

Questions

1. Which of the following is true about BGP neighbor formation?
 - A. Neighbors are dynamically discovered via multicast Hello messages
 - B. A neighbor's IP address must be statically configured
 - C. A UDP session is established between neighbors
 - D. By default, BGP neighbors can be as many as 255 hops away from one another

2. Select the correct order of path selection criteria considered by BGP.
 - A. Weight, Local Preference, Origin Type, AS Path Length, Originate MED, Paths, Router ID
 - B. Router ID, Weight, Local Preference, Originate, AS Path Length, Origin Type, MED, Paths
 - C. Local Preference, Weight, Originate, AS Path Length, Origin Type, MED, Paths, Router ID
 - D. Weight, Local Preference, Originate, AS Path Length, Origin Type, MED, Paths, Router ID

3. Identify the statement that is NOT true concerning iBGP connections.
 - A. By default, a route received from an iBGP neighbor is not advertised to other iBGP neighbors.
 - B. When a router receives a route from an eBGP neighbor and advertises that route to an iBGP neighbor, the NEXT-HOP attribute is not updated.
 - C. When configuring an iBGP neighbor, the "neighbor" command uses the "local-as" parameter instead of the "remote-as" parameter.
 - D. A route reflector is often used within an autonomous system if there is not a full mesh of iBGP neighborships.

4. Which configuration is often used to influence outbound path selection on a BGP router with two or more neighbors in different autonomous systems?
 - A. Assigning a higher Local Preference value to routes coming in from a preferred neighbor
 - B. Assigning a lower Local Preference value to routes coming in from a preferred neighbor
 - C. Assigning a shorter AS Path value to routes coming in from a preferred neighbor
 - D. Assigning a longer AS Path value to routes coming in from a preferred neighbor

5. When configuring Multiprotocol BGP to advertise IPv6 routes over an IPv4 BGP session, what extra configuration step is needed?
 - A. You need to enable Cisco Express Forwarding (CEF).

- B. You need to apply a route map to your IPv6 neighbor specifying the next-hop IPv6 address.
 - C. You need to apply a route map to your IPv6 neighbor specifying the next-hop IPv4 address.
 - D. You need to disable Cisco Express Forwarding (CEF).
6. If you're configuring Multiprotocol BGP, where IPv4 routes are advertised over an IPv4 session and IPv6 routes are advertised over an IPv6 session, what step must be manually configured for an IPv6 neighbor that is automatically configured for an IP4 neighbor?
- A. The "ebpg-multihop" value must be specified.
 - B. The remote AS of a neighbor must be configured in IPv6 address family configuration mode.
 - C. A route-map must be configured to advertise the IPv6 next-hop address.
 - D. The IPv6 neighbor needs to be activated.

Questions and Answers

1. Which of the following is true about BGP neighbor formation?

- A. Neighbors are dynamically discovered via multicast Hello messages
- B. A neighbor's IP address must be statically configured
- C. A UDP session is established between neighbors
- D. By default, BGP neighbors can be as many as 255 hops away from one another

Answer: B

Explanation: BGP neighbors must be configured with one another's IP addresses, as opposed to dynamically discovering each other with multicast Hello messages, which are used by EIGRP and OSPF. BGP neighbors form a TCP session between themselves, rather than a UDP session. Also, even though BGP neighbors can be a maximum of 255 hops away from one another (using the "ebgp-multihop" command), by default, BGP neighbors must be adjacent to one another.

Video Reference: 3.3.1 Fundamental BGP Concepts

2. Select the correct order of path selection criteria considered by BGP.

- A. Weight, Local Preference, Origin Type, AS Path Length, Originate MED, Paths, Router ID
- B. Router ID, Weight, Local Preference, Originate, AS Path Length, Origin Type, MED, Paths
- C. Local Preference, Weight, Originate, AS Path Length, Origin Type, MED, Paths, Router ID
- D. Weight, Local Preference, Originate, AS Path Length, Origin Type, MED, Paths, Router ID

Answer: D

Explanation: The correct order of BGP path selection criteria is: Weight, Local Preference, Originate, AS Path Length, Origin Type, MED, Paths, and Router ID. A memory aid for remembering this order is the acronym: "We Love Oranges AS Oranges Mean Pure Refreshment." The main challenge with this memory aid is correctly ordering the "Originate" and "Origin Type" criteria, because they both begin with a "O."

Video Reference: 3.3.2 BGP Path Selection Criteria

3. Identify the statement that is NOT true concerning iBGP connections.

- A. By default, a route received from an iBGP neighbor is not advertised to other iBGP neighbors.
- B. When a router receives a route from an eBGP neighbor and advertises that route to an iBGP neighbor, the NEXT-HOP attribute is not updated.
- C. When configuring an iBGP neighbor, the "neighbor" command uses the "local-as" parameter instead of the "remote-as" parameter.
- D. A route reflector is often used within an autonomous system if there is not a full mesh of iBGP neighborships.

Answer: C

Explanation: An iBGP (Internal BGP) neighborship is formed between two routers within an autonomous system (AS). An eBGP (External BGP) neighborship is formed between two routers in different autonomous systems. When a router receives a route from an eBGP neighbor, it advertises that route to any iBGP neighbors without updating the NEXT-HOP attribute (which can be addressed by configuring the NEXT-HOP-SELF option). Also, when a router receives a route advertisement from an iBGP neighbor, the router does not advertise that route to other iBGP neighbors (which can be addressed using a Route Reflector or by configuring a full mesh of iBGP neighborships). Interestingly, the “neighbor remote-as” command is used to form a neighborship between routers in different autonomous systems as well as between routers in the same autonomous system.

Video Reference: 3.3.3 Routing IPv4 with BGP - Part 1

4. Which configuration is often used to influence outbound path selection on a BGP router with two or more neighbors in different autonomous systems?
- A. Assigning a higher Local Preference value to routes coming in from a preferred neighbor
 - B. Assigning a lower Local Preference value to routes coming in from a preferred neighbor
 - C. Assigning a shorter AS Path value to routes coming in from a preferred neighbor
 - D. Assigning a longer AS Path value to routes coming in from a preferred neighbor

Answer: A

Explanation: The Local Preference path selection parameter is commonly used for influencing outbound path selection decisions, with higher values being preferred. The AS Path attribute is commonly used for influencing inbound path selection decisions, with shorter AS Paths being preferred.

Video Reference: 3.3.4 Routing IPv4 with BGP - Part 2

5. When configuring Multiprotocol BGP to advertise IPv6 routes over an IPv4 BGP session, what extra configuration step is needed?
- A. You need to enable Cisco Express Forwarding (CEF).
 - B. You need to apply a route map to your IPv6 neighbor specifying the next-hop IPv6 address.
 - C. You need to apply a route map to your IPv6 neighbor specifying the next-hop IPv4 address.
 - D. You need to disable Cisco Express Forwarding (CEF).

Answer: B

Explanation: IPv6 routes can be advertised over either an IPv4 or an IPv6 session with Multiprotocol BGP. However, if an IPv4 session is used, the receiving BGP neighbor doesn't

learn the IPv6 address of the router sending the IPv6 route advertisement. To overcome this issue, you can configure a route map to add the IPv6 next-hop address to IPv6 route advertisements.

Video Reference: 3.3.6 Routing IPv6 with BGP Over an IPv4 Session Configuration

6. If you're configuring Multiprotocol BGP, where IPv4 routes are advertised over an IPv4 session and IPv6 routes are advertised over an IPv6 session, what step must be manually configured for an IPv6 neighbor that is automatically configured for an IP4 neighbor?
- A. The "ebpg-multihop" value must be specified.
 - B. The remote AS of a neighbor must be configured in IPv6 address family configuration mode.
 - C. A route-map must be configured to advertise the IPv6 next-hop address.
 - D. The IPv6 neighbor needs to be activated.

Answer: D

Explanation: When configuring Multiprotocol BGP, neighbors are specified under router configuration mode. Then, under router-address-family configuration mode, the neighbors need to be activated. Interestingly, the "neighbor [neighbor_ip_address] activate" command is automatically entered for the IPv4 address family but must be manually configured for the IPv6 address family.

Video Reference: 3.3.7 Routing IPv6 with BGP Over an IPv6 Session Configuration