



Welcome to Cert007 - Your Ultimate IT Certification Partner



- Real Exam Questions
- Free Updates
- Expert Support
- Instant Access
- Money-Back Guarantee



Visit us at <https://www.cert007.com/> for more information

Exam : **JN0-664**

Title : Service Provider
Professional (JNCIP-SP)

Version : DEMO

1.Which two statements are correct about reflecting inet-vpn unicast prefixes in BGP route reflection?
(Choose two.)


- A. Route reflectors do not change any existing BGP attributes by default when advertising routes.
- B. A BGP peer does not require any configuration changes to become a route reflector client.
- C. Clients add their originator ID when advertising routes to their route reflector
- D. Route reflectors add their cluster ID to the AS path when readvertising client routes.

Answer: A,B

Explanation:

Route reflection is a BGP feature that allows a router to reflect routes learned from one IBGP peer to another IBGP peer, without requiring a full-mesh IBGP topology. Route reflectors do not change any existing BGP attributes by default when advertising routes, unless explicitly configured to do so. A BGP peer does not require any configuration changes to become a route reflector client, only the route reflector needs to be configured with the client parameter under [edit protocols bgp group group-name neighbor neighbor-address] hierarchy level.

2.Exhibit



Click the Exhibit button-Referring to the exhibit, which two statements are correct about BGP routes on R3 that are learned from the ISP-A neighbor? (Choose two.)

- A. By default, the next-hop value for these routes is not changed by ISP-A before being sent to R3.
- B. The BGP local-preference value that is used by ISP-A is not advertised to R3.
- C. All BGP attribute values must be removed before receiving the routes.
- D. The next-hop value for these routes is changed by ISP-A before being sent to R3.

Answer: A,B

Explanation:

BGP is an exterior gateway protocol that uses path vector routing to exchange routing information among autonomous systems. BGP uses various attributes to select the best path to each destination and to propagate routing policies. Some of the common BGP attributes are AS path, next hop, local preference, MED, origin, weight, and community. BGP attributes can be classified into four categories: well-known mandatory, well-known discretionary, optional transitive, and optional nontransitive. Well-known mandatory attributes are attributes that must be present in every BGP update message and must be recognized by every BGP speaker. Well-known discretionary attributes are attributes that may or may not be present in a BGP update message but must be recognized by every BGP speaker. Optional transitive attributes are attributes that may or may not be present in a BGP update message and may or may not be recognized by a BGP speaker. If an optional transitive attribute is not recognized by a BGP speaker, it is passed along to the next BGP speaker. Optional nontransitive attributes are attributes that may or may not be present in a BGP update message and may or may not be recognized by a BGP speaker. If an optional nontransitive attribute is not recognized by a BGP speaker, it is not passed along to the next BGP speaker. In this question, we have four routers (R1, R2, R3, and R4) that are connected in a full mesh topology and running IBGP. R3 receives the 192.168.0.0/16 route from its EBGP neighbor and advertises it to R1 and R4 with different BGP attribute values. We are asked which statements are correct about the BGP routes on R3 that are learned from the ISP-A neighbor.

Based on the information given, we can infer that the correct statements are:

- ☞ By default, the next-hop value for these routes is not changed by ISP-A before being sent to R3. This is because the default behavior of EBGP is to preserve the next-hop attribute of the routes received from another EBGP neighbor. The next-hop attribute indicates the IP address of the router that should be used as the next hop to reach the destination network.
- ☞ The BGP local-preference value that is used by ISP-A is not advertised to R3. This is because the local-preference attribute is a well-known discretionary attribute that is used to influence the outbound traffic from an autonomous system. The local-preference attribute is only propagated within an autonomous system and is not advertised to external neighbors.

References:

- <https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13753-25.html>:
- <https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13762-40.html>:
- <https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13759-37.html>

3.A packet is received on an interface configured with transmission scheduling. One of the configured queues.

In this scenario, which two actions will be taken by default on a Junos device? (Choose two.)

- A. The excess traffic will be discarded
- B. The exceeding queue will be considered to have negative bandwidth credit.
- C. The excess traffic will use bandwidth available from other queues
- D. The exceeding queue will be considered to have positive bandwidth credit

Answer: A,C

Explanation:

In Junos devices, when a packet is received on an interface configured with transmission scheduling, and one of the configured queues is exceeding its allocated bandwidth, the typical actions taken are based on the scheduling configuration and congestion management mechanisms in place. Here are the

two likely default actions:

The excess traffic will be discarded. When a queue exceeds its configured bandwidth, and if there are no other congestion management mechanisms in place (like buffer or RED profiles), the excess traffic could be dropped by default.

The excess traffic will use bandwidth available from other queues. If excess bandwidth is available from other queues and the scheduling configuration allows for it, excess traffic may utilize unused bandwidth from other queues. This is typical behavior in scenarios where queues are configured with some form of shared bandwidth allocation or where one queue can borrow unused bandwidth from others.

4.Which two statements are correct about VPLS tunnels? (Choose two.)

- A. LDP-signaled VPLS tunnels only support control bit 0.
- B. LDP-signaled VPLS tunnels use auto-discovery to provision sites
- C. BGP-signaled VPLS tunnels can use either RSVP or LDP between the PE routers.
- D. BGP-signaled VPLS tunnels require manual provisioning of sites.

Answer: B,D


Explanation:

In the context of Virtual Private LAN Service (VPLS) and the signaling protocols used to establish VPLS tunnels:

LDP-signaled VPLS tunnels use auto-discovery to provision sites. In LDP-signaled VPLS, auto-discovery is used to discover other PE routers that are part of the same VPLS instance. This is typically done through the exchange of LDP messages that carry VPLS labels.

BGP-signaled VPLS tunnels require manual provisioning of sites. When using BGP for signaling in VPLS (also known as BGP-based VPLS), each site needs to be manually provisioned. This includes configuring the site identifier and the parameters for the VPLS instance on the PE router.

5.Exhibit



The network shown in the exhibit is based on IS-IS

Which statement is correct in this scenario?

- A. The NSEL byte for Area 0001 is 00.
- B. The area address is two bytes.
- C. The routers are using unnumbered interfaces
- D. The system ID of R1_2 is 192.168.16.1

Answer: A

Explanation:

IS-IS is an interior gateway protocol that uses link-state routing to exchange routing information among routers within a single autonomous system. IS-IS uses two types of addresses to identify routers and areas: system ID and area address. The system ID is a unique identifier for each router in an IS-IS domain. The system ID is 6 octets long and can be derived from the MAC address or manually configured. The area address is a variable-length identifier for each area in an IS-IS domain. The area address can be 1 to 13 octets long and is composed of high-order octets of the address. An IS-IS instance may be assigned multiple area addresses, which are considered synonymous. Multiple synonymous area addresses are useful when merging or splitting areas in the domain¹. In this question, we have a network based on IS-IS with four routers (R1_1, R1_2, R2_1, and R2_2) belonging to area 0001. The area address for area 0001 is 49.0001. The NSEL byte for area 0001 is the last octet of the address, which is 01. The NSEL byte stands for Network Service Access Point Selector (NSAP Selector) and indicates the type of service requested from the network layer². Therefore, the correct statement in this scenario is that the NSEL byte for area 0001 is 01.

References:

1: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_isis/configuration/xr-16/irs-xr-16-book/irs-ovrvw-cf.html

2: <https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/concept/is-is-routing-overview.html>