c00.0c36 (Ethernet1), cir type 3, cir idBBBB.BBBB.BBBB.03
Debugging IS-IS adjacency packets provides you with knowledge of what Hello PDUs the router is sending and receiving. (IIH is an abbreviation for IS-IS Hello.)
```

---

**QUESTION 535:**

Which statement is true about BGP policy-based routing?

- A. If the next-hop router goes down and no alternative path is in place, policy routing will route to null 0.
- B. If the next-hop router goes down and no alternative path is in place, policy routing will default to another BGP path.
- C. If the next-hop router goes down and no alternative path is in place, policy routing will deny all traffic to that destination.
- D. If the next-hop router goes down and no alternative path is in place, policy routing will default to dynamic routing decisions.

Answer: D

Explanation:

Input and output policies generally are defined using route maps. Route maps are used

with BGP to control and modify routing information. Route maps are used to define how the routes are redistributed between routing domains.

Recall from Module 8, Route Optimization, that the route-map command is entered using the following syntax:

```
Router(config)#route-map map-tag [permit | deny] [sequence-number]
```

Notice that map-tag is a name that identifies the route map. The sequence-number indicates the position that an instance of the route map is to have in relation to other instances of the same route map. Instances are ordered sequentially, starting with the number ten by default.

---

**QUESTION 536:**

Exhibit:

- 192.168.8.0
- 192.168.9.0
- 192.168.10.0
- 192.168.11.0
- 192.168.12.0
- 192.168.13.0
- 192.168.14.0
- 192.168.15.0

Which address will successfully summarize only the following 8 class C addresses?

- A. 192.168.4.0/20
- B. 192.168.8.0/20
- C. 192.168.8.0/21
- D. 192.168.12.0/20
- E. 192.168.12.0/21
- F. 192.168.12.0/22

Answer: C

Explanation: To get the summarize address, your should convert into binary and take the common digits. See the example


| IP Network Address | Binary Equivalent           |
|--------------------|-----------------------------|
| 172.16.100.0       | 10101100.0001000.01100100.0 |
| 172.16.101.0       | 10101100.0001000.01100101.0 |
| 172.16.102.0       | 10101100.0001000.01100110.0 |
| 172.16.103.0       | 10101100.0001000.01100111.0 |
| 172.16.104.0       | 10101100.0001000.01101000.0 |
| 172.16.105.0       | 10101100.0001000.01101001.0 |
| 172.16.106.0       | 10101100.0001000.01101010.0 |

In this example summarize address is 172.16.96.0/20. Similarly you should convert all

given addresses into equivalent binary digits, you will get 192.168.8.0/21 summarize address.

**QUESTION 537:**

Exhibit:



Refer to the exhibit. You have been assigned the network address 64.186.80.0 /24 from your service provider. Which three statements would most efficiently address the IP addressing requirements outlined in the exhibit? (Choose three.)

- A. The network address 64.186.80.128 /27 could be assigned to the Fa0 interface of the Certkiller 1 router.
- B. The network address 64.186.80.128 /28 could be assigned to the Fa0 interface of the Certkiller 1 router.
- C. The network address 64.186.80.192 /28 could be assigned to the E0 interface of the Certkiller 3 router.
- D. The network address 64.186.80.192 /29 could be assigned to the E0 interface of the Certkiller 3 router.
- E. The network address 64.186.80.0 /24 could be assigned to the Fa0/0 interface of Certkiller 2 router.
- F. The network address 64.186.80.0 /25 could be assigned to the Fa0/0 interface of Certkiller 2 router.

Answer: A,C,F


Explanation: We require the 20 IP Addresses for the network connected to fa0 interface of Certkiller 1 router, while we use 27 bits network address, we will get 30 usable IP address. So answer B is correct.

10 usable IP addresses are required for Certkiller 3 connected through E0 interface. While we use the 28 bits network address, we will get 14 usable IP addresses, four IP address can use for further extend.

100 IP addresses are required for the network connected to fa0/0 port of Certkiller 2 router. When 25 bits use for network, we will get 126 usable host address so 26 host address can use for further extend.

**QUESTION 538:**

Exhibit:



Refer to the exhibit. IS-IS is configured on all routers. What two NET addresses could be used to support an IS-IS adjacency between routers Certkiller 1 and Certkiller 2?

- A. Certkiller 1 - 49.0002.1111.1111.1111.00  
Certkiller 2 - 49.0002.2222.2222.2222.00
- B. Certkiller 1 - 49.0005.1111.1111.1111.00  
Certkiller 2 - 49.0002.2222.2222.2222.00
- C. Certkiller 1 - 49.0005.1111.1111.1111.00  
Certkiller 2 - 47.0002.2222.2222.2222.00
- D. Certkiller 1 - 49.0005.1111.1111.1111.00  
Certkiller 2 - 47.0005.2222.2222.2222.00

Answer: B

Explanation:

**Area ID** The area ID is typically a one- or two-octet field that precedes the system ID. The area ID is used to signify the area the router belongs to. The area ID can span up to two octets if need be. For the BSCI, we will concentrate on a one-octet area ID. For the NSAP addressing scheme used with IS-IS, the area ID is everything to the left of the system ID.

**System ID** The system ID is used to identify the router. It is similar to the router ID in OSPF. The system ID can be up to eight octets in length. Cisco supports only a six-octet field for the system ID. Whether or not the vendor you are using supports more than six octets, the number of octets used must be the same throughout the IS-IS routing domain. Normally, the system ID will be set to one of the MAC addresses on the router. This is not a rule, and the system ID can be represented in a number of ways. The system ID that is used for a router must be unique throughout the IS-IS routing domain. There must be only one system ID set per router, regardless of the number of routing processes. The system ID must be common across all routing processes on the same router.


**SEL** The NSAP selector byte (SEL) is a one-octet field that represents the service being offered at the network level of the device. For our implementation of IS-IS, the SEL will always be set to 00. 00 represents the router. A good comparison to the SEL is the TCP or UDP port number being associated with an IP address, forming a socket that identifies a unique application instance on the internetwork. Basically, it tells you what service is being referenced for the particular address.

As previously mentioned, the NET can come in and be entered in a variety of NSAP formats, the last eight bytes of each serving the same function. Let's go ahead and take a look at the three most common formats

---

**QUESTION 539:**

Exhibit:



Refer to the exhibit. Routing information is being exchanged between the routers using the RIP protocol. Given the information in the exhibit, which statement is true?

- A. The command ip access-group 1 in has been configured on the Serial 0/0 interface on the Certkiller 2 router.
- B. The command ip access-group 1 out has been configured on the Serial 0/0 interface on the Certkiller 2 router.
- C. The command distribute-list 1 in serial 0/0 has been configured in router configuration mode on the Certkiller 2 router.
- D. The command distribute-list 1 out serial 0/0 has been configured in router configuration mode on the Certkiller 2 router.
- E. The command passive-interface s0/0 has been configured in router configuration mode on the Certkiller 2 router.

Answer: D

Explanation: According to output, distribute list 1 out is applied in serial 0/0 interface.

---

**QUESTION 540:**

Which three are benefits of IS-IS over OSPF? (Choose three.)

- A. supports more routers in an area
- B. does not require Hello packets to establish neighbor relationships

- C. produces fewer link state advertisements for a given network
- D. supports route tags
- E. supports network layer protocols other than IP
- F. requires fewer neighbor relationships in a broadcast multiaccess network

Answer: A,C,E

Explanation:

Similarities of IS-IS and OSPF

1. Classless
2. Link-state database and Dijkstra's algorithm
3. Hello Packets to form and maintain adjacencies
4. Use areas to form hierarchical topologies
5. Support address summarization between areas
6. Elect Designates Router on Multi-access network
7. Link-state representation, aging and metrics
8. Update, decision and flooding processes
9. Convergence capabilities
10. Deployed on ISP backbones

OSPF and IS-IS

1. IS-IS does not elect a Backup Designated Router
2. IS-IS repeats the election process whenever a new router become active
3. Every time the DR changes, a new set of LSPs must be flooded
4. IS-IS routers form adjacencies with all neighbors, not just the DR.
5. Each IS-IS router belongs to exactly one Level 2 area
6. IS-IS allows a more flexible approach to extending the backbone
7. One link-state packets per IS-IS router in one area, including redistributed prefixes, compared to many OSPF LSA.
8. OSPF supports 50 routers per area.
9. IS-IS supports 1000 routers per area.

---

**QUESTION 541:**

Which is the most effective technique to contain EIGRP queries?

- A. route summarization
- B. configuring route filters
- C. using a hierarchical addressing scheme
- D. establishing separate autonomous systems

Answer: A

Explanation:

EIGRP also enables manual configuration of a prefix to use as a summary address. Manual summary routes are configured on a per-interface basis. The interface that will propagate the route summary must first be selected and then defined with the ip summary-address eigrp

command, which has the following syntax:

```
Router(config-if)#ip summary-address eigrp autonomous-system-number  
ip-address mask administrative-distance
```

D 2.1.0.0/16 is a summary, 00:00:22, Null0

Notice that the summary route is sourced from Null0, and not an actual interface,

---

**QUESTION 542:**

Which two statements about BGP policy-based routing are true? (Choose two.)

- A. Policy routing is configured on the inbound interface.
- B. Policy routing can be used to alter the final destination of the packet.
- C. Policy routing can be used to alter the next hop in the path to the destination.
- D. Policy routing does not allow traffic to be directed based on the source address.

Answer: A,C

Explanation:

Input and output policies generally are defined using route maps. Route maps are used with BGP to control and modify routing information. Route maps are used to define how the routes are redistributed between routing domains.

Recall from Module 8, Route Optimization, that the route-map command is entered using the following syntax:

```
Router(config)#route-map map-tag [permit | deny] [sequence-number]
```

Notice that map-tag is a name that identifies the route map. The sequence-number indicates the position that an instance of the route map is to have in relation to other instances of the same route map. Instances are ordered sequentially, starting with the number ten by default.

---

**QUESTION 543:**

Which show command will display only the Type 5 LSAs in the OSPF topology database?

- A. show ip route
- B. show ip route ospf
- C. show ip ospf database summary
- D. show ip ospf database nssa-external
- E. show ip ospf database external

Answer: E

Explanation:

show ip ospf database - Displays the contents of the topological database maintained by the router. Several keywords can be used with this command to get specific information about the following links:

show ip ospf [process-id area-id] database [router] - Displays individual router link-state information.

show ip ospf [process-id area-id] database [network] - Displays network link-state information. The area ID is the area number associated with the OSPF address range defined in the network router configuration command when defining a particular area.

show ip ospf [process-id area-id] database [summary] - Displays summary information about ABR link states.

show ip ospf [process-id area-id] database [asbr-summary] - Displays information about ASBR link-states.

show ip ospf [process-id area-id] database [external] - Displays information about autonomous system external link states.

show ip ospf [process-id area-id] database [database-summary] - Displays database summary information and totals.

---

### QUESTION 544:

Exhibit:

CertKiller1

```
C 10.1.3.0/24 is directly connected, Serial2
D 10.1.2.0/24 [90/10537472] via 10.1.1.2, 00:23:24, Serial1
D 10.0.0.0/8 is a summary, 00:23:20, Null0
C 10.1.1.0/24 is directly connected, Serial1
S 192.168.20.0/24 is directly connected, Ethernet0
```

Refer to the exhibit. What happens when the router stops receiving advertisements for the 10.1.2.0/24 network?

- A. The summary route will be removed from the table.
- B. The summary route will remain in the table.
- C. The more specific routes will be advertised from the table.
- D. 10.1.2.0/24 will still be advertised but packets destined for it will be dropped when they reach this router.


Answer: B

Explanation: If router stops receiving advertisements also summary routing remain in routing table.

---

### QUESTION 545:

Exhibit:



Refer to the exhibit and the partial configuration on router Certkiller 2. On router Certkiller 4, all RIP routes are redistributed into the OSPF domain. A second redistribution is configured on router Certkiller 2 using a route map. Based on the configuration on Certkiller 2, which EIGRP external routes will be present in the routing table of Certkiller 1?

- A. Only routes originating in the OSPF routing domain.
- B. The routes originating from the RIP routing domain.
- C. All routes originating from RIP and OSPF routing domains.
- D. There will be no EIGRP external routes in the routing table of Certkiller 1.

Answer: A

Explanation: According to policy routing configuration, only the routes originating in the OSPF domain presents in the routing table.

### QUESTION 546:

DRAG DROP

Your boss at Certkiller , Mrs. Certkiller, is curious about BGP attributes. In which order are they used for determining a route?

| BGP Attributes, Select from these | Place      |
|-----------------------------------|------------|
| Originate route                   | Place      |
| AS_Path                           | Place seco |
| weight                            | Place thir |
| Local preference                  | Place four |
| MED                               | Place 5th  |

Answer:

Place here

|                  |
|------------------|
| weight           |
| Local preference |
| Originate route  |
| AS_Path          |
| MED              |

Explanation:

BGP bases its decision process on the attribute values. When faced with multiple routes to the same destination, BGP chooses the best route for routing traffic toward the destination. The following process summarizes how BGP chooses the best route:

1. If the next hop is inaccessible, the route is ignored. This is why it is important to have an IGP route to the next hop.
2. The BGP router will prefer the path with the largest weight. Weight is a Cisco proprietary parameter.
3. If the weights are the same, the BGP router will prefer the route with the largest local preference.
4. If the routes have the same local preference, the BGP router will prefer the route that was locally originated by this router.
5. If the local preference is the same, the BGP router will prefer the route with the shortest AS\_Path.
6. If the AS\_Path length is the same, the BGP router will prefer the route with the lowest origin type. This is when IGP is lower than EGP and EGP is lower than Incomplete.
7. If the origin type is the same, the BGP router will prefer the route with the lowest MED.
8. If the routes have the same MED, the BGP router will prefer the route in the following manner. External (EBGP) is better than confederation external, which is better than IBGP.
9. If all the preceding scenarios are identical, the BGP router will prefer the route that can

be reached by way of the closest IGP neighbor. That means taking the shortest internal path inside the AS to reach the destination by following the shortest path to the BGP Next Hop.

10. If the internal path is the same, the BGP router ID will be a tiebreaker. The BGP router will prefer the route coming from the BGP router with the lowest router ID. The router ID is usually the highest IP address on the router or the loopback address.

---

**QUESTION 547:**

Exhibit:

```
SW1#show ip eigrp topology all-links
E-IGRP Topology Table for AS(65001)/ID(172.16.10.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       - - reply status is - sin Status

P 172.16.3.128/25, 2 successors, FD is 30720, serno 9
  via 172.16.10.2 (30720/28160), FastEthernet0/1
  via 172.16.3.2 (30720/28160), FastEthernet0/3
P 152.168.0.0/24, 1 successors, FD is 156160, serno 16
  via 172.16.3.2 (156160/120256), FastEthernet0/3
  via 172.16.10.2 (158720/156160), FastEthernet0/1
P 172.16.10.0/24, 1 successors, FD is 28160, serno 1
  via Connected, FastEthernet0/1
P 172.16.0.0/30, 1 successors, FD is 20514560, serno 15
  via 172.16.1.1 (20514560/20512000), FastEthernet0/2
  via 172.16.10.2 (20517120/20514560), FastEthernet0/1
P 172.16.1.0/24, 1 successors, FD is 28160, serno 2
  via Connected, FastEthernet0/2
P 172.16.2.0/24, 1 successors, FD is 30720, serno 8
  via 172.16.10.2 (30720/28160), FastEthernet0/1
  via 172.16.3.2 (33280/30720), FastEthernet0/3
P 172.16.3.0/25, 1 successors, FD is 28160, serno 3
  via Connected, FastEthernet0/3

SW1#
```

Refer to the exhibit. Safword Incorporated is a worldwide insurance provider. The network uses EIGRP as its routing protocol throughout the corporation. The network administrator does not understand the convergence of EIGRP. Using the output of the show ip eigrp topology all-links command, answer the administrator's questions.

Which three routes will be installed for the 172.16.3.128/25 networks? (choose three)

- A. 172.16.3.128/25 [90/28 160] via 172.16.1.2, 01:26:35, FastEthernet0/2
- B. 172.16.3.128/25 [90/30720] via 172.16.3.2, 01:26:35, FastEthernet0/3
- C. 172.16.3.128/25 [90/30720] via 172.16.10.2, 01:26:35, FastEthernet0/2
- D. 172.16.2.0/24 [90/30720] via 172.16.10.2,02:10:11, FastEthernet0/1
- E. 172.16.2.0/24 [90/28160] via 172.16.10.2,02:10:11, FastEthernet0/1
- F. 172.16.2.0/24 [90/33280] via 172.16.3.2,02:10:11, FastEthernet0/3

Answer: B, C, D;

Explanation: show ip eigrp topology all-links command display all routes, not just feasible successor in the eigrp topology. So, According to exhibit, B,C and D is correct answer.

---

**QUESTION 548:**

Exhibit:

```
SW1#show ip eigrp topology all-links
IP-EIGRP Topology Table for AS(65001)/ID(172.16.10.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       - - reply status is - sin Status

P 172.16.3.128/25, 2 successors, FD is 30720, serno 9
  via 172.16.10.2 (30720/28160), FastEthernet0/1
  via 172.16.3.2 (30720/28160), FastEthernet0/3
P 192.168.0.0/24, 1 successors, FD is 156160, serno 16
  via 172.16.3.2 (156160/120256), FastEthernet0/3
  via 172.16.10.2 (158720/156160), FastEthernet0/1
P 172.16.10.0/24, 1 successors, FD is 28160, serno 1
  via Connected, FastEthernet0/1
P 172.16.0.0/30, 1 successors, FD is 20514560, serno 15
  via 172.16.1.1 (20514560/20512000), FastEthernet0/2
  via 172.16.10.2 (20517120/20514560), FastEthernet0/1
P 172.16.1.0/24, 1 successors, FD is 28160, serno 2
  via Connected, FastEthernet0/2
P 172.16.2.0/24, 1 successors, FD is 30720, serno 8
  via 172.16.10.2 (30720/28160), FastEthernet0/1
  via 172.16.3.2 (33280/30720), FastEthernet0/3
P 172.16.3.0/25, 1 successors, FD is 28160, serno 3
  via Connected, FastEthernet0/3

SW1#
```

Refer to the exhibit. Safword Incorporated is a worldwide insurance provider. The network uses EIGRP as its routing protocol throughout the corporation. The network administrator does not understand the convergence of EIGRP. Using the output of the show ip eigrp topology all-links command, answer the administrator's questions.

Which three networks does the SW1 device have feasible successor for? (Choose three)

- A. 172.16.0.0/30
- B. 172.16.1.0/24
- C. 172.16.2.0/24
- D. 172.16.3.0/25
- E. 172.16.3.128/25
- F. 192.168.0.0/24

Answer: A, C, F

Explanation: show ip eigrp topology all-links command display all routes, not just feasible successor in the eigrp topology. So, According to exhibit, B,C and D is correct answer.

Feasible successor - A feasible successor is a backup route. These routes are selected at the same time the successors are identified, but are kept in the topology table. Multiple feasible successors for a destination can be retained in the topology table. 172.16.0.0/30, 172.16.2.0/24, 192.168.0.0/24 are feasible successor.

---

### **QUESTION 549:**

Exhibit:

```
SW1#show ip eigrp topology all-links
IP-EIGRP Topology Table for AS(65001)/ID(172.16.10.1)
Codes: P - Passive, A - Active, U - Update, O - Query, R - Reply,
       r - reply status is - sin Status
P 172.16.3.128/25, 2 successors, FD is 30720, serno 9
   via 172.16.10.2 (30720/28160), FastEthernet0/1
   via 172.16.3.2 (30720/28160), FastEthernet0/3
P 192.168.0.0/24, 1 successors, FD is 156160, serno 16
   via 172.16.0.2 (156160/120256), FastEthernet0/0
   via 172.16.10.2 (158720/156160), FastEthernet0/1
P 172.16.1.0/24, 1 successors, FD is 28160, serno 1
   via Connected, FastEthernet0/1
P 172.16.0.0/30, 1 successors, FD is 20514560, serno 15
   via 172.16.1.1 (20514560/20512000), FastEthernet0/2
   via 172.16.10.2 (20517120/20514560), FastEthernet0/1
P 172.16.1.0/24, 1 successors, FD is 28160, serno 2
   via Connected, FastEthernet0/2
P 172.16.2.0/24, 1 successors, FD is 30720, serno 8
   via 172.16.10.2 (30720/28160), FastEthernet0/1
   via 172.16.3.2 (33280/30720), FastEthernet0/3
P 172.16.3.0/25, 1 successors, FD is 28160, serno 3
   via Connected, FastEthernet0/3
SW1#
```


Which three networks is the router at 172.16.10.2 directly connected to? (Choose three)

- A. 172.16.0.0/30
- B. 172.16.1.0/24
- C. 172.16.2.0/24
- D. 172.16.3.0/25
- E. 172.16.3.128/25
- F. 192.168.0.0/24

Answer: C, E, F

**QUESTION 550:**

Network topology exhibit:



You work as a network administrator at Certkiller .com. Study the network topology carefully. Based on the configuration and the default behavior of the routing protocols, which network will be present in the routing table of Certkiller 4?

- A. Network A, which is automatically redistributed Between two different EIGRP Autonomous Systems.
- B. Network B, which is redistributed using the redistribute statement under EIGRP 100.

C. Network A, which is automatically redistributed between two different EIGRP autonomous systems and network B which is redistributed using the redistribute statement.


D. Only Networks belonging to EIGRP 100 will be present in the routing table of Certkiller 4.

Answer: D

---

**QUESTION 551:**

Exhibit:



You work as a network administrator at Certkiller .com. Study the network topology carefully. Based on the configuration and the default behavior of the routing protocols, which network will be present in the routing table of Certkiller 4?


- A. Network A, Which is automatically redistributed Between two different EIGRP Autonomous Systems.
- B. Network B, Which is redistributed using the redistribute statement under EIGRP 100.
- C. Network A, Which is automatically redistributed between two different EIGRP autonomous systems and network B which is redistributed using the redistribute statement.
- D. Only Networks belonging to EIGRP 100 will be present in the routing table of R4.

Answer: B

---

**QUESTION 552:**

Exhibit:



You work as a network administrator at Certkiller .com. Study the network topology carefully. Networks 10.10.2.0/24, 192.168.13.0/24, and 192.168.14.0/24 participate in EIGRP. Network 172.16.15.0/30 does not participate in EIGRP or any routing protocol. Which of the following commands should be used on the Certkiller 1 router to ensure that the traffic to destination network 192.168.13.0/24 from network 10.10.2.0/24 uses network 192.168.14.0/24 instead of the ISDN network?

- A. ip route 192.168.13.0 255.255.255.0 172.16.15.2 120
- B. ip route 192.168.13.0 255.255.255.0 172.16.15.2
- C. ip route 192.168.13.0 255.255.255.0 172.16.15.1 100
- D. ip route 192.168.13.0 255.255.255.0 172.16.15.1 90

Answer: D

---

**QUESTION 553:**

An Administrator types in the command `router ospf 1` and receives the error message: "OSPF process 1 can't start". What should be done to correctly set up OSPF?

- A. Ensure that an interface has been configured with an IP Address,
- B. Ensure that an interface has been configured with an IP Address and is up.
- C. Ensure that IP Classless is enabled.
- D. Ensure that the interface can ping their directly connected neighbors.

Answer: A

---

**QUESTION 554:**

You are the network administrator at Certkiller . You have configured multiple IP routing protocols on a single router on the Certkiller network. Which command lists the filters applied to routing updates on a routing protocol basis?

- A. show ip
- B. show ip route
- C. show ip protocols
- D. show ip interface

Answer: C

Explanation:

The show ip protocols command will display the IP routing protocols configured on the router and will also show what each routing process is redistributing. In addition, it will list the redistribution filters applied to interfaces. Specifically, the output will show:

Routing protocol and process ID  
Update frequency  
Hold down timers  
Incoming and outgoing filters  
Default distribution metric  
Redistribution parameters  
Chapter: 1

---

**QUESTION 555:**

Using route summarization, which two of these networks fall into the 174.69.16.0/20 range?

- A. 174.69.33.0/24
- B. 174.69.31.0/24
- C. 174.69.17.0/24
- D. 174.69.32.0/24

Answer: B, C

Explanation:

The valid 24-bit subnets from the address 174.69.16.0/20 are:

174.69.16.0/24 174.69.24.0/24  
174.69.17.0/24 174.69.25.0/24  
174.69.18.0/24 174.69.26.0/24  
174.69.19.0/24 174.69.27.0/24  
174.69.20.0/24 174.69.28.0/24  
174.69.21.0/24 174.69.29.0/24  
174.69.22.0/24 174.69.30.0/24  
174.69.23.0/24 174.69.31.0/24

Chapter: 1

---

**QUESTION 556:**

Which of the following statement is true when a static route is configured on a

router and that static route is advertised throughout the network?

- A. The router automatically advertises static routes to all routers
- B. You should configure redistribution using the redistribute static command
- C. You should enable static advertisements using the advertise static route command
- D. You should include the static route in a distribution list and specify which interface to use when redistributing the route.

Answer: B

Explanation:

To redistribute static routes that have been created on the local router to other routers in the network, use the redistribute static command.

Chapter: 1

---

**QUESTION 557:**

You are the network administrator at Certkiller . You are configuring redistribution to advertise OSPF routes into EIGRP on a boundary router on the Certkiller network. You specify a seed metric with the default-metric command. What is the format of the metric being specified?

- A. hop-count
- B. hop-count ticks
- C. bandwidth delay hop-count load
- D. load delay hop-count reliability mtu
- E. bandwidth delay reliability load mtu

Answer: E

Explanation:

When redistributing static routes or other protocols within EIGRP, metrics can be set for these routes using the default-metric command. The range of values for each parameter is listed below:

bandwidth - 0 to 4,294,967,295 in Kbps

delay - 0 to 4,294,967,295 in 10-microsecond units

reliability - 0 to 255 with 255 being the most reliable

load - 0 to 255 with 255 being a saturated link

MTU - 0 to 4,294,967,295

Chapter: 1

---

**QUESTION 558:**

Which switching mode is enabled by default on a router running Cisco IOS 11.2 or later, to forward packets that match the established policy routing?

- A. fast
- B. wire-speed
- C. NetFast
- D. packet

Answer: A

Explanation:

Fast switching is the default switching mechanism on all Cisco router platforms. It is accomplished by maintaining a cache of recently switched destinations therefore reducing the number of full route table lookups. It also allows the information required for MAC header rewrites to be stored in cache rather than being recalculated.

Chapter: 1

---

**QUESTION 559:**

You are the network administrator at Certkiller . Router CK1 is configured as follows:

```
router igrp 300
network 192.168.20.0
network 192.168.24.0
network 192.168.27.0
redistribute rip
default-metric 10 100 255 1 1500
distance 140 0.0.0.0 255.255.255.255 9
access-list 9 permit 192.168.20.0
access-list 9 permit 192.168.24.0
access-list 9 permit 192.168.27.0
```

Which of the following statements are true? (Choose all that apply.)

- A. Networks 192.168.20.0, 192.168.24.0, and 192.168.27.0 are allowed into the routing table
- B. The RIP learned routes to networks 192.168.20.0, 192.168.24.0, and 192.168.27.0, will be assigned an administrative distance of 140
- C. The IGRP learned routes to networks 192.168.20.0, 192.168.24.0, and 192.168.27.0, will be assigned an administrative distance of 140
- D. Changing the administrative distance to a number larger than the default value makes networks 192.168.20.0, 192.168.24.0, and 192.168.27.0 unreachable

Answer: A, B

Explanation:

The networks listed in under the IGRP section are advertised to other routers on the network and installed into the routing table. The routes that are redistributed from RIP from those same networks are assigned an administrative distance of 140 because of the distance command listed above. The distance command is used to define an

administrative distance for routes learned from other routing protocols. The last argument (9) specifies that access-list 9 be used to permit/deny networks.

Chapter: 1

---

**QUESTION 560:**

Which command could you use to verify proper operation of multiple routing protocols that are sharing routes?

- A. ping
- B. show ip route
- C. show cdp neighbor
- D. show ip ospf neighbor

Answer: B

Explanation:

This is the only command listed that can show any information regarding the state of routes or routing protocols. Answer-A will not show any information regarding route selection or the route to a target address. Ping can be used to verify connectivity to another IP address.

Chapter: 1

---

**QUESTION 561:**

Which Cisco IOS command can be used to display the route maps configured on an interface?

- A. show interface
- B. show route-map
- C. show ip policy
- D. show ip route map

Answer: C

Explanation:

"show ip policy" will show you which route maps are configured on which interfaces, which is what the question seems to be asking.

Not B: WE believe this one to be wrong as the output from "show route-map" will indeed show you the route maps configured as a list or individually.

---

**QUESTION 562:**

What happens due to the implicit deny at the end of a route-map?

- A. Packets that reach the end of the route map are discarded

- B. Packets are forwarded to the null interface for special handling
- C. Packets that reach the end of the route map are routed in normal fashion
- D. Packets that reach the end of the route map are returned to the originating interface

Answer: C

Explanation:

The use of route maps for policy-based routing is a little different than other application of route maps. When used for policy-based routing, if a packet does not match the criteria specified in the route map or a matched route map statement specifies deny, then the packet is not dropped. It is sent to the routing process and routed normally, by destination, as if it had never encountered a route map. If your intention is to drop packets that do not match the criteria, it is necessary to use the set command to route packets to the null interface as the last entry in the route map.

Source: Self-Study CCNP BSCI Exam Certification Guide Third Edition P.674

Topic: Understanding Policy-Based Routing

---

### **QUESTION 563:**

You are the network administrator at Certkiller . You want to redistribute and advertise EIGRP routes into OSPF on a boundary router. The router has the following configuration:

```
router ospf 1
redistribute eigrp 1 metric 25 subnets
```

What does the 25 parameter in the redistribute command specify?

- A. It specifies the seed cost to be applied to the redistributed routes
- B. It specifies the administrative distance on the redistributed routes
- C. It specifies the metric limit of 25 subnets in each OSPF route advertisement
- D. It specifies the process-id for the pseudo process that injects the EIGRP routes into OSPF

Answer: A

Explanation:

The metric {value} command specifies the seed metric for use in redistributed routes.

Reference: Building Scalable Cisco Networks (Ciscopress) page 456

---

### **QUESTION 564:**

You are the network administrator at Certkiller . A router on the Certkiller network has one serial interface and one Ethernet interface. Given the serial interface to a WAN configuration:

```
interface serial 0.122 point-to-point
ip address 192.168.1.2 255.255.255.0
encapsulation frame-relay
```

frame-relay interface-dlci 122

Which command prevents routing protocol information from being sent on the Ethernet interface?

- A. interface serial 0.122 point-to-point  
passive-interface ethernet 0
- B. interface Ethernet 0  
ip address 192.168.12.1 255.255.255.0  
passive interface
- C. router ospf 102  
area 1 ospf  
network 192.168.1.0 0.0.0.255 area 0  
network 192.168.12.0 0.0.0.255 area 1
- D. router ospf 102  
passive-interface Ethernet 0  
network 192.168.1.0 0.0.0.255 area 0  
network 192.168.12.0 0.0.0.255 area 1

Answer: D

Explanation:

When a passive interface is defined for any routing process, then updates are not sent on the specified interface by that routing process. Passive interfaces must be defined for each routing protocol (process). The passive interface command is not a valid interface configuration command.

Chapter: 1

---

**QUESTION 565:**

Which of the following commands would produce output that can be used to verify route redistribution? (Choose all that apply.)

- A. debug
- B. traceroute
- C. show tech-support
- D. show ip route

Answer: A, B, D

Explanation:

Debug can be used to view routing protocol information exchanged between routers. Traceroute can be used to determine the path an IP packet will take when traversing the network. The show ip route command will display all known routes and indicate the source of the route (Static, OSPF, RIP, etc)

Chapter: 1

**QUESTION 566:**

Which command forces manually entered route entries are injected into the routing process?

- A. inject static
- B. inject permanent
- C. redistribute all
- D. redistribute static

Answer: D

Explanation:

The redistribute static command is used to inject static routes into the routing protocol's route table and subsequent updates.

Chapter: 1

---

**QUESTION 567:**

You are the network administrator at Certkiller . The Certkiller network includes Router CK1 . Router CK1 is configured as follows:

```
interface serial 0
ip address 10.1.1.1 255.255.255.0
encapsulation frame-relay
ip ospf network point-to-multipoint
router ospf 7
network 10.1.1.0 0.0.0.255 area 0
```

Which of the following statements are true? (Choose all that apply.)

- A. DR/BDR elections do not take place
- B. Neighbor statements are required
- C. Communication between neighbors is broadcast to 255.255.255.255
- D. The area 0 NBMA cloud is configured as more than one subnet

Answer: A, B

Explanation:

When configuring OSPF in a point-to-multipoint environment, DR/BDR elections do not take place. Neighbor statements must be statically defined due to the NBMA architecture. The point-to-multipoint environment removes the assumption that there is a full mesh and communication between neighbors is done via unicast.

Chapter: 1

---

**QUESTION 568:**

When configuring a router to participate in an OSPF area, what is the default

priority used in DR/BDR elections?

- A. 0
- B. 1
- C. 16
- D. 255

Answer: B

Explanation:

The ip ospf priority command can be used to administer which router becomes the DR. This number ranges from 0-255 and defaults to 1. A router configured with a priority of 0 can never be elected DR.

Chapter: 1

---

**QUESTION 569:**

Which two are benefits of using OSPF over RIP as a routing protocol in a large network?

(Choose all that apply.)

- A. OSPF has fewer tables to manage
- B. OSPF is a simpler protocol than RIP
- C. OSPF has virtually no reachability limits
- D. OSPF uses a metric that is based on bandwidth to select a path through a network

Answer: C, D

Explanation:

OSPF has virtually no limits with regard to scalability in large networks because of its hierarchal design. RIP uses a hop count limit (15 hops) to prevent routing loops. It is possible in a very large network to outgrow this limitation with RIP. RIP also uses hop count as its metric for selecting the best route. Cisco's implantation of OSPF calculates link cost based on bandwidth ( $10^8 / \text{Interface Bandwidth}$ ) to determine path selection across the network.

Chapter: 1

---

**QUESTION 570:**

You are a technician at Certkiller . You want to assign an OSPF router ID of 172.16.20.127.

Which series of commands should you use?

- A. ospf loopback 0  
ip address 172.16.20.127 255.255.255.0
- B. router loopback 0

```
ip address 172.16.20.127 255.255.255.0
C. interface loopback 0
ip address 172.16.20.127 255.255.255.0
D. ospf interface loopback 0
ip address 172.16.20.127 255.255.255.0
```

Answer: C

Explanation:

In Cisco's OSPF implementation, the Loopback interface address is used as the Router ID. If the Loopback interface is not configured with an IP address, the highest IP address configured on any router interface is used.

Chapter: 1

---

### **QUESTION 571:**

What is used to determine which router that will become the DR in an OSPF network?

- A. the lowest router ID
- B. the highest priority value
- C. the first router to attach to the network
- D. a router that is connected to more than one OSPF area and designated ASBR

Answer: B

Explanation:

The router with the highest priority value is elected as the DR. The second highest priority value becomes the BDR. The ip ospf priority command can be used to administer which router becomes the DR. This number ranges from 0-255 and defaults to 1. A router configured with a priority of 0 can never be elected DR.

Chapter: 1

---

### **QUESTION 572:**

Before an running OSPF can route traffic to another OSPF neighbor, what state must the router be in to route traffic?

- A. full state
- B. INIT state
- C. 2wy
- D. forwarding state

Answer: A

Explanation:

In order to route traffic to an OSPF neighbor router, the adjacency must be established before any traffic can be passed. The adjacency is not established until DR/BDR elections are completed and link-state information is exchanged (full routing information).

Chapter: 1

---

**QUESTION 573:**

In OSPF, what is defined using the network command? (Choose all that apply.)

- A. the OSPF area ID
- B. the OSPF router ID
- C. the OSPF process ID
- D. which interface is in which OSPF area

Answer: A, D

Explanation:

In OSPF, the network area command defines the interfaces on which OSPF runs and the area ID for those interfaces. The syntax for the command is below: network [address] [wildcard-mask] area [area-id]

Chapter: 1

---

**QUESTION 574:**

How does OSPF simulate a broadcast environment in an NBMA point-to-multipoint configuration for routed traffic?

- A. by creating adjacencies with each endpoint
- B. by sending replicated traffic to each neighbor
- C. by using the 224.0.0.5 multicast address on serial links
- D. by separating out each endpoint using the hello protocol

Answer: A

Explanation:

In a point-to-multipoint configuration, neighbors must be statically defined and communication are done via unicast instead of multicast.

Chapter: 1

---

**QUESTION 575:**

Which command can be used to verify when out-of-date routes will be removed from the topological database?

- A. show ip ospf

- B. show ip route
- C. show ip ospf interface
- D. show ip ospf topo-database

Answer: A

Explanation:

Displays the Link State Update Interval and the Link State Age Interval and when an update is due. Each route is flooded throughout the area via an LS

A. Each LSA has an age field that is incremented while it is contained in the database or as it gets flooded throughout the area. When an LSA reaches a Maxage it gets flushed from the database if that LSA is not on any neighbors retransmission list.

Router# show ip ospf

Routing Process "ospf 201" with ID 192.42.110.200

Supports only single TOS(TOS0) route

It is an area border and autonomous system boundary router

Summary Link update interval is 0:30:00 and the update due in 0:16:26

External Link update interval is 0:30:00 and the update due in 0:16:27

Redistributing External Routes from,

igrp 200 with metric mapped to 2, includes subnets in redistribution

rip with metric mapped to 2

igrp 2 with metric mapped to 100

igrp 32 with metric mapped to 1

Number of areas in this router is 3

Area 192.42.110.0

Number of interfaces in this area is 1

Area has simple password authentication

SPF algorithm executed 6 times

Area ranges are

Link State Update Interval is 0:30:00 and due in 0:16:55

Link State Age Interval is 0:20:00 and due in 0:06:55

Chapter: 1

---

**QUESTION 576:**

At a minimum, which two configuration commands are required to configure OSPF on a single internal router? (Choose all that apply.)

- A. network
- B. neighbor
- C. router ospf dr 1
- D. router ospf

Answer: A, D

Explanation:

OSPF is enabled on a router by specifying an OSPF process ID and defining the network, interfaces used, and area-id that will be included in the OSPF process. The network command defines the interfaces on which OSPF runs and the area ID for those interfaces.

Chapter: 1

---

**QUESTION 577:**

You are a trainee technician at Certkiller . Your instructor shows you the following router configuration:

```
interface serial 0
ip address 172.14.12.1 255.255.255.224
encapsulation frame-relay
ip ospf network non-broadcast
!
router ospf
network 172.14.12.0 31.255.255.255
neighbor 172.14.12.2
neighbor 172.14.12.3
```

Your instructor wants to know which of the following statements are true. What would your reply be? (Choose all that apply.)

- A. DR/BDR elections are not held
- B. This is a point-to-multipoint configuration
- C. The network type is non-broadcast multi-access (NBMA)
- D. The DR and BDR require a static list of neighbors

Answer: C, D

Explanation:

Because the network type is defined as non-broadcast, DR/BDR election take place based on statically defined neighbors. In addition, communication between neighbors is done via unicast instead of multicast. If this were a point-to-multipoint configuration, the ip ospf network point-to-multipoint command would have been used on the serial interface.

Chapter: 1

---

**QUESTION 578:**

You are a technician at Certkiller . Certkiller has an OSPF network. Your newly appointed Certkiller trainee wants to know what is used to send link-state information to all other routers within an OSPF area.

What would your reply be?

- A. LSA - router link, type 1
- B. LSA - network link, type2
- C. LSA - external link, type 5

D. NSA - network summary link, type 3

Answer: B

Explanation:

Network Link Advertisements are sent during the adjacency process to inform the neighbor of its network links. When a link changes state or a new link added on an existing router, the router that owns the link generates a new LSA.

Chapter: 1

---

**QUESTION 579:**

Which of the following features require the subnet mask to be carried within OSPF routing protocol updates? (Choose all that apply.)

- A. VLSM
- B. NBMA
- C. summarization
- D. SPF route calculation

Answer: A, C

Explanation:

OSPF like all classless routing protocols, carry the network number and mask in its updates. This is required to provide support for VLSM, route summarization, and super netting.

Chapter: 1

---

**QUESTION 580:**

If an OSPF router has interfaces connected in two or more areas, what kind of router is it considered?

- A. ABR
- B. ASBR
- C. MAR
- D. backbone router

Answer: A

Explanation:

An ABR is a router that has multiple interfaces with at least two interfaces in two different OSPF areas. An ASBR is a router with at least one interface connected to an external network or AS.

Chapter: 1

---

**QUESTION 581:**

Which is true of an OSPF area with too many routers?

- A. The second BDR cannot keep all the LSA information up to date
- B. Convergence time can be slower
- C. A second backbone area must be created to split the traffic into two areas.
- D. Route processing time is decreased because the information is dispersed among all routers in the area

Answer: B

Explanation:

With too many routers will take longer to converge. Answer should be B. Also not possible to have two backbone areas in ospf. Only one is allowed.

---

**QUESTION 582:**

You are a technician at Certkiller . Your newly appointed Certkiller trainee wants to know what is used to connect a new OSPF area to area 0  
What would your reply be?

- A. external router
- B. DR
- C. BDR
- D. backbone router

Answer: D

Explanation:

Any router that will have any interface connected to an OSPF backbone area is considered a backbone router.

Chapter: 1

---

**QUESTION 583:**

You are the network administrator at Certkiller . You are using OSPF as your IGP throughout the Certkiller network. You want to connect the network to the outside world or to a different routing protocol.  
Which kind of OSPF router must you configure?

- A. ABR
- B. BDR
- C. ASBR
- D. neighbor border router
- E. backbone router

Answer: C

Explanation:

ASBR's are used to connect two separate autonomous systems together. The role of the ASBR is to exchange routing information between the two routing processes.

Chapter: 1

---

**QUESTION 584:**

In order to summarize routes, which configuration requirement does an ASBR have that an ABR does not?

- A. area range command
- B. ospf summarize command
- C. aggregate-route command
- D. summary-address command

Answer: D

Explanation:

The router subordinate command summary-address is used on ASBR's to consolidate external routes. It can be used in combination with the stub areas or used stand-alone

Chapter: 1

---

**QUESTION 585:**

Which command should you use to verify which process is responsible for routing which network?

- A. show ospf
- B. show ip route
- C. show ip protocols
- D. show ip ospf database

Answer: D

Explanation:

Sample output is shown below. Note the router ID and process ID in the first line. Each link ID is

representative of a network route.

```
router#show ip ospf database
```

```
OSPF Router with id(190.20.239.66) (Process ID 300)
```

```
Displaying Router Link States(Area 0.0.0.0)
```

```
Link ID ADV Router Age Seq# Checksum Link count
```

```
155.187.21.6 155.187.21.6 1731 0x80002CFB 0x69BC 8
```

155.187.21.5 155.187.21.5 1112 0x800009D2 0xA2B8 5  
155.187.1.2 155.187.1.2 1662 0x80000A98 0x4CB6 9  
155.187.1.1 155.187.1.1 1115 0x800009B6 0x5F2C 1  
155.187.1.5 155.187.1.5 1691 0x80002BC 0x2A1A 5  
155.187.65.6 155.187.65.6 1395 0x80001947 0xEEE1 4  
155.187.241.5 155.187.241.5 1161 0x8000007C 0x7C70 1  
155.187.27.6 155.187.27.6 1723 0x80000548 0x8641 4  
155.187.70.6 155.187.70.6 1485 0x80000B97 0xEB84 6

Chapter: 1

---

**QUESTION 586:**

You are the network administrator at Certkiller . Certkiller has an address range of 172.16.20.192 to 172.16.20.223. You want to configure the area 3 border router for network summarization

Which configuration command must you use? Select two.

- A. summarize 172.16.20.192 0.0.0.31 area 3
- B. area 3 range 172.16.20.192 172.16.20.223
- B. area 3 range 172.16.20.192 255.255.255.224
- C. D. network 172.16.20.192 255.255.255.224 area 3

Answer: A, C

Explanation:

The area range command is used on ABR's to summarize and advertise routes. A network statement must also be defined, but answer D does not use a wildcard mask and therefore is an invalid command.

Chapter: 1

---

**QUESTION 587:**

You are a technician at Certkiller . Certkiller has an OSPF network. Your newly appointed Certkiller trainee wants to know why VLSM is used in an OSPF network.

What would your reply be? (Choose all that apply.)

- A. to allow for address summarization
- B. to allow use of the all zero's subnet
- C. to make efficient use of available addresses
- D. it is required for a point-to-multipoint nonbroadcast network

Answer: A, C

Explanation:

VLSM is utilized in OSPF for address summarization and it also allows for more

efficient use of networks due to support for classless boundaries.

Chapter: 1

---

**QUESTION 588:**

What does an ABR connect in an OSPF network?

- A. multiple OSPF areas
- B. OSPF and RIP networks
- C. multiple designated routers
- D. multiple autonomous systems

Answer: A

Explanation:

An ABR is a router that has multiple interfaces with at least two interfaces in two different OSPF areas. It is used to connect different OSPF areas.

Chapter: 1

---

**QUESTION 589:**

Which OSPF router is responsible for flooding an of OSPF area with type 2 link LSAs?

- A. DR
- B. ABR
- C. BDR
- D. ASBR

Answer: A

Explanation:

The DR is responsible for flooding the network with network LSA's when a change occurs within the OSPF area.

Chapter: 1

---

**QUESTION 590:**

Which statement regarding route summarization within OSPF is true?

- A. Summarization must be performed by every router within an OSPF area.
- B. Summarization prevents type 1 link LSAs from being propagated into the backbone area 0.
- C. Route summarization can be performed at any point in the network where enough contiguous addresses are present
- D. Route summarization reduces the amount of bandwidth, CPU, and memory resources

consumed by the OSPF process.

Answer: D

Explanation:

Route summarization reduces the amount of bandwidth, CPU, and memory resources consumed by the OSPF process. Each additional network requires an individual entry in the routing table and must be propagated throughout the network and added to each router's table.

Chapter: 1

---

**QUESTION 591:**

Which of the following Cisco IOS commands that can be used to view neighbor adjacencies?

(Choose all that apply.)

- A. show ip ospf database
- B. show ip ospf neighbor
- C. show ip ospf protocols
- D. show ip ospf interface ethernet 0

Answer: B, D

Explanation:

Both of these commands display neighbor adjacency information.

```
Router# show ip ospf interface ethernet 0
```

```
Ethernet 0 is up, line protocol is up
```

```
Internet Address 131.119.254.202, Mask 255.255.255.0, Area 0.0.0.0
```

```
AS 201, Router ID 192.77.99.1, Network Type BROADCAST, Cost: 10
```

```
Transmit Delay is 1 sec, State OTHER, Priority 1
```

```
Designated Router id 131.119.254.10, Interface address 131.119.254.10
```

```
Backup Designated router id 131.119.254.28, Interface addr 131.119.254.28
```

```
Timer intervals configured, Hello 10, Dead 60, Wait 40, Retransmit 5
```

```
Hello due in 0:00:05
```

```
Neighbor Count is 8, Adjacent neighbor count is 2
```

```
Adjacent with neighbor 131.119.254.28 (Backup Designated Router)
```

```
Adjacent with neighbor 131.119.254.10 (Designated Router)
```

```
Router# show ip ospf neighbor
```

```
ID Pri State Dead Time Address Interface
```

```
199.199.199.137 1 FULL/DR 0:00:31 160.89.80.37 Ethernet0
```

```
192.31.48.1 1 FULL/DROTHER 0:00:33 192.31.48.1 Fddi0
```

```
192.31.48.200 1 FULL/DROTHER 0:00:33 192.31.48.200 Fddi0
```

```
199.199.199.137 5 FULL/DR 0:00:33 192.31.48.189 Fddi0
```

Chapter: 1

---

**QUESTION 592:**

What term is used to describe a BGP autonomous system connected to two different BGP autonomous systems for increased reliability?

- A. multi-exit
- B. multisource
- C. multihomed
- D. multi-neighbor

Answer: C

Explanation:

Multihoming refers to a single network or AS as having more than one connection to another network or AS to improve reliability and/or performance.

Chapter: 1

---

**QUESTION 593:**

You are the network administrator at Certkiller . Certkiller has subscribed to multiple ISPs. You use BGP to connect to multiple ISPs. You want to force outbound Internet traffic to one ISP unless there is a link failure.

Which tool would you use?

- A. configure weight
- B. enable route reflector
- C. create a distribute list
- D. enable the Longer Autonomous System path option

Answer: A

Explanation:

The weight attribute is a Cisco proprietary attribute used by BGP in path selection. This allows as administrator to "prefer" one path over one or more paths to the same destination. In the event the preferred route or link fails, the secondary route will automatically be used by outbound traffic.

Chapter: 1

---

**QUESTION 594:**

You are a technician at Certkiller . Your newly appointed Certkiller trainee wants to know which command displays both the configured iBGP and eBGP neighbors.

What would your reply be?

- A. show bgp neighbors

- B. show ip bgp paths
- C. show ip bgp peers
- D. show ip bgp summary

Answer: D

Explanation:

This is the only valid command that is listed that will show any information about BGP neighbors. The show ip bgp neighbors will show detailed information about each neighbor but the syntax listed is incorrect. The show ip bgp summary will show the status of all configured BGP connections.

Chapter: 1

---

**QUESTION 595:**

Which method makes it possible to receive BGP routes from multiple ISPs?

- A. accept only IGP routes from the ISPs
- B. accept an external route from the ISPs
- C. accept only default routes from the ISPs
- D. accept only redistributed routes from the ISPs

Answer: B

Explanation:

This is a major difference between iBGP and eBGP. A BGP router will never forward a path learned from an iBGP peer to another iBGP peer. However, eBGP peers always forward routes learned from one eBGP peer to both eBGP and iBGP peers. Thus receiving routes from 2 upstream ISP's (eBGP peers) is allowed.

Chapter: 1

---

**QUESTION 596:**

With regard to BGP updates, which of the following statements is true?

- A. A BGP router will forward a learned path from an iBGP peer to another iBGP peer.
- B. A eBGP peer will never forward a learned path to an iBGP peer
- C. BGP Route reflectors will propagate a route learned from an iBGP peer to eBGP peers
- D. If a BGP route was learned via an update from an iBGP peer, it will propagate this information to iBGP and eBGP peers.
- E. If a BGP route was learned via an update from an eBGP peer, it will propagate this information to iBGP and eBGP peers.

Answer: E

Explanation:

This is a major difference between iBGP and eBGP. A BGP router will never forward a path learned from an iBGP peer to another iBGP peer. However, eBGP peers always forward routes learned from one eBGP peer to both eBGP and iBGP peers. Thus receiving routes from 2 upstream ISP's (eBGP peers) is allowed.

Chapter: 1

---

**QUESTION 597:**

You are a technician at Certkiller . Your newly appointed Certkiller trainee wants to know why it is necessary to redistribute or advertise IGP routes into BGP. What would your reply be?

- A. so BGP can propagate this information to other IGP neighbors
- B. so BGP can propagate this information to other iBGP neighbors
- C. so BGP can propagate this information to other OSPF neighbors
- D. so BGP can propagate this information to other eBGP neighbors

Answer: B

Explanation:

iBGP will advertise these redistributed routes to all configured iBGP neighbors on the network. This will provide the entire AS with the routing information required to reach networks that are routed with different routing protocols.

Chapter: 1

---

**QUESTION 598:**

You are a technician at Certkiller . Your newly appointed Certkiller trainee wants to know why an iBGP router must be peered with all iBGP routers within an AS. What would your reply be? (Choose two.)

- A. iBGP routes are not propagated to other eBGP peers
- B. iBGP routes that a router originates are propagated to other iBGP peers
- C. iBGP routes are propagated to other iBGP speakers in the AS that are not peers
- D. iBGP routes that are learned from an eBGP neighbor are propagated to only eBGP peers

Answer: B, D

Explanation:

A BGP router will never forward a path learned from an iBGP peer to another iBGP peer. So, in order for all routers to know about the routes originated on any router in the network, they must maintain the full mesh because another iBGP peer will never send the update to its iBGP neighbors.

Chapter: 1

---

**QUESTION 599:**

You are the network administrator at Certkiller . You have limited router memory. Under these conditions, what does Cisco suggest as the best way to connect to multiple ISPs using BGP?

- A. receive only default routes
- B. receive only external BGP routes
- C. receive only internal BGP routes
- D. receive only redistributed routes

Answer: A

Explanation:

When working with limited router resources, use default routes instead of BGP routes from the ISP. Your internal AS decides which ISP to use and sends the traffic to the appropriate ISP.

Chapter: 1

---

**QUESTION 600:**

You are a trainee technician at Certkiller . Your instructor shows you the following router configuration:

```
router bgp 6500
redistribute static
ip route 164.20.0.0 255.255.0.0 null 0
```

Your instructor wants to know which of the following statements are true. What would your reply be?

- A. It allows BGP to advertise the 164.20.0.0/16 network
- B. It results in all traffic for all subnets of 172.16.0.0 being dropped at this router
- C. Cisco prefers that you use the aggregate-address command to distribute IGP routes into BGP
- D. Cisco prefers this method of distributing IGP routes into BGP over using the network command

Answer: C

Explanation:

Redistribution of static routes configured to the null 0 interface into BGP is done to advertise aggregate routes rather than specific routes from the IP table. However, Cisco recommends the use of the aggregate-address-command instead.

---

**QUESTION 601:**

You are a trainee technician at Certkiller . Your instructor shows you the following

partial information from the output of a BGP command on Router CK1 :

```
Network Next Hop Metric LocPrf Weight Path
192.168.2.0 10.15.10.2 0 100 65250 65000 i
10.15.20.2 0 120 65200 65000 i
10.15.30.2 0 130 65000 i
10.15.40.2 0 140 65000 i
```

Your instructor wants to know which next-hop address Router CK1 uses to send data destined for the network 192.168.2.0.

What would your reply be? (Choose all that apply.)

- A. 10.15.30.2
- B. 10.15.40.2
- C. 10.15.10.2
- D. 10.15.20.2

Answer: B

Explanation:

The next hop router 10.15.40.2 has the highest weight value. The weight attribute is Cisco proprietary and when multiple paths to the same destination exist, the connection with the highest weight value is used. This allows an administrator to prefer one path over one or more others to the same destination.

Chapter: 1

---

### **QUESTION 602:**

What is the correct command to create a BGP prefix list that will permit all prefixes between /8 and /24 for the 10.0.0.0 network?

- A. ip prefix-list 10.0.0.0/8 ge 8 le 24
- B. ip prefix-list 10.0.0.0/8 ge 24 le 8
- C. ip prefix-list 10.0.0.0/24 ge 24 le 8
- D. ip prefix-list 10.0.0.0/24 ge 8 le 24

Answer: A

Explanation:

The prefix list optional syntaxes ge-value and le-value are used when you need to specify a range of the prefix that is more specific than identified in the network/len syntax. Use the following rule when specifying these values:

len < ge-value < le-value <=32.

Chapter: 1

---

### **QUESTION 603:**

You are a technician at Certkiller . Your newly appointed Certkiller trainee wants to

know what can be used to advertise iBGP learned routes to other iBGP neighbors within the AS.

What would your reply be?

- A. client router
- B. EBGP peer
- C. route reflector
- D. community router

Answer: C

Explanation:

Route reflectors are the only BGP routers that can propagate iBGP routes to other iBGP peers. By configuring route reflectors, you reduce the number of neighbor peering relationships in an AS. This creates a central source for updates to the route reflector clients and eliminates the need for a fully meshed iBGP network.

Chapter: 1

---

#### **QUESTION 604:**

You are the network administrator at Certkiller . The Certkiller network has three configured BGP route reflectors. Each route reflector has a minimum of 2 clients. Your newly appointed Certkiller trainee wants to know what action a BGP route reflector takes if it receives updates from a peer in another autonomous system. What would your reply be?

- A. It discards the route
- B. It sends the update to all iBGP peers
- C. It sends the update only to nonclients
- D. It sends the update only to its configured route reflector clients
- E. It send the update to all routers in the autonomous system

Answer: B

Explanation:

When a route reflector receives an update from a peer in an external AS, the routes are only advertised to the reflector's peers (established BGP sessions). This would mean that the new route would be sent to the reflector's clients and/or other configured route reflectors within the AS. The reflector that receives the update would NOT send to all routers within the AS.

Chapter: 2

---

#### **QUESTION 605:**

Which methods advertises your internal networks to external ISPs via BGP? Select two.

- A. using aggregate routes
- B. disabling synchronization
- C. forcing the next-hop address
- D. defining routes via the network statement

Answer: A, D

Explanation:

BGP will advertise the network number and mask specified in the network statement unless the community attribute is changed to NO ADVERTISE or a route filter is used to block the advertisement.

Chapter: 2

---

**QUESTION 606:**

You are the network administrator at Certkiller . Router CK1 is the headquarters router in a hub and spoke topology supporting 24 remote offices. Point-to-point Frame Relay EIGRP network is deployed between the headquarters and the remote offices. The CIR for each Frame Relay PVC is different and that there is no bandwidth command configured under either the major serial interface nor the subinterfaces on Router CK1 .

What is a possible fix for the potential EIGRP packet pacing problem because of the different CIR each PVC has?

- A. convert each Frame Relay PVC to a point-to-multipoint connection
- B. manually configure the bandwidth of the major interface to the lowest CIRx24
- C. manually configure the bandwidth of the major interface to the highest CIRx24
- D. manually configure the bandwidth of each of these PVCs to equal to their respective CIR

Answer: D

Explanation:

EIGRP assumes that all serial interfaces operate at T-1 speed. By configuring a bandwidth for each subinterface, EIGRP can identify slow links ( < T-1 ) and will not generate packets faster than the configured line speed.

Chapter: 2

---

**QUESTION 607:**

Which type of packet is used by EIGRP routers build a neighbor table?

- A. hello
- B. ACK
- C. LSA-Type 1

- D. query
- E. update

Answer: A

Explanation:

The hello protocol is used to establish neighbor relationships on a common network. Two routers become neighbors (establish adjacency) when they acknowledge each other's hello packets and their K values match

Chapter: 2

---

**QUESTION 608:**

You are a trainee technician at Certkiller . Your instructor shows you the following configuration commands:

- 1) router eigrp 200
- 2) network 172.16.0.0
- 3) network 3.0.0.0

Your instructor wants to know which of the following statements are true.

What would your reply be? (Choose all that apply.)

- A. Line 1 defines EIGRP as an ip routing process in area 200
- B. The command network 172.16.0.0 causes this router to become the access point for the default network
- C. Line 2 causes all interfaces connected to network 172.16.0.0 to send and receive EIGRP updates to/from other EIGRP routers
- D. The number at the end of line 1 restricts this EIGRP routing process to only communication with other EIGRP routing process that have the same number

Answer: C, D

Explanation:

EIGRP will send/receive updates on all interfaces that are included on the specified network. The number at the end of the line is the autonomous system number and routers will only exchange protocol information with routers within the same autonomous system.

Chapter: 2

---

**QUESTION 609:**

What appears in a routing table after EIGRP route summarization is configured on a router's Serial0 interface summarizing routes learned from Ethernet0 interface?

- A. a summary route pointing to the Null0 interface
- B. a summary route pointing to the Serial0 interface
- C. a summary route pointing to the Ethernet0 interface

D. a summary route pointing to the Loopback0 interface

Answer: A

Explanation:

EIGRP creates a null route that matches the summary-address network and mask entry to prevent routing loops.

Chapter: 2

---

**QUESTION 610:**

You are the network administrator at Certkiller . An EIGRP router on the Certkiller network has not established adjacency with a neighbor.

What are the possible causes for this? (Choose all that apply.)

- A. K-values do not match
- B. Hold times do not match
- C. Hello times do not match
- D. AS numbers do not match

Answer: A, D

Explanation:

If the autonomous system (AS) numbers do not match, the routers will not form an adjacency. When EIGRP is enabled on an interface, the router begins sending hellos to a multicast address. The hello packet includes the configured EIGRP metric K values. The two routers become adjacent if their K values match.

Chapter: 2

---

**QUESTION 611:**

You are a technician at Certkiller . Your newly appointed Certkiller trainee wants to know what the correct command format to configure EIGRP summary route is.

What would your reply be?

- A. ip auto-summary as-number address mask
- B. ip summary-address eigrp as-number address mask
- C. ip auto-summary eigrp as-number address mask
- D. ip summary-route eigrp as-number address mask

Answer: B

Explanation:

This is the correct syntax to summarize external routes in EIGRP.

Chapter: 2

---

**QUESTION 612:**

You are a technician at Certkiller . Your newly appointed Certkiller trainee wants to know how bandwidth information per neighbor is determined on a multipoint Frame Relay interface.  
What would your reply be?

- A. bandwidth command per neighbor
- B. the configured CIR per subinterface
- C. the configured CIR divided by the number of neighbors on that interface
- D. bandwidth of the main interface divided by the number of neighbors on that interface

Answer: B

Explanation:

The bandwidth is assumed to be 1.54Mbps on serial interfaces by EIGRP. The only way to overcome this assumption, is to configure the bandwidth on each subinterface.

Chapter: 2

---

**QUESTION 613:**

Which command shows the active or passive state of EIGRP routes, the number of successors, and the feasible distance to the destination?

- A. show ip route eigrp
- B. show ip eigrp traffic
- C. show ip eigrp neighbors
- D. show ip eigrp topology

Answer: D

Explanation:

If the show ip route command were issued, only the current route would appear in the routing table. The EIGRP traffic and neighbor command do not show any information about the routes in the topology database.

IP-EIGRP Topology Table for process 77

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,  
r - Reply status

P 172.16.90.0 255.255.255.0, 2 successors, FD is 0  
via 172.16.80.28 (46251776/46226176), Ethernet0  
via 172.16.81.28 (46251776/46226176), Ethernet1  
via 172.16.80.31 (46277376/46251776), Ethernet0

P 172.16.81.0 255.255.255.0, 1 successors, FD is 307200  
via Connected, Ethernet1  
via 172.16.81.28 (307200/281600), Ethernet1  
via 172.16.80.28 (307200/281600), Ethernet0

via 172.16.80.31 (332800/307200), Ethernet0

Chapter: 2

---

**QUESTION 614:**

What are two classless routing protocol features supported by EIGRP? (Choose all that apply.)

- A. triggered updates
- B. variable length subnet masks
- C. periodic update announcements
- D. unequal path-cost load balancing

Answer: A, B

Explanation:

Eigrp does not send out periodic updates only triggered.

---

**QUESTION 615:**

When point-to-point Frame Relay sub-interfaces are used on a router running EIGRP, what is the default line speed used in calculating routes associated with the WAN interfaces?

- A. 256 Kbps
- B. 1.544 Mbps
- C. set by the PVC
- D. set by the DLCI

Answer: B

Explanation:

EIGRP assumes that all serial interfaces are operating at T-1 speed. If the actual line speed is different, the interface line speed should be specified using the bandwidth command.

Chapter: 2

---

**QUESTION 616:**

You are the network administrator at Certkiller . Router CK1 is the central router in a hub and spoke topology supporting 24 remote locations. Point-to-point Frame Relay EIGRP networks are deployed between the central router and the remote locations. There is no bandwidth command configured under either the major serial interface or the subinterfaces on the central router.

What does EIGRP perceive as the bandwidth of each Frame Relay connection?

- A. 64 kbps
- B. 128 kbps
- C. 512 kbps
- D. 1.544 Mbps

Answer: D

Explanation:

EIGRP assumes that all serial interfaces are operating at T-1 speed. If the actual line speed is different, the interface line speed should be specified using the bandwidth command.

Chapter: 2

---

**QUESTION 617:**

What is a use of VLSM in EIGRP?

- A. disjointed networks
- B. address cumulative
- C. address aggregation
- D. contiguous networks

Answer: A

Explanation:

Variable Length subnet masks allows the router to separate networks according to the amount of Hosts required per network. Disjointed networks would be the result of using EIGRP. EIGRP does not require Contiguous Network Assignment, as it does not require hierarchical Address.

Chapter: 2

---

**QUESTION 618:**

You are a technician at Certkiller . Your newly appointed Certkiller trainee wants to know what could possibly cause EIGRP Stuck-In-Active routes.

What would your reply be? (Choose all that apply.)

- A. The neighboring router has a better route than the active route on this router
- B. The neighboring router starts receiving route updates from this router
- C. The neighboring router is having memory problems and cannot allocate the memory to process the query or build the reply packet
- D. Packets are being dropped because EIGRP is sending packets faster than the actual line speed of the circuit.

Answer: C, D

Explanation:

An EIGRP SIA message means that an EIGRP router hasn't received a reply to a query from one or more neighbors within the allotted time. When this happens, EIGRP clears the neighbors that didn't send a reply and logs a DUAL-3-SIA error message for the route that went active.

Chapter: 2

---

**QUESTION 619:**

Which EIGRP information is added to a routing table?

- A. successor only
- B. feasible successor only
- C. successor and back up successor
- D. successor and feasible successor

Answer: A

Explanation:

EIGRP maintains its own topology database where up to six routes for every destination can be stored. Only the best route (successor) is installed into the routing table. A backup is registered with the routing table maintenance process, but not installed into the routing table. In the event the route in the routing table fails, the routing table maintenance process calls each routing protocol process that has registered a backup route, and asks them to reinstall the route in the routing table. Then, the route with the preferred route is chosen based administrative distance.

Chapter: 2

---

**QUESTION 620:**

Which is the most effective technique to contain EIGRP queries?

- A. route summarization
- B. configuring route filters
- C. using a hierarchical addressing scheme
- D. establishing separate autonomous systems

Answer: A

Explanation:

After you determine the minimum routing requirements, you can make EIGRP more scalable. Two of the best options are the following: 1. Configure route summarization using the ip summary-address eigrp command on the outbound interface of the appropriate routers 2. Configure the remote routers as stub EIGRP routers. Summarizing routes limits the queries scope by limiting a routers knowledge of networks subnets. If a subnet goes down, queries go only as far as the routers that have knowledge of that

subnet.

Source: CCNP Self-Study Second Edition P.185

Topic: Limiting the EIGRP Query Range

---

**QUESTION 621:**

What is the default hold time for EIGRP hellos on NBMA media?

- A. 30 seconds
- B. 60 seconds
- C. 90 seconds
- D. 180 seconds

Answer: D

Explanation:

EIGRP sends hello packets every 5 seconds on high bandwidth links and every 60 seconds on low bandwidth NBMA media. The default hold time is three times the hello interval or 180 seconds for NBMA 5-second hello:

broadcast media, such as Ethernet, Token Ring, and FDDI point-to-point serial links, such as PPP or HDLC leased circuits, Frame Relay point-to-point subinterfaces, and ATM point-to-point subinterfaces high bandwidth (greater than T1) multipoint circuits, such as ISDN PRI and Frame Relay

Chapter: 2

60-second hello:

multipoint circuits T1 bandwidth or slower, such as Frame Relay multipoint interfaces, ATM multipoint interfaces, ATM switched virtual circuits, and ISDN BRIs

Chapter: 2

---

**QUESTION 622:**

With regard to iBGP routers, which of the following statements is true?

- A. They are level-1 routers
- B. They are level-2 routers
- C. They are in the same AS
- D. They are in a different AS

Answer: C

Explanation:

BGP routers that are all within the same AS are considered to be internal BGP routers (iBGP).

Chapter: 2

---

**QUESTION 623:**

You are a technician at Certkiller . Your newly appointed Certkiller trainee wants to know what the characteristics of an autonomous system in a BGP network is. What would your reply be? (Choose all that apply.)

- A. It used only Interior Gateway Protocols (IGPs)
- B. EGPs are used to connect different autonomous systems
- C. It is a set of routers under a single technical administration
- D. It uses EGPs to route packets to other autonomous systems and IGPs to route packets within the autonomous system
- E. It uses IGPs to route packets to other autonomous systems and EGPs to route packets within the autonomous system

Answer: B, C, D

Explanation:

EGP (Exterior Gateway Protocol) is used to exchange routing information between two different autonomous systems. IGP (Interior Gateway Protocol) is used to exchange routing information within an autonomous system. Routers within the same AS are normally under a single technical administration.

Chapter: 2

---

**QUESTION 624:**

You are a trainee technician at Certkiller . Your instructor shows you the following router configuration:

```
router bgp 65000
neighbor 172.16.1.1 remote-as 65000
neighbor 10.1.1.2 remote-as 64550
network 192.168.1.192 mask 255.255.255.224
```

Your instructor wants to know which of the following statements is true if interfaces 192.168.1.193 172.16.1.1 and 10.1.1.2 are active.

What would your reply be? (Choose all that apply.)

- A. Router A is in autonomous system 64550
- B. Router A advertises network 192.168.1.192/27
- C. Router A forms an iBGP relationship with neighbor 10.1.1.2
- D. Router A forms and eBGP relationship with neighbor 172.16.1.1

Answer: B

Explanation:

Router A is in AS 65000 as indicated on line 1.

Line 2 specifies 172.16.1.1 as a neighbor and belonging to the same AS (iBGP peer).

Line 3 specifies 10.1.1.2 as a neighbor and belonging to a different AS (eBGP peer).

Line 4 specifies that 192.168.1.192/27 be advertised

Chapter: 2

---

**QUESTION 625:**

In which instance is it appropriate to use BGP?

- A. If there is single connection to the Internet
- B. If you have limited understanding of route filtering
- C. If there is a low-bandwidth connection between autonomous systems
- D. If route selection to routes outside of your autonomous system is not a concern
- E. If an autonomous system allows packets to transit through it to reach other autonomous systems

Answer: E

Explanation:

This is the only reason listed that would require you to run BGP. If you are serving as a transit AS for other downstream AS's, then you must run BGP so that all paths are known into the downstream AS's.

Chapter: 2

---

**QUESTION 626:**

You are a trainee technician at Certkiller . Your instructor shows you the following router configuration on Router CK1 :

```
router bgp 65000
network 10.0.0.0
neighbor 172.17.1.1 remote-as 65000
```

Your instructor wants to know what type of relationship is neighbor 172.17.1.1 to Router CK1 .

What would your reply be? (Choose all that apply.)

- A. a peer router running iBGP
- B. a peer router running eBGP
- C. a community member running iBGP
- D. a peer group member running iBGP
- E. a peer group member running eBGP

Answer: A

Explanation:

Line 1 specifies that Router A belongs to AS 65000 and the neighbor statement on line 3 indicates that the neighbor is also a member of AS 65000. Therefore, these routers are considered to be iBGP peers.

Chapter: 2

**QUESTION 627:**

Which two statements are true about BGP peering? (Choose two.)

- A. Periodic keepalives are used to verify connectivity
- B. Incremental keepalives are used to verify connectivity
- C. It provides a "best effort" connection between two BGP providers
- D. It provides a reliable connection between two BGP providers

Answer: A, D

Explanation:

BGP uses periodic keepalives to maintain connectivity. The interval can be changed to suit your needs (fast fail-over). Each neighbor sessions runs over TCP (port 179) and ensures reliable delivery of routing information.

Chapter: 2

---

**QUESTION 628:**

You are a trainee technician at Certkiller . Your instructor shows you the following router configuration:

```
interface serial 0
ip address 172.16.1.1 255.255.255.0
!
interface ethernet 0
ip address 10.1.1.1 255.255.255.0
!
router rip
network 10.0.0.0
!
ip route 0.0.0.0 0.0.0.0 serial0
!
```

Your instructor wants to know which of the following statements is true.

What would your reply be?

- A. RIP updates are sent and received on interface serial 0 of the router
- B. A default route is sent to neighbors on interface serial 0 of the router
- C. A default route is sent to neighbors on interface ethernet 0 of the router
- D. RIP updates are sent and received on interfaces serial 0 and ethernet 0 of the router

Answer: D

Explanation:

RIP is a true classful routing protocol. One of the problems with classful routing protocols is that periodic routing updates are sent out all active interfaces on every router.

Even if the router is not running RIP, it will still receive broadcast RIP packets. RIPv2 is a classless routing protocol and instead of using broadcast updates it sends multicast packets to a multicast address of 224.0.0.9.

Chapter: 2

---

**QUESTION 629:**

When using BGP policy-based routing, which two statements are true? (Choose two.)

- A. Policy routing cannot be used to modify the AS-path
- B. Policy routing can be used to alter the final destination of the IP packet
- C. Policy routing allows traffic to be directed based on the source address
- D. Policy routing can influence which router will be used as the next-hop router for a given packet

Answer: C, D

Explanation:

Answers A and B are false. With the use of route maps, BGP AS paths may be lengthened by adding fictitious AS numbers. This technique is called AS path pre-pending. The final destination of a packet cannot be altered by a routing protocol or a routing protocol policy.

Chapter: 2

---

**QUESTION 630:**

When should BGP synchronization be unnecessary?

- A. when only the edge routers in the AS will be running BGP
- B. when traffic from a different AS passes through an AS to a third AS
- C. when traffic from a different AS will not pass through an AS to a third AS
- D. when sending and receiving of external BGP updates is controlled by using a number of different filtering methods

Answer: C

Explanation:

All of the other items require BGP synchronization. When traffic from a different AS passes through an AS to a third AS, BGP will not advertise the route until all routers within the AS have learned of the route through IGP.

Chapter: 2

---

**QUESTION 631:**

Identify two statements regarding BGP peer groups from the choices below?

(Choose all that apply.)

- A. The peer group name is passed to other routers in the peer group during routing updates.
- B. A peer group is a group of BGP neighbors with different update policies.
- C. The peer group name is only local to the router on which it is configured.
- D. A peer group allows options that affect outbound updates to be overridden.
- E. BGP configurations can be placed on one peer group router and the configuration is applied to all members of the peer group.

Answer: C, E

Explanation:

A BGP peer group is a group of routers that share similar configurations. Every configuration line supplied to a peer group definition is applied to each peer group member. A peer group name may be specified, but it is not passed to any other router; the name is local only to the router it is configured on.

Chapter: 2

---

### **QUESTION 632:**

With regard to BGP attributes, which of the following statements are true? (Choose all that apply.)

- A. Med is an optional attribute
- B. Origin is an optional attribute
- C. Next-hop is an optional attribute
- D. Local Preference is an optional attribute
- E. AS-Path is a well-known mandatory attribute
- F. Community is a well-known mandatory attribute

Answer: A, E

Explanation:

MED is an optional attribute

ORIGIN is a well-known mandatory attribute

NEXT\_HOP is a well-known mandatory attribute

LOCAL\_PREF is a well-known discretionary attribute

AS\_PATH is a well-known mandatory attribute

COMMUNITY is an optional attribute

Chapter: 2

Cisco Press BSCI Third edition, Pages 526-527

---

### **QUESTION 633:**

With regard to the network command on a BGP router, which of the following

statements is true?

- A. The local route matching the network command can be learned dynamically
- B. The local route matching the network command are blocked from the BGP routing table
- C. The route to a neighbor autonomous system must have the correct MED applied to be installed into BGP's routing table
- D. The specified network is identified as a transit AS and traffic must pass through this AS to reach its final destination.

Answer: A

Explanation:

The network command only specifies which networks are to be advertised by BGP. This can be learned from static routes, other routing protocols, or directly connected interfaces.

Chapter: 2

---

#### **QUESTION 634:**

With regard to BGP community attributes, which of the following statements are true?

- A. Communities are tagged by default in outgoing updates
- B. Communities are local to the autonomous system where specified and can only be used within that autonomous system
- C. Communities are a means of tagging routes to ensure consistent filtering
- D. Communities perform summarization of blocks of contiguous network prefixes

Answer: C

Explanation:

The BGP COMMUNITY attribute is used to tag/mark routes. Once these routes are marked, route maps can be used to limit the distribution and acceptance of routes with a particular mark. Some commonly used communities:

No-Export: The route will not be passed outside the AS

No-Advertise: The route will not be advertised to other routers

No-Export-Subconfed: Routes will not be advertised to eBGP peers (including eBGP peers in the same confederation)

Chapter: 2

---

#### **QUESTION 635:**

With regard to BGP policy-based routing, which of the following statements is true?

- A. If the next-hop router goes down and no alternative path is in place, policy routing

will route to null 0

B. If the next-hop router goes down and no alternative path is in place, policy routing will default to another BGP path

C. If the next-hop router goes down and no alternative path is in place, policy routing will deny all traffic to that destination

D. If the next-hop router goes down and no alternative path is in place, policy routing will default to dynamic routing decisions.

Answer: D

Explanation:

Some things should be considered before arbitrarily deciding to implement policy-based routing. Understand that any additional configurations require additional CPU, particularly when every packet characteristic must be examined. It is also wise to have a backup path in place in case the defined next-hop router goes down. If there is no alternative defined, policy-routing will default to dynamic routing decisions.

Source: Self-Study CCNP BSCI Exam Certification Guide Third Edition P.523

Topic: BGP and Policy-Based Routing

A backup path should be defined in case the defined next-hop router goes down. If there is no alternative defined, policy-based routing uses the IP routing table

Source: Self-Study CCNP BSCI Exam Certification Guide Third Edition P.675

Topic: Disadvantages of Policy-Based Routing

---

### **QUESTION 636:**

You are a trainee technician at Certkiller . Your instructor shows you the following router configuration for Router Certkiller 1:

```
ROUTER Certkiller 1
```

```
router bgp 500
```

```
neighbor 190.225.11.1 remote-as 500
```

```
neighbor 190.225.11.1 update-source loopback 1
```

```
ROUTER Certkiller 2
```

```
router bgp 500
```

```
neighbor 150.212.1.1 remote-as 500
```

Your instructor wants to know which of the following statements are true.

What would your reply be? (Choose all that apply.)

A. ROUTER Certkiller 1 and ROUTER Certkiller 2 are running iBGP inside as 500

B. The IP address of ROUTER Certkiller 1's loopback 1 interface is 150.212.1.1

C. The IP address of ROUTER Certkiller 1's loopback 1 interface is 190.225.11.1

D. ROUTER Certkiller 1 and ROUTER Certkiller 2 are running eBGP between autonomous systems

E. ROUTER Certkiller 1 is configured to use the loopback IP address as the source in the BGP neighbor connection with neighbor 190.225.11.1

Answer: A, B, E

Explanation:

We believe that the answers should A B and E as the address of router Certkiller 1's loopback interface is shown in the output of router B as 150.212.1.1 and not 190.225.11.1 We understand what the question is getting at, it is testing the readers knowledge to see if they understand that the "update-source loopback" command forces the BGP to use the loopback address of the neighbour router. Therefore if answer C asked for the loopback address of router Certkiller 2 it would be correct

---

**QUESTION 637:**

When the default-information originate always command used?

- A. It is required whenever you want to propagate a default route into a RIP autonomous system
- B. It is required whenever you want to propagate a default route into an IGRP autonomous system
- C. It is required whenever you want to propagate a default route into an OSPF autonomous system
- D. It is required whenever you want to propagate a default route into an EIGRP autonomous system

Answer: C

Explanation:

This command is used to direct all other OSPF routers to place a default route into its routing table. This command is not used with IGRP/EIGRP. A similar command is used in RIP to specify some interfaces as the default route even when the local router does not have a default route in its routing table.

Chapter: 2

---

**QUESTION 638:**

Which routing protocol uses the Diffusing Update Algorithm (DUAL) for route calculation?

- A. BGP
- B. OSPF
- C. EIGRP
- D. RIPv2

Answer: C

Explanation:

DUAL is used by EIGRP. This allows EIGRP to achieve fast, loop-free convergence with

little impact on CPU cost and overhead. DUAL takes corrective action when topology changes occur and in doing so, only involves the routers that are affected.

Chapter: 2

---

**QUESTION 639:**

Which fields is included in a RIP version 2 routing update packet that is not included in RIP version 1?

- A. metric
- B. next hop
- C. subnet mask
- D. autonomous system number

Answer: C

Explanation:

RIPv1 only carries the next hop address and its associated metric in its routing updates.

One of the major improvements in RIPv2 is that is a classless routing protocol. This means it carries the subnet mask along with the network number in its routing updates.

Chapter: 2

---

**QUESTION 640:**

You are a technician at Certkiller . Your newly appointed Certkiller trainee wants to know what a router must determine in order to route data.

What would your reply be?

- A. the route age of the next-hop device
- B. the subnet mask of the source network
- C. the cost metric of the path to the destination
- D. the outbound interface of the best path to the destination

Answer: D

Explanation:

The router must know which interface that the data will be forwarded to. The other items listed are not used to determine if the router can route the data.

Chapter: 2

---

**QUESTION 641:**

Which routing protocol does not have a hop count limit?

- A. IGRP
- B. OSPF

- C. EIGRP
- D. RIPv1
- E. RIPv2

Answer: B

Explanation:

RIP has a maximum allowable hop count of 15 meaning a value of 16 is considered unreachable. IGRP and EIGRP have a limit of 255 (100 by default).

Chapter: 2

---

**QUESTION 642:**

You are the network administrator at Certkiller . Certkiller uses the RIPv2 routing protocol. Due to a dropped connection, your core router has not received an update from a neighboring router for four minutes and the route is marked invalid.

What is the next step for the core router?

- A. It recalculates the network topology
- B. It purges that link from its routing table
- C. It places a hold-down on the routes from that link
- D. An invalid timer is started

Answer: D

Explanation:

RIP versions 1 and 2 use timed updates to respond to changes in topology. Updates are sent every 30 seconds by all participating routers. If a neighbor does not respond for 180 seconds, the router assumes the link or neighbor is down. The router then marks the route as invalid and an invalid timer is started (180seconds). When the invalid timer expires, a hold-down time begins. During this period, the route is marked as possibly down and the metric is set and advertised with infinity. If a new route was received during the invalid or hold-down periods, the router begins advertising the new route. If not, a request is sent to neighbors for an alternate route to the destination. The last timer is the flush timer. It is activated when the invalid timer expires and last 60 seconds after the hold-down timer expires. When the flush timer expires (240 seconds), the route is finally deleted if no alternate route was found.

Chapter: 2

---

**QUESTION 643:**

Which routing protocol provides support for load balancing over unequal cost paths?

- A. IGRP
- B. OSPF

- C. RIPv1
- D. RIPv2

Answer: A

Explanation:

IGRP and EIGRP provide support for load balancing via unequal cost paths using the variance command. In order for RIP to perform balancing, the hop count must be the same as RIP looks at hop count only when determining the best path to a destination. IGRP and EIGRP use a composite metric to determine the cost to a remote network. The variance command controls the load balancing between the best path and the worst acceptable path.

Chapter: 2

---

**QUESTION 644:**

An LSA is used by which routing protocol to send topology change information across the network?

- A. RIP
- B. IGRP
- C. OSPF
- D. EIGRP

Answer: C

Explanation:

The OSPF routing protocol uses different types of LSA to send information to other OSPF routers. An OSPF router's link state database is comprised of LSA's Listed below are the five types of LSA's used by the OSPF protocol:

- LSA Type 1 Router link state
- LSA Type 2 Network link state
- LSA Type 3 Summary link state (type 3)
- LSA Type 4 Summary link state (type 4)
- LSA Type 5 External link state

Chapter: 2

---

**QUESTION 645:**

Which of the following is a characteristic of link-state protocols?

- A. A network topology change generates an update that is broadcast to every router on the network.
- B. The periodic routing update interval is slightly different in each router on a subnet.
- C. Each router creates a routing table that only includes its directly connected neighbors.
- D. Link-state protocols send routing updates only when there is a change in the network

topology.

Answer: D

Explanation:

One of the differences between distance-vector and link-state routing protocols is the way topology changes are conveyed across the network. Distance-vector protocols use periodic updates and send out complete routing tables. If a link goes down somewhere in the network, routers running only distance-vector protocols will not know about the change until the next update.

Link-state routing protocols send updates only when there is a topology change and the change is also included within the update. In addition, only the change is sent in the update. This means that each router does not have to rebuild its entire routing table with each update. The update is either added or removed from the routing table.

Chapter: 2

---

**QUESTION 646:**

What is the term associated with synchronizing the routing tables after a topology change occurs?

- A. flooding
- B. broadcasting
- C. convergence
- D. summarization

Answer: C

Explanation:

Convergence is the time required for all routers on the network to agree on the network topology after a change has occurred. Simply stated, all routers synchronize the routing tables with the same information.

Chapter: 2

---

**QUESTION 647:**

Which of the following must be determined first by the router in order to route the data?

- A. the distance metric of the data
- B. the source address of the data
- C. which routing protocol is used by the data
- D. whether the protocol suite of the data is active

Answer: C

Explanation:

Before a routing decision can be made three major decisions must be made in the following order:

- 1) Is the logical destination address a known protocol and is the protocol active/enabled on the router?
- 2) Is the destination address in the routing table?
- 3) If the destination address is in the routing table, which interface will be used to forward the data?

Chapter: 2

---

**QUESTION 648:**

You are a technician at Certkiller . Certkiller has the subnet/mask of 172.29.100.0/26 set aside for small remote locations. Each location will have 5 IP devices connected to the network.

Which VLSM mask will provide the minimum number of hosts you require?

- A. /27
- B. /28
- C. /29
- D. /30
- E. /31

Answer: C

Explanation:

/29 mask is equal to 255.255.255.248 in decimal format. If 29 bits are used for the network portion of the address, then 3 bits are left for the host portion of the address.

Once the reserved network and the broadcast addresses are removed, there are six usable addresses for hosts on the subnet.  $2 \times 2 \times 2 = 8$  minus network and broadcast address = 6 valid host addresses (Requirement from question is 5).

Chapter: 2

---

**QUESTION 649:**

You are a trainee technician at Certkiller . Your instructor shows you the following route summarization entry:

192.168.134.0/22

Your instructor wants to know how many class C addresses are contained in this summarization.

What would your reply be?

- A. 1
- B. 2
- C. 4
- D. 8

E. 16

Answer: C

Explanation:

The valid class C addresses are:

192.168.134.0

192.168.135.0

192.168.136.0

192.168.137.0

Chapter: 2

---

**QUESTION 650:**

Which two UDP ports are NOT enabled automatically when the ip helper-address command is used on a router? (Choose all that apply.)

A. 53 (DNS)

B. 69 (TFTP)

C. 515 (LPR)

D. 161 (SNMP)

E. 49 (TACACS)

Answer: C, D

Explanation:

By default, if an IP helper address is specified, UDP forwarding is enabled on certain ports. Broadcast packets destined to the following port numbers are forwarded by default:

Trivial File Transfer Protocol (TFTP) (port 69)

Domain Naming System (port 53)

Time service (port 37)

NetBIOS Name Server (port 137)

NetBIOS Datagram Server (port 138)

Boot Protocol (BOOTP) client and server datagrams (ports 67 and 68)

TACACS service (port 49)

IEN-116 Name Service (port 42)

Chapter: 2

---

**QUESTION 651:**

You are a technician at Certkiller . Your newly appointed Certkiller trainee wants to know what the benefit of hierarchical IP addressing is.

What would your reply be?

A. support for network address translation

B. efficient address allocation

- C. translation of private addresses
- D. support for link-state routing protocols

Answer: B

Explanation:

The only listed benefit of hierarchical IP address is efficient address allocation. Another major benefit in large enterprise networks is the reduced number of entries in the routing table. Combined with link-state routing protocols, overall CPU and memory requirements are reduced on deployed routers.

Chapter: 2

---

**QUESTION 652:**

What entry would be used to summarize all the address space between 172.18.129.0 and 172.18.133.255?

- A. 172.18.128.0/21
- B. 172.18.128.0/22
- C. 172.18.130.0/22
- D. 172.18.132.0/20

Answer: A

Explanation:

172.18.128.0/21 is the only network/mask combination that will cover the entire listed range. Using this network/mask, the valid subnets would be (partial listing):

172.18.0.0 172.18.104.0 172.18.208.0  
172.18.8.0 172.18.112.0 172.18.216.0  
172.18.16.0 172.18.120.0 172.18.224.0  
172.18.24.0 172.18.128.0 172.18.232.0  
172.18.32.0 172.18.136.0 172.18.240.0  
172.18.40.0 172.18.144.0 172.18.248.0  
172.18.48.0 172.18.152.0  
172.18.56.0 172.18.160.0  
172.18.64.0 172.18.168.0  
172.18.72.0 172.18.176.0  
172.18.80.0 172.18.184.0  
172.18.88.0 172.18.192.0  
172.18.96.0 172.18.200.0

Chapter: 2

---

**QUESTION 653:**

With regard to route filtering, which of the following statements are true?

- A. It uses the same technique as packet filtering; it's based on standard access-lists
- B. It is required when using EIGRP and OSPF at the same time
- C. Unlike packet filtering (which uses inverted subnet masks, route filtering uses normal subnet masks
- D. Route filtering is used only while performing route redistribution

Answer: A

Explanation:

Routes can be filtered so that only a subset of routes can be advertised or received. Two configuration commands are necessary to filter a route. First, a standard access list is created that specifies which networks are allowed/denied. Second, the distribute-list command is configured for the routing protocol.

```
!  
access-list 2 permit 192.168.1.0 0.0.0.255  
access-list 2 permit 192.168.2.0 0.0.0.255  
!  
router eigrp 100  
distribute-list 2 in
```

```
!  
Chapter: 2
```

---

**QUESTION 654:**

Which of these are benefits of incorporating hierarchical addressing?

- A. You can summarize multiple routes into a single route summaries, making the network more scaleable
- B. A contiguous address assignment allows the most efficient use of address
- C. Reduction in the number of routing table entries for participating routers
- D. A more efficient allocation of addresses network-wide

Answer: B, C, D

Explanation:

Incorporating hierarchical and contiguous addressing minimizes the number of wasted addresses, reduces the number of entries in routing tables (summarizing), and simplifies network design and administration.

Chapter: 3

---

**QUESTION 655:**

How can a single IP statement indicate many IP addresses?

- A. Source-route bridge

- B. Route summarization
- C. Helper address
- D. Default gateway
- E. Passive interface

Answer: B

Explanation:

A single IP network entry can represent multiple subnets using route summarization and VLSM. Classless routing protocols carry/distribute the mask information when managing routing tables and sending updates to topology changes.

Chapter: 3

---

**QUESTION 656:**

What is NOT a reason reasons makes OSPF a better and/or more efficient routing protocol over RIP?

- A. OSPF can select paths based on bandwidth
- B. Link-State protocols generally have faster convergence than Distance-Vector
- C. OSPF sends its topology database to configured neighbors on a regular basis to ensure that all routers have the same view of the network
- D. OSPF supports VLSM
- E. OSPF has no hop count limitation because it is based on path cost

Answer: C

Explanation:

RIP periodically (60 seconds) broadcasts its entire routing table to all routers in the network. RIP also employs a series of timers that must expire before alternate routes can be discovered or routes removed in the event of a link failure. These are considered drawbacks to distance-vector protocols as it increases the time for the network to convergence.

Chapter: 3

---

**QUESTION 657:**

With regard to EIGRP, which of the following are true?

- A. It only sends change-updates to the systems that are affected by the event
- B. It can provide client address conflict resolution
- C. It is considerably more complex than OSPF
- D. It supports Variable Length Subnet Masking
- E. It supports compression over the WAN link

Answer: A, D

Explanation:

EIGRP is considered an advanced distance-vector routing protocol. EIGRP uses the DUAL (Distributed Update Algorithm) to ensure that updates and queries are not propagated beyond affected routers. EIGRP is also a classless routing protocol meaning that the network mask information is retained and advertised by the protocol. By default, EIGRP advertises the natural classful network boundary for all EIGRP internal routes. This behavior can be modified using the no auto-summary command and EIGRP will leave the network mask unchanged.

```
!  
router eigrp 93  
no auto-summary  
!  
Chapter: 3
```

---

**QUESTION 658:**

Which of the statements below correctly describes the BGP regular expression ^200\_[0-9]\*\$ ?

- A. Matches AS path 200 only
- B. Matches all AS's
- C. Matches AS path that is originated in AS 200 regardless of location
- D. Matches AS path that is originated in AS 200 and AS's that are directly connected to AS 200

Answer: D

Explanation:

You can use regular expressions in the ip as-path access-list command with BGP. This allows you to setup filters to allow incoming and outgoing updates based on the BGP autonomous system paths. In the example above, the ^ starts the input string and designates "AS". The \_ signifies a null string or space follows 200 and [0-9]\*\$ indicates that any connected AS with a valid AS string can pass the filter. The \$ matches the end of the input string.

Chapter: 3

---

**QUESTION 659:**

You are a technician at Certkiller . Your newly appointed Certkiller trainee wants to know what a BGP peer is.

What would your reply be?

- A. A BGP neighbor that the router has formed a TCP connection with
- B. Another BGP router that the router is currently speaking with
- C. Another router on the network running BGP

D. A BGP neighbor with a broken TCP connection

Answer: A, B

Explanation:

If BGP cannot establish the TCP connection with its specified neighbor, the routers are not considered to be peering and thus not peers. A peering relationship only exists between routers that explicitly identify other routers as neighbors. BGP does not "auto-discover" neighbors as other routing protocols such as EIGRP.

Chapter: 3

---

**QUESTION 660:**

With regard to BGP, which of the following statements are true?

- A. iBGP is used between the AS
- B. iBGP is used within the AS
- C. eBGP is used between the AS
- D. eBGP is used within the AS

Answer: B, C

Explanation:

Internal BGP (iBGP) sessions are contained within the autonomous system. The active session is between routers with the same AS number.

External BGP (eBGP) sessions are connections between two different autonomous systems. The active session is between routers with different AS numbers.

Chapter: 3

---

**QUESTION 661:**

OSPF stub or totally stubby Areas can only exist if adjacent to which of the following?

- A. LSA
- B. LSP
- C. BDR
- D. ABR
- E. DR
- F. LSU

Answer: D

Explanation:

There are only two possible answers - BDR and ABR. The function of the BDR is not directly related to stub or stubby areas. The BDR is a hot standby for the DR and would

assume the role of the DR in the event the DR failed. The purpose of a DR is to minimize the number of adjacencies formed and disseminate/receive routing updates throughout the network. An ABR has multiple area assignments and two or more interfaces in two or more areas. An ABR would be used to connect a stub or totally stubby area to the rest of the network.

Chapter: 3

---

**QUESTION 662:**

You are the network administrator at Certkiller . You stop sending RIP broadcasts from all router interfaces that do not require it. You notice that RIP is broadcasting on Ethernet0, Serial0, and Serial1 on the router in question. You determine that Serial0 does not need to participate in RIP updates.

What command would you use to disable RIP on this interface?

- A. Configure RIP using the 'passive interface' command
- B. Configure RIP using the 'no router rip' command
- C. Configure Serial0 using the 'shutdown rip' command
- D. Configure the serial interface using the passive interface command

Answer: A

Explanation:

The passive interface would be defined in the RIP configuration as follows:

```
!  
router rip  
network 192.168.1.0  
passive interface serial 0  
!
```

Chapter: 3

---

**QUESTION 663:**

Why will an EIGRP route get stuck in the Active state?

- A. The EIGRP router on the far end of the link has crashed
- B. The reply to a query caused by a change in a route takes too long and times out
- C. The physical link between the two AS's has gone down
- D. The network is unstable and probably has a routing loop

Answer: B

Explanation:

EIGRP sends out queries to neighbors when it detects a change in topology. In some situations, the response to the query could be delayed longer than the router issuing the query is willing to wait. When this happens, the route in question is stuck in the active

state. This means that EIGRP is manipulating the route and cannot route data to the remote network. Eventually, the router sending the query will give up and clear its connection with the neighbor. This can be caused by slow links, bad connections, dropped packets, or an overloaded router.

Chapter: 3

---

**QUESTION 664:**

You are a technician at Certkiller . Your newly appointed Certkiller trainee wants to know what the most effective method is to reduce large routing tables. What would your reply be?

- A. Route filters
- B. CIDR summary blocks
- C. Compression
- D. Switching

Answer: B

Explanation:

The only valid choice is to use CIDR (Classless Interdomain Routing) summaries whenever possible. With CIDR, masks in address space are grouped together to form one update. The lack of midsize address space (between a Class C and Class B), the growth of the Internet routing tables, and the eventual exhaustion of the 32-bit IP address space all fueled the creation of CIDR.

Chapter: 3

---

**QUESTION 665:**

You are a technician at Certkiller . You need to create filter for EIGRP network 198.30.64.0/24.

What mask should be used to create the filter?

- A. 0.0.0.255
- B. 0.0.255.255
- C. 255.255.1.0
- D. 255.255.254.0

Answer: A

Explanation:

A distribution list is used to block (filter) routing updates or queries and relies on an access-list to match the criteria. The configuration commands are listed below to allow only route information for the 198.30.60.0/24 network:

```
!  
router eigrp 1
```

```
network 198.30.64.0
distribute-list 2 in
!
access-list 2 permit 198.30.64.0 0.0.0.255
Chapter: 3
```

---

**QUESTION 666:**

Which of the following best describes a Type 3 or 4 LSA?

- A. Summary link entry which shows routing tables as IA for OSPF inter-area
- B. Autonomous System External type-1 which shows in routing tables as 'E1' for OSPF External-1
- C. Network link entry that shows in routing tables as '0' for OSPF
- D. Router link entry/record which show in routing tables as '0' for OSPF

Answer: A

Explanation:

Type 3 and 4 LSA's are generated by ABR's and sent to all routers within an area. These type LSA's advertise intra-area routes to the backbone area and both intra-area and inter-area routes to non-backbone areas.

Chapter: 3

---

**QUESTION 667:**

You are a technician at Certkiller . Your newly appointed Certkiller trainee wants to know how long OSPF will wait after receiving no updates before it sends a new LSA.

What would your reply be?

- A. 10 minutes
- B. 20 minutes
- C. 30 minutes
- D. 1 hour

Answer: C

Explanation:

The fixed OSPF constant "Link State Refresh" is defined as the maximum amount of time between distinct origination of the same LS

A. When the LSA age reaches this interval, the router must originate a new instance of the same LSA, keeping everything the same. The value of this constant is 30 minutes.

Chapter: 3

---

**QUESTION 668:**

By default, how does BGP handle IGP routes?

- A. BGP automatically redistributes all IGP routes
- B. BGP advertisements are independent of IGP route propagation
- C. BGP must wait until the IGP has propagated routing information across the autonomous system
- D. BGP can immediately advertise routes without waiting until the IGP has propagated information across the autonomous system

Answer: C

Explanation:

BGP synchronization rules dictate when traffic from a different AS passes through an AS to a third AS, BGP will not advertise the route until all routers within the AS have learned of the route through IGP.

Chapter: 3

---

**QUESTION 669:**

What methods does EIGRP use to determine if a previous neighbor is dead?

- A. Unicast
- B. Hold Time
- C. Multicast
- D. Broadcast

Answer: B, C

Explanation:

EIGRP multicasts hello packets to all neighbors at regular intervals (5 or 60 seconds). Every neighbor must send these hello packets. By default, if three consecutive hello packets are not received, the hold time expires and the neighbor is declared dead.

Chapter: 3

---

**QUESTION 670:**

Select the valid subnet number and broadcast address combinations for mask /26:

- A. Subnet 10.0.0.128, broadcast 10.0.0.192
- B. Subnet 10.0.0.64, broadcast 10.0.0.127
- C. Subnet 10.0.0.128, broadcast 10.0.0.191
- D. Subnet 10.0.0.0, broadcast 10.255.255.255

Answer: B, C

Explanation:

If 26 bits are used for the network portion of the address, then 6 bits are used in the mask portion.

$2 \times 2 \times 2 \times 2 \times 2 \times 2 = 64$  addresses per subnet. The broadcast address is always the all 1's address or the last address within the subnet.

Subnet 10.0.0.0 broadcast 10.0.0.63

Subnet 10.0.0.64 broadcast 10.0.0.127

Subnet 10.0.0.128 broadcast 10.0.0.191

Subnet 10.0.0.192 broadcast 10.0.0.255

Chapter: 3

---

**QUESTION 671:**

Policy-Based routing has which of these criteria?

- A. Can Deny based on subnet mask
- B. Can Deny based on RIP version
- C. Can Match based on the IP address
- D. Can Match based on Next Hop address

Answer: C

Explanation:

Policy-based routing relies on standard or extended access list to define the traffic matching criteria. Using standard or extended access lists, we cannot match subnet mask, RIP version, or next hop address. IP addresses can easily be matched using access lists.

Chapter: 3

---

**QUESTION 672:**

Which of the following BGP Path Attributes (metrics) is a proprietary enhancement?

- A. Community
- B. Origin
- C. Aggregator
- D. Weight

Answer: D

Explanation:

The weight attribute is a Cisco proprietary attribute used for path selection. This allows as administrator to "prefer" one path over one or more paths to the same destination.

Chapter: 3

---

**QUESTION 673:**

Which of the following are needed for an OSPF to form an adjacency?

- A. Autonomous System Number
- B. Area ID
- C. Hello and Dead intervals
- D. Stub Area Flag

Answer: C

Explanation:

In order to form an adjacency hello packets are exchanged. All routers connected to a common network must agree on all the parameters from the hello packet listed below:

Network Mask

Hello Interval

Dead Interval

Neighbor

Chapter: 3

---

**QUESTION 674:**

Which of the following correctly describes an OSPF Area Border Router?

- A. It is used to connect various OSPF Areas
- B. It is used only at the outside edge of an OSPF network for protocol translation
- C. It is used only internally for OSPF to EIGRP redistribution
- D. It is used for routing protocol redistribution

Answer: A

Explanation:

An ABR is defined as having multiple area assignments and two or more interfaces in two or more areas. An ABR would be used to connect a different, stub or totally stubby area to the rest of the network.

Chapter: 3

---

**QUESTION 675:**

You are a technician at Certkiller . You are troubleshooting a problem on the Certkiller network. You issue the debug ip ospf adj command. Your newly appointed Certkiller trainee wants to know what the purpose of this command is. What would your reply be?

- A. LSA type 1
- B. Elections

- C. LSA type 3 and 4
- D. OSPF neighbor relationships

Answer: D

Explanation:

The output from this command can be very useful in troubleshooting neighbor relationships. A sample output from the command is listed below:

```
Router#debug ip ospf adj
OSPF: Receive dbd from 70.70.70.70 seq 0x14B
OSPF: 2 Way Communication to neighbor 70.70.70.70
OSPF: send DBD packet to 192.16.64.2 seq 0x1797
OSPF: Receive dbd from 70.70.70.70 seq 0x1797
OSPF: NBR Negotiation Done We are the MASTER
OSPF: send DBD packet to 192.16.64.2 seq 0x1798
OSPF: Database request to 70.70.70.70
OSPF: sent LS REQ packet to 192.16.64.2, length 12
OSPF: Receive dbd from 70.70.70.70 seq 0x1798
OSPF: send DBD packet to 192.16.64.2 seq 0x1799
OSPF: Receive dbd from 70.70.70.70 seq 0x1799
OSPF: Exchange Done with neighbor 70.70.70.70
OSPF: Synchronized with neighbor 70.70.70.70, state:FULL
OSPF: Build router LSA, router ID 172.16.13.1
Chapter: 3]
```

---

**QUESTION 676:**

What are BGP routers with the same AS number exchange information known as?

- A. BGP
- B. eBGP
- C. IGRP
- D. iBGP

Answer: D

Explanation:

Internal BGP (iBGP) sessions are contained within the autonomous system. The active session is between routers with the same AS number.

External BGP (eBGP) sessions are connections between two different autonomous systems. The active session is between routers with different AS numbers.

Chapter: 3

---

**QUESTION 677:**

BGP falls into which routing protocol category below?

- A. Link-state protocol
- B. Interior gateway protocol
- C. Exterior gateway protocol
- D. Distance-Path state protocol

Answer: C

Explanation:

BGP is an EGP (Exterior Gateway Protocol used to connect and find routes to/from autonomous systems.

Chapter: 3

---

**QUESTION 678:**

Which of the following is untrue of routers running internal BGP?

- A. Routers are running iBGP when two BGP speakers are in different AS's
- B. They must form a TCP port 179 connection with each other
- C. They do not have to be directly connected
- D. Routers are running iBGP when two BGP speaking routers are in the same AS

Answer: A

Explanation:

Internal BGP (iBGP) sessions are contained within the autonomous system. The active session is between routers with the same AS number.

External BGP (eBGP) sessions are connections between two different autonomous systems. The active session is between routers with different AS numbers.

Chapter: 3

---

**QUESTION 679:**

Which of the following are true of a network design with a high number of OSPF routers in a single area?

- A. Increased problems with routing table 'black holes'
- B. Increased SPF algorithm re-calculation times
- C. Decreased OSPF update traffic overhead
- D. Increased OSPF update traffic overhead

Answer: B, D

Explanation:

SPF path re-calculation times are increased as more routers (and paths) are added to the network. In addition, as more routers and paths are added additional LSA's must be

generated for each network and summary that exists within the area.

Chapter: 3

---

**QUESTION 680:**

Which of these values is required by EIGRP as a Seed Metric when performing Route Redistribution?

- A. Bandwidth, Delay, Reliability, Load, MTU
- B. Cost, Delay, Reliability
- C. Route-Tag
- D. Cost

Answer: A

Explanation:

When redistributing routes into IGRP or EIGRP, the bandwidth, delay, reliability, load and MTU (maximum transmission unit) must be specified.

!

```
router eigrp 91
network 10.10.108.0
redistribute static
redistribute rip
default-metric 10000 100 255 1 1500
```

Chapter: 3

---

**QUESTION 681:**

With regard to an OSPF internal router, which of the following statements are true?

- A. Internal routers are in more than one area
- B. At least one interface is in Area 0
- C. All interfaces are in the same Area
- D. Internal routers redistribute routing information with routers in other areas

Answer: C

Explanation:

An internal OSPF router is defined by having all its interfaces members of the same area.

Chapter: 3

---

**QUESTION 682:**

What is the inverted wildcard mask of /26?

- A. 63.255.255.255
- B. 0.0.0.63
- C. 192.255.255.255
- D. 0.0.0.192

Answer: B

Explanation:

The regular dotted decimal format of a /26 prefix is represented as: 255.255.255.192 or 11111111.11111111.11111111.11000000.

To get the wildcard mask, use  $256-192=64$ , which is the block size of the subnet mask.

The wildcard mask is always one less than the block size or 63 in this case. A wildcard of zero (0) means all network bits are on. So the answer is 0.0.0.63.

Chapter: 3

---

**QUESTION 683:**

You are the network administrator at Certkiller . You attempt to start an OSPF process on Router CK1 but you receive an 'Unable to allocate router ID' error. What could be the cause of this problem?

- A. The OSPF area is Stubby
- B. No IP address assigned to any interface
- C. Router ID is used by another router
- D. All interfaces are shut down

Answer: B, D

Explanation:

When OSPF is enabled on a router, the loopback interface IP address is used to assign the Router ID by default. If the loopback interface is not configured, then the highest IP address of all interfaces is assigned as the Router ID. At least one interface must be active/enabled and configured with an IP address to enable OSPF.

Chapter: 3

---

**QUESTION 684:**

Which types of LSA's are passed into a Totally Stubby Area?

- A. Summary Type 3
- B. Summary Type 4
- C. Router Link Advertisement Type 1
- D. External Link Advertisement Type 5

Answer: C

Explanation:

Only possible answer might be C. Type 3, 4 and 5 are all blocked from going into a Totally stubby area. So only intra-area and default.

Source: BSCI Self Study Cisco Press (2003) 2nd Edition, P190

---

**QUESTION 685:**

With regard to Route Reflectors, which of the following are true?

- A. They increase the size of the BGP routing table
- B. They enable the use of Route Summarization
- C. They reduce the number of TCP sessions
- D. A route reflector cannot propagate iBGP routes to iBGP peers.

Answer: C

Explanation:

One of the major reasons route reflectors are deployed is to relieve the iBGP full-mesh requirement. The BGP split-horizon rule is modified by route reflectors. It is modified by allowing a route reflector to be the only router that propagates routes learned by iBGP to other iBGP peers.

Chapter: 3

---

**QUESTION 686:**

Among these methods of sending an IGP route into BGP, which is not recommended?

- A. Redistribute the IGP into BGP
- B. Use the network command with the 'mask' option
- C. Redistribute BGP into the IGP
- D. Use the network command without the 'mask' option

Answer: A

Explanation:

Redistribution can cause routing loops and route flapping. In addition, BGP actually pulls information that other IGP's have learned about their environment. BGP handles the translation of one protocol to another if multiple routing protocols are used in the AS.

Chapter: 3

---

**QUESTION 687:**

Which of the following are reasons a network administrator would choose to run multiple routing protocols at the same time on the same network?

- A. Reduce routing protocol overhead on the network
- B. Optimize the route redistribution between areas
- C. Convert from an older routing protocol to a more modern one
- D. Acquisition of a company running a different routing protocol than your own

Answer: C

Explanation:

In general, multiple routing protocols are not needed simultaneously on the same network. However, when migrating to a new routing protocol, this is a good methodology as migration configurations can be tested without affecting existing traffic.

Chapter: 3

---

**QUESTION 688:**

You are a technician at Certkiller . Certkiller has a BGP network. Your newly appointed Certkiller trainee wants to know when BGP is not suitable for use. What would your reply be?

- A. Routing policy and route selection are not important for your AS
- B. A single connection to the Internet
- C. When you need to perform load-sharing to multiple ISP's simultaneously
- D. Low bandwidth between your AS and the other AS

Answer: A, C, D

Explanation:

"B" " A single connection to the Internet" is suitable for BGP. According to this Question, "ACD " is correct , because these are unsuitable responses to this Question.

---

**QUESTION 689:**

Which command can be used to view the state of the link, such as exstart, exchange, or full?

- A. show ospf neighbor
- B. show ip protocols
- C. show ip ospf neighbor
- D. show ospf interface

Answer: C

Explanation:

Answer D lists the state of the DR on the interface. Answer C shows all information about all OSPF neighbor routers including the state.

router2#show ip ospf neighbor

```
Neighbor ID Pri State Dead Time Address Interface  
170.170.5.1 1 INIT/- 00:00:34 170.170.1.1 Serial0  
router-2#  
Chapter: 3
```

---

**QUESTION 690:**

Which routing protocol features provide solutions to the problems associated with very large routing tables?

- A. Filtering
- B. Data Compression
- C. Route summarization
- D. Incremental updates

Answer: C, D

Explanation:

Route summarization reduces the number of entries in the routing tables. Incremental updates only send information about the network topology that has changed. Both of these features free up resources on routers, and bandwidth on the network.

Chapter: 3

---

**QUESTION 691:**

What is the router command to diagnose and display the entire route, including delays, given the path to a destination?

- A. routepath
- B. routetrace
- C. pathtrace
- D. traceroute

Answer: D

Explanation:

The traceroute command provides each hop and delay information about the path to a destination.

Chapter: 3

---

**QUESTION 692:**

You are the network administrator at Certkiller . You configure a static route on router running EIGRP. Your newly appointed Certkiller trainee wants to know what the default administrative distance of this static route would be. What would your reply be?

- A. 255
- B. 5
- C. 1
- D. 90

Answer: C

Explanation:

Static routes are assigned a default administrative distance of 1.

Chapter: 3

---

**QUESTION 693:**

With regard to an OSPF backbone router, which of the following statements is true?

- A. It redistributes routing information with routers in other areas
- B. All backbone routers must be in more than one area
- C. At least one interface is in Area 0
- D. All interfaces are in the same Area

Answer: C

Explanation:

A backbone router has at least one interface in the backbone area.

Chapter: 3

---

**QUESTION 694:**

Which of these packet types does an OSPF router use at bootup to receive information?

- A. SPF
- B. LCP
- C. Flood
- D. Hello

Answer: D

Explanation:

OSPF uses hello packets to build adjacencies when OSPF is enabled on a router or when a router boots up with OSPF already configured.

Chapter: 3

---

**QUESTION 695:**

You are a technician at Certkiller . You reboot the existing DR. This results in the BDR being promoted to DR.

With regard to this scenario, which of the following statements are true?

- A. Upon boot of the old DR, the newly elected DR will automatically demote itself back to BDR
- B. When the old DR comes back online, it will send out an LSA to override the new DR
- C. The newly elected OSPF DR will remain DR even when the old DR comes back on-line
- D. Upon boot of the old DR another election will occur

Answer: C

Explanation:

Once a DR and BDR have been elected, the DR/BDR is sent to all routers via hello packets. In this case, the existing DR is rebooted and the BDR is promoted to DR and a new BDR is elected. When the previous DR comes back on-line, it accepts the newly promoted/elected DR/BDR that is received in the hello packets.

Chapter: 3

---

**QUESTION 696:**

Which of the following supports multiple routed protocols?

- A. IP
- B. OSPF
- C. BGP
- D. EIGRP
- E. IGRP

Answer: D

Explanation:

One of the enhancements to EIGRP from IGRP is its support for all major Layer 3 routed protocols. EIGRP supports IP, IPX, and AppleTalk. All of the other routing protocols listed provide support for routing IP only.

Chapter: 3

---

**QUESTION 697:**

You are a technician at Certkiller . Your newly appointed Certkiller trainee wants to know what the BGP MED attribute is used for.

What would your reply be?

- A. Setting the route maps peer ID
- B. Setting IGP synchronization
- C. Setting the route reflectors peer ID
- D. Setting a preferred return-pathway back into the originating AS

Answer: D

Explanation:

The MED attribute in BGP is used to indicate the best entry point or path to reach a particular destination to a neighboring AS. A lower MED is preferred over a higher MED. The MED attribute is also non-transitive because the AS does not pass the MEDs it learns from one AS to another.

Chapter: 3

---

**QUESTION 698:**

With regard to RIPv1, which of the following statements are true? (Choose all that apply.)

- A. Maximum hop count is 16
- B. Update interval is 90 seconds
- C. Update interval is 30 seconds
- D. VLSM is not supported
- E. Maximum hop count is 15

Answer: C, D, E

Explanation:

Answer A is tricky as the maximum allowable hop count is 15. If the hop count is set to 16 (infinity), the destination is considered unreachable. The RIP periodic update interval is 30 seconds.

Chapter: 3

---

**QUESTION 699:**

How does the command 'ip bandwidth-percent eigrp' set EIGRP maximum bandwidth allocation?

- A. On a per AS basis
- B. On a per packet basis
- C. On a per interface basis
- D. Globally for all IGRP and EIGRP AS's on the router

Answer: C

Explanation:

This command is applied to an interface to set the maximum amount of bandwidth to be used on an interface for EIGRP traffic.

Chapter: 3

---

**QUESTION 700:**

Which of the following are possible reasons that EIGRP might not make an adjacency to a neighbor router?

- A. Both routers are not running DUAL
- B. Hop counts do not match on both routers
- C. The feasible successor is still in the Hello state
- D. Update packets have not been sent

Answer: D

Explanation:

When EIGRP begins making adjacencies, it multicasts hello packets out an interface(s). The neighbors become adjacent when they acknowledge each others hello packets and their K values match.

Chapter: 3

---

**QUESTION 701:**

With regard to OSPF, which of the following statements are correct? (Choose all that apply)

- A. OSPF computes cost based on the interface's bandwidth setting
- B. OSPF is not limited to a hop count metric
- C. OSPF uses LSA packets instead of broadcasts
- D. OSPF sends the complete routing table inside of each LSA packet

Answer: A, B, C

Explanation:

OSPF does not send its complete routing table inside each LS

A. An LSA is sent for each network that is included in the routing process.

Chapter: 3

---

**QUESTION 702:**

All BGP routers belong to which of these communities by default?

- A. no-export
- B. classless

- C. Internet
- D. no-advertised

Answer: C

Explanation:

All routers by default belong to the "Internet" community and can be used to advertise routes to all other routers. No-export tags the routes so that it will not be sent outside the AS. The No-advertise tag will prevent the route from being advertised to any other BGP router. Classless is not a valid BGP community.

Chapter: 3

---

**QUESTION 703:**

Which Administrative Distances listed below are correct? (Choose all that apply)

- A. BGP=20
- B. Static Route=1
- C. External IGRP=170
- D. EIGRP=120
- E. Summary EIGRP=5

Answer: A, B, E

Explanation:

Administrative distances are used to rate the reliability or trustworthiness of a route. Connected interfaces are assigned a value of 0, and static routes assigned a value of 1. These two are the most "trusted" routes, while an administrative distance of 255 indicates the route will never be used. Different routing protocols are assigned different values and can have different values for different types of routes for each protocol (internal, external and summary). Choice C is incorrect because IGRP does not distinguish between internal and external routes. EIGRP is assigned a default administrative distance of 90 not 120 as specified in choice D.

Chapter: 3

---

**QUESTION 704:**

With regard to RIPv1, which of the following statements are true? (Choose all that apply.)

- A. rip v.1 route updates include a subnet mask field.
- B. rip v.1 is considered a "classful" routing protocol.
- C. rip v.1 is considered a "classless" routing protocol.
- D. rip v.1 route updates do not include a subnet mask field.

Answer: B, D

Explanation:

Ok, set the "wayback" machine to the first day of networking class. On that day, we learned some interesting things about routing protocols and their behaviors. We know RIP v.1 is a "distance vector" routing protocol that relies on "hop count" for pathing decisions. We also know that RIP v.1 is a "classful" routing protocol, which means that it "observes" address class boundaries. When you advertise a network, using the "network" command, you are not given an opportunity to enter a subnet mask. RIP is going to assume the default mask for that address class (255.0.0.0 for a Class "A", 255.255.0.0 for a Class "B" and so on). RIP does NOT include subnet masks in its routing updates. Only a classless routing protocol like OSPF would need to furnish that kind of information.

```
Router(config)#router rip
Router(config-router)#network 192.168.1.10
Router(config-router)#network 192.168.2.10
Router(config-router)#network 192.168.3.10
Router(config-router)#^Z
```

Router#

Chapter: 4

---

**QUESTION 705:**

Which two routing protocols will most likely be used in a network with limited scalability in mind?

- A. IGRP
- B. OSPF
- C. EIGRP
- D. RIPv2
- E. RIPv1

Answer: A, E

Explanation:

This should be a relatively easy one to answer since there are only two "classful" routing protocols in the list of possible answer choices. Due to the fact that RIP v1 summarizes all known routes along classful network boundaries, using VLSM (Variable Length Subnet Mask) to extend your addressing scheme is out of the question. RIP limits network scalability thanks to its classful nature and its need to send updates every 30 seconds. Imagine how long it would take a network of 200+ routers to converge with RIP sending updates every 30 seconds.

Chapter: 4

---

**QUESTION 706:**

You are a technician at Certkiller . Certkiller has an OSPF network. Your newly appointed Certkiller trainee wants to know what he benefits of running a

"link-state" protocol like OSPF are.

What would your reply be? (Choose all that apply.)

- A. link-state protocols maintain a central network topology database
- B. link-state protocols react quickly to topology changes.
- C. link-state protocol updates are sent out every 30 seconds regardless of topology change
- D. link-state protocols use very little bandwidth in a stable network.

Answer: A, B, D

Explanation:

Link-state routing protocols like OSPF maintain a central network topology database that is built during convergence. They also react quickly to topology changes. If a router's interface goes down, that information is sent out at once in the form of an LSA (Link-State Advertisement), which informs the other routers of the change. Link-state routing protocols like OSPF (starting to notice a theme here?) use very little bandwidth as their routing updates are "incremental" in nature and only advertise which part of the network has changed. This behavior is the opposite of RIP, which sends a complete copy of its routing table with every update.

Chapter: 4

---

**QUESTION 707:**

The BGP routing protocol maintain which two types of tables?

- A. IP
- B. BGP topology
- C. BGP attributes
- D. BGP information sent to and received from other BGP routers
- E. shared table that combines UDP and BGP route information

Answer: A, D

Explanation:

The BGP routing protocol maintains two sets of tables, an IP table and a table composed of BGP information sent to and received from other BGP routers. The information in these tables is by default maintained separately, however you can configure the router to share the information if you so desire.

Chapter: 4

---

**QUESTION 708:**

If you want a permanently created static route entry picked up by the routing process, what command would you use?

- A. static insert
- B. static announce perm
- C. static redistribute
- D. redistribute static

Answer: D

Explanation:

When you create a static route that is not included in the list of the networks advertised by the IOS "network" command, it will not be picked up/learned/injected into the routing protocol unless you specifically advertise it. This feat is accomplished through the use of the "redistribute static" command.

Chapter: 4

---

**QUESTION 709:**

BGP peers communicate via which method?

- A. RIP
- B. TCP
- C. UDP
- D. ICMP

Answer: B

Explanation:

BGP peers communicate via TCP port 179. The peers initiate a TCP session in which they exchange a series of "OPEN" messages, which ultimately forms the connection. The newly formed connections are kept active by exchanging "KEEPALIVE" messages. Once all this has been completed, the peers exchange routing table, and periodic incremental updates.

Chapter: 4

---

**QUESTION 710:**

What BGP is considered to be?

- A. both an internal and external routing protocol
- B. an internal routing protocol only
- C. an external routing protocol only
- D. an independent routing protocol

Answer: A

Explanation:

BGP is truly a magical thing! It runs as both an internal routing protocol (iBGP) within

the AS and an external routing protocol (eBGP) between ASes. When BGP is used internally (iBGP) its routing duties do not leave the confines of the AS. iBGP learns of new routes within the AS and passes these updates to eBGP peers for distribution. eBGP learns and distributes routes between different ASes.

Chapter: 4

---

**QUESTION 711:**

What does BGP "ingress filtering" allow you to do?

- A. allows filtering of data packets based on protocol type
- B. allows you to decide which routes will be accepted from peers or neighbors
- C. allows you to decide which routes will be advertised to peers or neighbors
- D. allows filtering of data packets based on protocol number

Answer: B

Explanation:

AAPT has deployed "ingress filtering" on all customer, peer and provider links. This means that filters will be installed to prevent traffic sourced from customer networks not registered for routing with AAPT from entering the AAPT network and will thus prevent this traffic from entering the networks of our peers and providers via AAPT

---

**QUESTION 712:**

BGP communities perform what function?

- A. they group routes into a single administrative control group
- B. they delete routes based on administrative control group names
- C. they restrict routes based on network class
- D. they allow routers to filter incoming or outgoing BGP routes

Answer: D

Explanation:

BGP communities allow a router a more effective way to filter BGP route updates than would be possible using distribute lists and prefix lists. A BGP community is a tag that is applied to a route update indicating a destination route (or other common attribute) that is shared with other routes. The router can therefore make pathing decisions based on the commonality of a group of updates as opposed to individually assessing each route.

Chapter: 4

---

**QUESTION 713:**

You are the network administrator at Certkiller . You are logged into the ASBR in Area 0 (backbone). You want to advertise the address class 192.168.1.0.

Which of the following commands will make that route available to routers outside the area?

- A. network 192.168.1.0 0.0.0.255 area 0
- B. network 192.168.1.0 255.255.255.0 area 0
- C. adv address 192.168.1.0 255.255.255.0
- D. adv address 192.168.1.0 0.0.0.255

Answer: A

Explanation:

This should have been another easy one to answer. Since the "adv address" command is not used to advertise networks, you should have been able to exclude the answers that used it. That left us with the two commands containing the "network" command. One of the network commands uses a standard subnet mask "255.255.255.0" and one uses a "reverse" or "wildcard" mask. The "network" command uses the "wildcard" mask to indicate how far into the class it should advertise. The correct answer is "network 192.168.1.0 0.0.0.255 area 0"

Chapter: 4

---

**QUESTION 714:**

Which of the following is NOT a type of OSPF Area?

- A. Normal
- B. Stub
- C. Totally Stub
- D. Not-So-Stubby
- E. Not-So-Normal

Answer: E

Explanation:

There are only four OSPF area types (according to Cisco) and they are stub, totally stub, not-so-stubby and normal. The answer "Not-So-Normal" is the made up one.

Chapter: 4

---

**QUESTION 715:**

With regard to VLSM, which of the following statements are true? (Choose all that apply.)

- A. the area number must be carried with the routing update
- B. the addresses being summarized must be classful in nature
- C. the subnet mask is carried with the routing update
- D. addresses being summarized must contain the same "high-order" bits

Answer: C, D

Explanation:

Variable Length Subnet Masking or VLSM gives the router the ability to announce a series of addresses with a single routing table entry. Pretty neat trick huh?

The way that it does this, is it looks at the addresses and ascertains which of those addresses are similar out to a designated bit pattern, also known as "high-order bits". It then looks at the subnet mask accompanying the update and decides which series of addresses are considered matches. You could potentially advertise as many as five destination addresses with a single routing table entry.

Chapter: 4

---

**QUESTION 716:**

You are a technician at Certkiller . You are designing a multi-area OSPF network. What two things must you do?

- A. make sure that each area connects to the backbone (Area 0).
- B. assign each area a unique number
- C. configure Areas 1-4 as Stub Areas
- D. configure Areas 1-4 as Not-So-Stubby Areas

Answer: A, B

Explanation:

Creating a multi-area OSPF network requires a few design considerations such as, all areas must have a connection to Area 0 so that routing updates will have a central distribution point, all created areas must have a unique area number (that one is kind of obvious) to prevent routing updates from being confused by the fact that two areas are using the same number.

Chapter: 4

---

**QUESTION 717:**

What function does the "area range" command perform?

- A. instructs the ABR to summarize the routes for a specific area before advertising them
- B. instructs the ABR to segregate the routes for a specific area before advertising them
- C. specifies a range of areas to advertise exclusively
- D. specifies a range of areas to observe and report the routing status

Answer: A

Explanation:

The "area range" command instructs the ABR to summarize the addresses within the area

before sending them out in a routing update.

Chapter: 4

---

**QUESTION 718:**

You are a technician at Certkiller . You want to check the status of configured virtual links.

Which command should you issue?

- A. show ospf virtual-links
- B. show virtual-links ospf
- C. show ip ospf virtual-links
- D. show ip ospf v-links

Answer: C

Explanation:

The command "show ip ospf virtual-links" will show you information about configured virtual-links.

```
new-york router#show ip ospf virtual-links Interface address: 192.168.1.10 (POS 1/1/1)
cost: 1, state: P To P, transit area: 1.2.3.4 hello: 10, dead: 40, retrans: 5 nbr id:
```

```
192.168.1.20, nbr address: 192.168.1.20 nbr state: Full, nbr mode: Master, last hello: 38
```

Chapter: 4

---

**QUESTION 719:**

An OSPF router must be in what state in order to route traffic?

- A. mixed
- B. partial
- C. two-way
- D. full

Answer: D

Explanation:

An OSPF router must be in "full" state (down, attempt, init, 2-way, exstart, exchange, loading, or full). When a router is in "full" state, adjacencies are fully established and the router is ready to route traffic.

Chapter: 4

---

**QUESTION 720:**

You are the network administrator at Certkiller .Your newly appointed Certkiller network engineer has not quite grasped the concept of classless routing.

How would you explain the concept? (Choose all that apply.)

- A. classless routing protocols carry the subnet mask with updates
- B. classful routing protocols carry the subnet mask with updates
- C. classless routing enables the use of VLSM
- D. classless routing enables the use of high-order routing

Answer: A, C

Explanation:

Classless routing is a pretty complex topic for new engineers to fully grasp. The best way to explain the basic principles is to re-enforce the fact that the classless routing is not bound by address class limitations, routing updates carry the specialized subnet mask which makes VLSM possible.

Chapter: 4

---

**QUESTION 721:**

The "ip bandwidth-percentage eigrp" performs what function?

- A. it dictates what the maximum bandwidth percentage that EIGRP packets are allowed to use on a single router interface.
- B. it dictates what the maximum bandwidth percentage that IP packets are allowed to use on a single router interface
- C. it dictates what the maximum bandwidth percentage that EIGRP packets are allowed to use on all router interfaces.
- D. it dictates what the maximum bandwidth percentage that packets routed via EIGRP are allowed to use on a single router interface.

Answer: A

Explanation:

The ip bandwidth-percent eigrp command is used to configure the percentage of bandwidth that may be used by EIGRP on an interface. If you issue the "ip bandwidth-percentage eigrp" command at the command line, the router will only allow 50% of the interface's bandwidth (which is the default) to be used for EIGRP.

Chapter: 4

---

**QUESTION 722:**

What is the cost between the next hop router and the destination?

- A. reliable distance
- B. calculated distance
- C. reasonable distance
- D. advertised distance
- E. administrative distance

Answer: D

Explanation:

The Diffusing Update Algorithm (DUAL) is the "brains" behind the EIGRP routing protocol's path decision making process. When routes to a destination network go down, the DUAL algorithm calculates new paths to make sure that traffic gets to its intended destination. One of the metrics involved in this decision making process is "advertised distance" which is equal to the cost between the next hop router and the destination network.

Chapter: 4

---

**QUESTION 723:**

What is the cost between the local router and the next hop router?

- A. feasible distance
- B. partial distance
- C. next-hop distance
- D. agregate distance
- E. advertised distance

Answer: C

Explanation: Only possible answer might be C.

Not A: Feasible distance is the cost of the route from the current router to the remote network.

Source BSCI Self Study Cisco Press (2003) 2nd Edition P362-363.

---

**QUESTION 724:**

In order to determine the "best" path to a destination network, you need to add which two path costs together? (Select two.)

- A. feasible distance
- B. partial distance
- C. advertised distance.
- D. agregate distance
- E. next hop distance

Answer: C, E

Explanation:

Two path costs would be advertised distance plus the next hop distance.

Source BSCI Self Study Cisco Press (2003) 2nd Edition P362-363.

---

**QUESTION 725:**

What could cause EIGRP routes to be considered "stuck-in-active"?

- A. the route being reported by the "stuck-in-active" message no longer exists.
- B. the route being reported by the "stuck-in-active" message has been activated.
- C. a query for the route generated by an EIGRP neighbor has not yet been replied to.
- D. a query for the route generated by an OSPF neighbor has not yet been replied to.

Answer: A, C

Explanation:

The "stuck-in-active" message refers to the condition in which a primary route is no longer available and no feasible successor is available. The EIGRP router has sent out a query to its neighbors and no reply has been heard back for three minutes.

Chapter: 4

---

**QUESTION 726:**

What are Query, Update and Reply known as?

- A. EIGRP renewable packets
- B. EIGRP reliable packets
- C. TCP/IP reliable packets
- D. QueryAll packets
- E. OSPF reliable packets

Answer: B

Explanation:

EIGRP utilizes five different types of packets:

- 1) ACK
- 2) hello
- 3) query
- 4) reply
- 5) update

These packets are instrumental in EIGRP's ability to reliably route packets from source to destination network. The "reliable packets" are query, update, and reply.

Chapter: 4

---

**QUESTION 727:**

You are a technician at Certkiller . Certkiller has an IS-IS network. Your newly appointed Certkiller trainee wants to know what IS-IS stand for.

What would your reply be?

- A. intercontinental station-to-intercontinental station
- B. intermediate section-to-intermediate section
- C. intermediate station-to-intermediate station
- D. intermediate system-to-intermediate system

Answer: D

Explanation:

The intermediate system-to-intermediate system routing protocol, more commonly known as IS-IS was developed by Digital Equipment Corporation.

Chapter: 4

---

**QUESTION 728:**

Which of the following attributes are common to both OSPF and IS-IS?

- A. they both maintain a link-state database.
- B. they both run the SPF algorithm to determine the shortest path to destination networks
- C. they both utilize the concept of "areas" to maintain hierarchical network topologies
- D. they both use RIP as their data transfer medium
- E. they both communicate via port 80 when sending routing updates

Answer: A, B, C

Explanation:

IS-IS and OSPF share a number of common features; link-state databases, the running of the SPF algorithm, the use of "areas" among other similarities.

Chapter: 4

---

**QUESTION 729:**

Which of the following subnet masks is the most optimized for providing five internet addresses?

- A. 225.255.255.192
- B. 225.255.255.248
- C. 225.255.255.255
- D. 225.255.255.254

Answer: B

Explanation:

Do the math and win a prize! This should have taken you about 5 minutes to answer. The subnet mask 255.255.255.248 will give you 32 subnets with 6 hosts per subnet. Sounds pretty efficient to me. If you want a great shortcut for figuring out number of hosts per subnet, subtract 248 from 256, that leaves you with 8 right?

Subtract two from 8 (can't use the first and last addresses in the range) and you are left with 6. There are six hosts in each of the subnets. It works every time.

The other subnet masks:

255.255.255.192 = 4 subnets with 62 hosts per subnet, way too many hosts.

255.255.255.255 = Gong! I hope no one fell for this one

255.255.255.254 = 128 subnets with two hosts per subnet, won't work if you need five hosts. (\*TIP\* use this for point to point serial connections)

Chapter: 4

---

**QUESTION 730:**

On a point-to-point network, what acknowledges each LSP that it receives.

- A. PSNP
- B. POPN
- C. BPDU
- D. LMNOP

Answer: A

Explanation:

On a point-to-point IS-IS network, an LSP(Link State Packet) is generated to announce a router's links and the status of the those links. When the LSP is received, a PSNP (Partial Sequence Number Packet) is sent to acknowledge each of the LSPs that are received.

Chapter: 4

---

**QUESTION 731:**

What is the IS-IS equivalent of the OSPF backbone Area 0?

- A. level-0 backbone
- B. area 0 backbone
- C. level-1 backbone
- D. level-2 backbone

Answer: D

Explanation:

In order to answer this question, we are only going to deal with Level-1 and Level-2 routers (I see the guy with his hand raised...Yes, there are Level1/2 routers in IS-IS as well). The Level-1 routers are similar to OSPF areas; they are independent of each other and require a concentration point to share routing updates.

It is in this capacity that the Level-2 backbone functions. Level-2 routers provide the common connection point through which route updates from the other Level-1 routers flow. This concept has been highly simplified, but with good reason...hint-hint.

Chapter: 4

**QUESTION 732:**

With which of the following can Level-2 IS routers establish an adjacency?

- A. Level-1 IS in the same area
- B. Level-2 IS in a different area
- C. Level-2 IS in the same area
- D. Level-1 IS in a different area

Answer: B, C

Explanation:

L2 router only establish adjacency with other L1/L2 or L2 routers - - never with a plain L1 router

---

**QUESTION 733:**

Which of the following are most like an OSPF Area Border Router (ABR)?

- A. Level 1 IS
- B. Level 2 IS
- C. Level1/Level2 IS
- D. Level2/Level3 IS
- E. IS-IS Media Border Router (IMBR)

Answer: C

Explanation:

IS-IS L1/L2 routers are a hybrid router type (which is the default setting when activating an IS-IS router). The L1/L2 router is most like the OSPF ABR because L1/L2 routers can connect to both the backbone (L2) and a router in a different area (L1), which as we know an ABR is also capable of doing.

Chapter: 4

---

**QUESTION 734:**

From the answer choices below, choose the statements that are true regarding IS-IS.

- A. L1 IS routers can be used to connect areas together
- B. L2 IS routers can be used to connect areas together
- C. IS-IS routers can either be an L1 or L2
- D. IS-IS routers can be both an L1 and L2 at the same time

Answer: B, D

Explanation:

IS-IS L1/L2 routers are a hybrid router type (which is the default setting when activating an IS-IS router). The L1/L2 router is most like the OSPF ABR because L1/L2 routers can connect to both the backbone (L2) and a router in a different area (L1).

Chapter: 4

---

**QUESTION 735:**

What will an OSPF ABR connect to when introduced into an existing network?

- A. one or more OSPF areas
- B. a single IS-IS area
- C. L2 IS router
- D. L1/L2 IS router

Answer: A

Explanation:

The OSPF Area Border Router's job is to sit on the border of one or more OSPF areas and provide a communication conduit between the areas and the backbone. The ABR runs the SPF algorithm and maintains the data for each area that it is connected to.

Chapter: 4

---

**QUESTION 736:**

Level 1/2 IS routers provide which service for IS-IS networks?

- A. intra-area routing services
- B. inter-area routing services
- C. intra-level routing services
- D. consolidation of OSPF and IS-IS routing updates

Answer: B

Explanation:

IS-IS L1/L2 routers are a hybrid router type (which is the default setting when activating an IS-IS router). The L1/L2 router is most like the OSPF ABR because L1/L2 routers can connect to both the backbone (L2) and a router in a different area (L1). Much in the same way that an OSPF ABR sits on the borders of one or more areas so does the L1/L2 router. Since the L1/L2 router straddles the area "fence", it is capable of providing inter-area routing services.

Chapter: 4

---

**QUESTION 737:**

OSPF can handle VLSM because it has what capability?

- A. OSPF organizes the network hierarchy when the SPF algorithm is run
- B. OSPF's nature as a link-state database automatically supports VLSM
- C. OSPF carries the subnet mask within its routing updates
- D. OSPF organizes areas and their ABRs into hierarchical groups

Answer: C

Explanation:

In order for VLSM to be deployed on a network infrastructure, you need to use a routing protocol that is capable of sending specialized subnet masks as part of the routing table updates. When you configure OSPF, one of the steps is to enter the network addresses with corresponding subnet masks.

Chapter: 4

---

**QUESTION 738:**

Which of the following are metrics available to IS-IS?

- A. delay
- B. error
- C. expense
- D. hop count

Answer: A, B, C

Explanation:

The IS-IS routing protocol incorporates three metrics: delay, expense, and error.

delay - measures the amount of delay on the link

expense - measures cost in resources of using the link

error - measures the number of errors occurring on the link

Chapter: 4

---

**QUESTION 739:**

What is the purpose of the "show isis route" command?

- A. It displays the Level-1 routing table for integrated IS-IS
- B. It displays the Level-2 routing table for integrated IS-IS
- C. It displays the routes most recently taken to a Level-1 ABR
- D. It displays the hops taken by a packet leaving the nearest Level-1 router

Answer: A

Explanation:

The "show isis route" command will show the Level-1 routing table as seen in the

example below:

```
Certkiller #show isis route
IS-IS Level-1 Routing Table - Version X
System Id Next-Hop SNPA Interface Metric State
BB00.0400.020C BB00.0400.020C bb00.0400.020c Serial0 10 Up
0800.2BB1.4434 0000.0000.0000 -- -- 0 Up
0800.2BB3.785B 0800.2BB3.785B bb00.0400.020c Serial0 10 Up
Chapter: 4
```

---

**QUESTION 740:**

You are a technician at Certkiller . Certkiller has an OSPF network. Your newly appointed Certkiller trainee wants to know what an OSPF router does when a link-state change occurs.  
What would your reply be?

- A. it shuts down and runs the SPF algorithm
- B. it multicasts the link-state change on ip address 224.0.0.6 to the DRs and BDRs via an LSU
- C. it multicasts the link-state change on ip address 255.255.255.224 to the DRs and BDRs via an LSU
- D. it unicasts the link-state change on ip address 255.255.255.224 to the ABRs and ASBRs via an LSU

Answer: B

Explanation:

An OSPF router reacts to link-state changes in the follow manner:

- 1) Link-state change is detected by a router
- 2) The router multicasts an LSU (Link-state update) to DRs and BDRs on address 224.0.0.6
- 3) The DR confirms receipt of the LSU and then floods it to other routers via address 224.0.0.5

There is more to the story, but this gives you the basic idea.

Chapter: 4

---

**QUESTION 741:**

You are a technician at Certkiller . Certkiller has an OSPF network. Your newly appointed Certkiller trainee wants to know OSPF which multi-cast addresses are used on the Certkiller network.

What would your reply be? (Choose all that apply.)

- A. 224.0.0.6
- B. 224.0.0.5
- C. 225.0.0.3

D. 255.255.255.224

Answer: A, B

Explanation:

An OSPF router reacts to link-state changes in the follow manner:

- 1) Link-state change is detected by a router
- 2) The router multicasts an LSU (Link-state update) to DRs and BDRs on address 224.0.0.6
- 3) The DR confirms receipt of the LSU and then floods it to other routers via address 224.0.0.5

These addresses are defined in RFC-1583.

224.0.0.5 OSPFIGP OSPFIGP All Routers

224.0.0.6 OSPFIGP OSPFIGP Designated Routers

Chapter: 4

---

### **QUESTION 742:**

You are a technician at Certkiller . Certkiller has an OSPF network. Your newly appointed Certkiller trainee wants to know why OSPF scales better than other routing protocols.

What would your reply be? (Choose all that apply.)

- A. OSPF converges faster
- B. OSPF uses a faster update protocol
- C. OSPF sends updates in a compressed format
- D. OSPF updates are incremental in nature, thus smaller in size

Answer: A, D

Explanation:

OSPF can scale better than some of its contemporaries due to a couple of factors. OSPF updates are based on topology changes, as opposed to RIP, which sends updates regardless of network topology status. When OSPF does send out an update, the update only contains elements that are different due to a linkstate change. RIP sends out a copy of its entire routing table whether anything has changed or not. OSPF is also aware of link costs, which means it won't send updates out over a slow link.

Chapter: 4

---

### **QUESTION 743:**

You are a technician at Certkiller . Your newly appointed Certkiller trainee wants to know why route summarization is so desirable

What would your reply be?

- A. summarization results in smaller routing tables

- B. summarization results in larger, more complex routing tables
- C. summarization is easier on router cpu due to decreased number of routes to manage
- D. summarization makes packets travel faster due to summarization compression routines

Answer: A, C

Explanation:

Route summarization is the process of condensing a series of routing table entries into a single entry that is maintained by the router. When you summarize, you compare a series of addresses and try to determine which ones have the same "high order" bit pattern (also known as significant bits). If you have a number of addresses with similar "high order" bit patterns, you can represent them with a single routing table entry as opposed to listing them separately. An end result of route summarization is that multiple "down stream" routes can be represented by a single table entry.

Chapter: 4

---

**QUESTION 744:**

You are a technician at Certkiller . The Certkiller network is connected to the internet via more than one ISP. Your newly appointed Certkiller trainee wants to know what this type of network is known as.

What would your reply be?

- A. multi-gnomed
- B. mega-homed
- C. multi-homed
- D. maxi-homed

Answer: C

Explanation:

If your company requires high availability to the internet you might consider having more than one ISP provide your internet connectivity. This arrangement is referred to as being "multi-homed". Many businesses that do a large amount of their business via "the net" utilize two or more ISP in order to provide fault tolerance should one service provider experience difficulties.

Chapter: 4

---

**QUESTION 745:**

Which of the following routing protocols support VLSM?

- A. EIGRP
- B. IGRP
- C. OSPF
- D. RIP version 1

E. RIP version 2

Answer: A, C, E

Explanation:

EIGRP, RIP v.2 and OSPF are classless routing protocols, which means that they are capable of including a specialized subnet mask in routing updates. Any protocol that carries a subnet mask in updates is a candidate for supporting VLSM. IGRP and RIP v.1 are classful routing protocols, which they are excluded from VLSM membership.

Chapter: 4

---

**QUESTION 746:**

You are a trainee technician at Certkiller . The Certkiller network implements the Border Gateway Protocol (BGP). Your instructor tells you that BGP comes in two flavors. He wants to what these two flavors are.

What would your reply be? (Choose all that apply.)

- A. reliable gateway protocol (RGP)
- B. interior gateway protocol (IGP)
- C. silent gateway protocol (SGP)
- D. exterior gateway protocol (EGP)

Answer: B, D

Explanation:

BGP runs in either of two capacities, as an IGP where it maintains routing tables inside the AS and EGP in which it maintains routing tables between AS's.

Chapter: 4

---

**QUESTION 747:**

You are a network engineer at Certkiller . Your newly appointed Certkiller trainee wants to know what command she should issue to view EIGRP entries in the routing table.

What would your reply be?

- A. show ip eigrp current
- B. show ip eigrp route recent
- C. show current ip eigrp route
- D. show ip eigrp route

Answer: D

Explanation:

The "show ip eigrp route" command will show only the routing table entries that are

EIGRP relevant.  
Chapter: 4

---

**QUESTION 748:**

With regard to BGP Prefix Lists, which of the following statements are true?

- A. prefix list sequence numbers by default start at 5 and increment by 5 as more are added
- B. if a router matches a prefix list entry, processing of that list ceases
- C. if a router matches a prefix list entry, processing of that list continues till another match is made
- D. the command "no ip prefix-list sequence-number" deletes the specified prefix list from memory

Answer: A, B

Explanation:

BGP prefix lists function much in the same way as ACLs, but with a few differences. Prefix lists allow increment updating, while ACLs do not. A router will cease processing a prefix list once a match is made; ACLs get processed all the way to the end. If you do not specify a BGP Prefix List sequence-numbering scheme, the list entries will increment by 5 as you enter new statements.

Chapter: 4

---

**QUESTION 749:**

You are a network engineer at Certkiller . Your newly appointed Certkiller trainee wants to know what the term "feasible distance" means.  
What would your reply be?

- A. It is the distance from the web server to the mail server
- B. It is the sum of the cost of the route from the next-hop to the next-hop
- C. It is the distance from the firewall to the ISP
- D. It is the distance from the next-hop router to the destination
- E. None of the above answers is correct

Answer: E

---

**QUESTION 750:**

What is the relevance of "priority" when dealing with DR/BDR elections?

- A. a priority of 1 means the router is eligible to become a DR
- B. a priority of -1 means the router is eligible to become a DR
- C. a priority of 0 means the router is ineligible to become a DR or BDR

D. DR/BDR elections run at regular intervals regardless of network topology status.

Answer: A, C

Explanation:

DR and BDR elections are handled as follows:

- Highest priority will become the DR
- Second Highest priority will become the BDR
- Only the malfunction of a DR or BDR will cause an election.
- A priority of "1" indicates eligibility to become a DR
- A priority of "0" indicates ineligibility to become a DR or BDR

Chapter: 4

---

### **QUESTION 751:**

You are the network administrator at Certkiller . Your newly appointed Certkiller trainee wants to know when the use of BGP would be appropriate. What would your reply be? (Choose all that apply.)

- A. if the border router in your AS is behind your corporate firewall
- B. if your AS is at the edge of another routing protocol's border
- C. if your network is multi-homed to the internet
- D. if you have a need to manage the traffic entering and exiting your autonomous system
- E. if your AS is a transit area for packets destined for another ASs

Answer: C, D, E

Explanation:

When to use BGP is a tough question.

The rule of thumb for when to use BGP usage is as follows:

- 1) if your network is multi-homed to the internet
- 2) if you have a need to manage the traffic entering and exiting your autonomous system
- 3) if your AS is a transit area for packets destined for another ASs

Chapter: 4

---

### **QUESTION 752:**

You are the network administrator at Certkiller . Your newly appointed Certkiller trainee wants to know when the use of BGP would NOT be appropriate. What would your reply be? (Choose all that apply.)

- A. if your network only has a single connection to the internet
- B. if you have a low bandwidth link between ASs
- C. if the selection of routes to outside ASs is not a high priority
- D. if the AS connecting you to the internet charges by the packet
- E. if the router connecting to your external AS is running IOS version 11.2 or earlier

Answer: A, B, C

Explanation:

When not to use BGP is an even tougher question.

The rule of thumb for when NOT to use BGP usage is as follows:

- 1) if your network only has a single connection to the internet
- 2) if you have a low bandwidth link between ASs
- 3) if the selection of routes to outside ASs is not a high priority

Most of the choices for not to BGP relate to the fact that BGP can easily overwhelm an underpowered link or router. So you should only use BGP if your situation really warrants the configuration and management that comes along with it.

Chapter: 4

---

**QUESTION 753:**

What attribute must all BGP implementations recognize?

- A. customized
- B. synchronized
- C. well-known
- D. well-adjusted
- E. optional

Answer: C

Explanation:

A BGP "well-known" attribute is one that must be recognized by all implementations of BGP and must be included in every update message. There are other types of updates that may or may not need to be included in the update messages.

The following BGP attributes are the most commonly known:

- 1) Well-known mandatory.
- 2) Well-known discretionary.
- 3) Optional transitive.
- 4) Optional non-transitive.

Chapter: 4