

Actualtests.com

The Power of Knowing



Exam : 642-801

Title : Building Scalable Cisco Internetworks (BSCI)

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 the 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 ipdirected-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:



```
Certkiller1# show run
hostname Certkiller1
.
ip route 10.0.0.0 255.0.0.0 interface fa0/0
ip route 10.0.0.0 255.0.0.0 172.17.20.3
```

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

NoteUp 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

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: Although link state protocols use multicasts for routing updates, they are sent using hellos. OSPF uses multicast IP addresses 224.0.0.5 and 224.0.0.6. LSAs are not sent every 90 seconds.

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 (RIPv2) - 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:

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:

Distance Vector

Sends complete routing tables at regular intervals

Based on the Bellman-Ford algorithm.

Gains knowledge of the network based only on information learned from directly connected neighbors.

Proprietary protocol type offered by Cisco

Link State

Routing updates occur only when a change occurs

Based on the Dijkstra algorithm.

Requires more memory and processing power.
