



Juniper

Exam Questions JN0-664

Service Provider - Professional (JNCIP-SP)

NEW QUESTION 1

When building an interprovider VPN, you notice on the PE router that you have hidden routes which are received from your BGP peer with family inet labeled-unicast configured.

Which parameter must you configure to solve this problem?

- A. Under the family inet labeled-unicast hierarchy, add the explicit null parameter.
- B. Under the protocols ospf hierarchy, add the traffic-engineering parameter.
- C. Under the family inet labeled-unicast hierarchy, add the resolve-vpn parameter.
- D. Under the protocols mpls hierarchy, add the traffic-engineering parameter

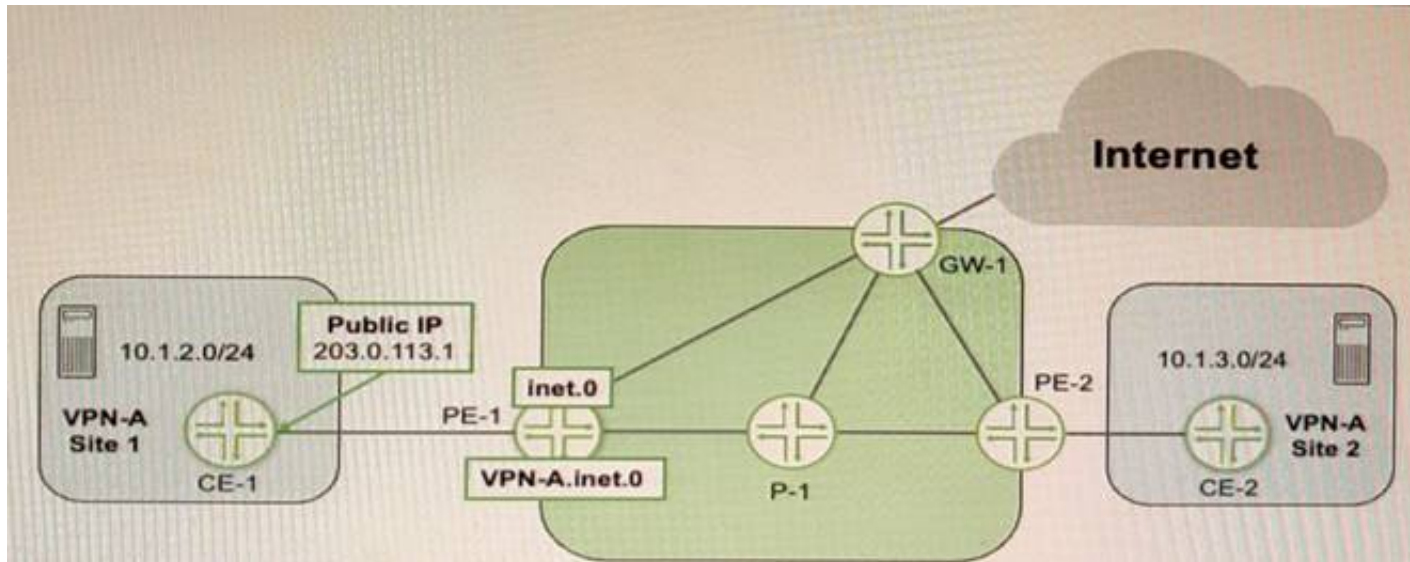
Answer: C

Explanation:

The resolve-vpn parameter is a BGP option that allows a router to resolve labeled VPN-IPv4 routes using unlabeled IPv4 routes received from another BGP peer with family inet labeled-unicast configured. This option enables interprovider VPNs without requiring MPLS labels between ASBRs or using VRF tables on ASBRs. In this scenario, you need to configure the resolve-vpn parameter under [edit protocols bgp group external family inet labeled-unicast] hierarchy level on both ASBRs.

NEW QUESTION 2

Exhibit



Referring to the exhibit, CE-1 is providing NAT services for the hosts at Site 1 and you must provide Internet access for those hosts. Which two statements are correct in this scenario? (Choose two.)

- A. You must configure a static route in the main routing instance for the 10.1.2.0/24 prefix that uses the VPN-A.inet.0 table as the next hop.
- B. You must configure a static route in the main routing instance for the 203.0.113.1/32 prefix that uses the VPN-A.inet.0 table as the next hop.
- C. You must configure a RIB group on PE-1 to leak a default route from the inet.0 table to the VPN-A.inet.0 table.
- D. You must configure a RIB group on PE-1 to leak the 10.1.2.0/24 prefix from the VPN-A.inet.0 table to the inet.0 table.


Answer: AB

Explanation:

To provide Internet access for the hosts at Site 1, you need to configure static routes in the main routing instance on PE-1 that point to the VPN-A.inet.0 table as the next hop. This allows PE-1 to forward traffic from the Internet to CE-1 using MPLS labels and vice versa. You need to configure two static routes: one for the 10.1.2.0/24 prefix that represents the private network of Site 1, and one for the 203.0.113.1/32 prefix that represents the public IP address of CE-1.

NEW QUESTION 3

Exhibit



You want Site 1 to access three VLANs that are located in Site 2 and Site 3. The customer-facing interface on the PE-1 router is configured for Ethernet-VLAN encapsulation.

What is the minimum number of L2VPN routing instances to be configured to accomplish this task?

- A. 1
- B. 3
- C. 2
- D. 6

Answer: B

Explanation:

To allow Site 1 to access three VLANs that are located in Site 2 and Site 3, you need to configure three L2VPN routing instances on PE-1, one for each VLAN. Each L2VPN routing instance will have a different VLAN ID and a different VNI for VXLAN encapsulation. Each L2VPN routing instance will also have a different vrf-target export value to identify which VPN routes belong to which VLAN. This way, PE-1 can forward traffic from Site 1 to Site 2 and Site 3 based on the VLAN tags and VNIs.

NEW QUESTION 4

Exhibit

```
user@router> show route extensive
...
2:192.168.101.5:65101::22031::02:00:31:06:00:01/304 MAC/IP (2 entries, 1
announced)
TSI:
Page 0 idx 0, (group IBGP-EVPN-Core type Internal) Type 1 val 0xb225964
(adv_entry)
  Advertised metrics:
    Nexthop: 192.168.101.5
    Localpref: 100
    AS path: [65101] I (Originator)
    Cluster list: 2.2.2.2
    Originator ID: 192.168.101.5
    Communities: target:65101:268457487 encapsulation:vxlan(0x8)
    Cluster ID: 3.3.3.3
    Advertise: 00000001
Path 2:192.168.101.5:65101::22031::02:00:31:06:00:01 from 192.168.101.3 Vector
len 4. Val: 0
  *BGP Preference: 170/-101
    Route Distinguisher: 192.168.101.5:65101
    Next hop type: Indirect, Next hop index: 0
    Address: 0xb2d3490
    Next-hop reference count: 10520
    Source: 192.168.101.3
    Protocol next hop: 192.168.101.5
    Indirect next hop: 0x2 no-forward INH Session ID: 0x0
    State: <Active Int Ext>
    Local AS: 65101 Peer AS: 65101
    Age: 3d 19:56:57 Metric2: 0
    Validation State: unverified
    Task: BGP_65101.192.168.101.3
    Announcement bits (1): 1-BGP_RT_Background
    AS path: I (Originator)
    Cluster list: 2.2.2.2
    Originator ID: 192.168.101.5
    Communities: target:65101:268457487 encapsulation:vxlan(0x8)
    Import Accepted
    Route Label: 22031
    ESI: 05:00:00:fe:4d:00:00:56:0f:00
    Localpref: 100
    Router ID: 192.168.101.3
    Secondary Tables: default-switch.evpn.0
    Indirect next hops: 1
      Protocol next hop: 192.168.101.5
      Indirect next hop: 0x2 no-forward INH Session ID: 0x0
      Indirect path forwarding next hops: 2
        Next hop type: Router
        Next hop: 10.0.2.12 via et-0/0/0.0
        Session Id: 0x0
        Next hop: 10.0.2.22 via et-0/0/1.0
        Session Id: 0x0

192.168.101.5/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 2
Nexthop: 10.0.2.12 via et-0/0/0.0
Session Id: 0
Nexthop: 10.0.2.22 via et-0/0/1.0
Session Id: 0
...
```

Referring to the exhibit, which two statements are true? (Choose two.)

- A. This route is learned through EBGp
- B. This is an EVpN Type-2 route.
- C. The device advertising this route into EVpN is 192.168.101.5.
- D. The devices advertising this route into EVpN are 10 0 2 12 and 10.0.2.22.

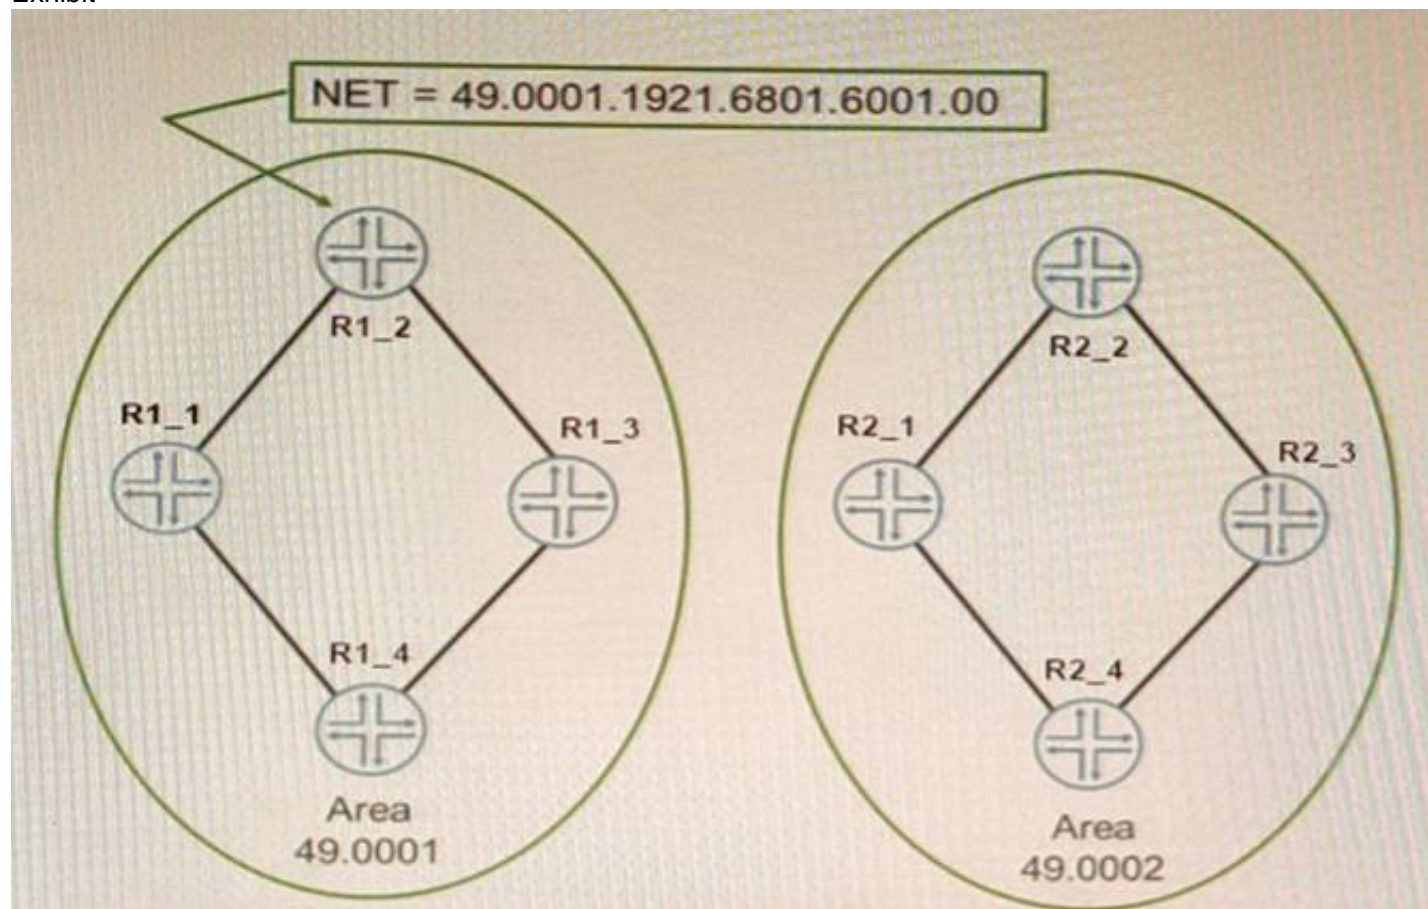
Answer: BC

Explanation:

This is an EVpN Type-2 route, also called a MAC/IP advertisement route, that is used to advertise host IP and MAC address information to other VTEPs in an EVpN network. The route type field in the EVpN NLRI has a value of 2, indicating a Type-2 route. The device advertising this route into EVpN is 192.168.101.5, which is the IP address of the VTEP that learned the host information from the local CE device. This IP address is carried in the MPLS label field of the route as part of the VXLAN encapsulation.

NEW QUESTION 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 domain. 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 layer. 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/ios-xr-16-book/ios-ovrvw-cf.html 2: <https://www.juniper.net/documentation/us/en/software/junos/isis/topics/concept/isis-routing-overview.html>

NEW QUESTION 6

You want to ensure that L1 IS-IS routers have only the most specific routes available from L2 IS-IS routers. Which action accomplishes this task?

- A. Configure the ignore-attached-bit parameter on all L2 routers.
- B. Configure all routers to allow wide metrics.
- C. Configure all routers to be L1.
- D. Configure the ignore-attached-bit parameter on all L1 routers


Answer: D

Explanation:

The attached bit is a flag in an IS-IS LSP that indicates whether a router is connected to another area or level (L2) of the network. By default, L2 routers set this bit when they advertise their LSPs to L1 routers, and L1 routers use this bit to select a default route to reach other areas or levels through L2 routers. However, this may result in suboptimal routing if there are multiple L2 routers with different paths to other areas or levels. To ensure that L1 routers have only the most specific routes available from L2 routers, you can configure the ignore-attached-bit parameter on all L1 routers. This makes L1 routers ignore the attached bit and install all interarea routes learned from L2 routers in their routing tables.

NEW QUESTION 7

Exhibit



Referring to the exhibit, you must provide Internet access for VPN-A using CE-1 as the hub CE. Which two statements are correct in this situation? (Choose two.)

- A. You must use RIB groups to leak routes between the inet
- B. o and vpn-
- C. ine
- D. o tables.
- E. RIB groups are not needed to leak routes between the ine
- F. 0 and VPN—
- G. ine
- H. 0 tables,
- I. Internet traffic from Site 2 takes the path of PE-2 -> PE-1 -> GW-1.
- J. Internet traffic from Site 2 takes the path of PE-2 -> PE-1 -> CE-1 -> PE-1 -> GW-1.

Answer: AD

Explanation:

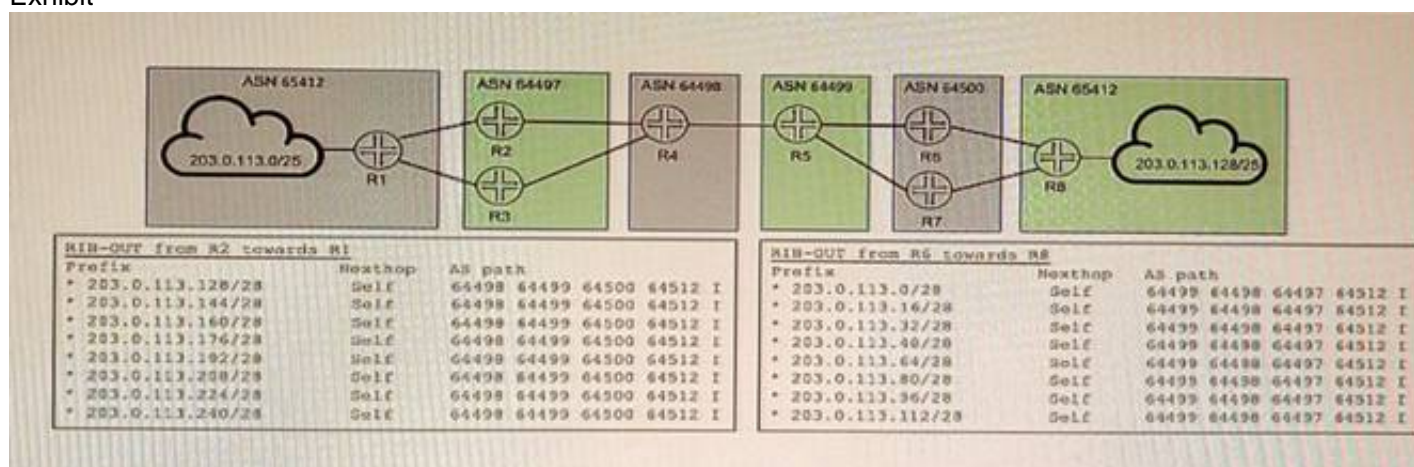
To provide Internet access for VPN-A using CE-1 as the hub CE, you need to do the following:

? You must use RIB groups to leak routes between the inet.0 and vpn-a.inet.0 tables on PE-1 and CE-1. RIB groups are routing options that allow you to import routes from one routing table into another routing table based on certain criteria. In this scenario, you need to configure RIB groups on PE-1 and CE-1 to import Internet routes from inet.0 into vpn-a.inet.0 and vice versa.

? Internet traffic from Site 2 takes the path of PE-2 -> PE-1 -> CE-1 -> PE-1 -> GW-1. This is because Site 2 does not have direct Internet access and needs to use CE-1 as its default gateway for Internet traffic. Site 2 sends its Internet traffic to PE-2, which forwards it to PE-1 based on VPN-A routes. PE-1 then sends it to CE-1 based on RIB group import policy. CE-1 then sends it back to PE-1 based on its default route pointing to GW-1. PE-1 then forwards it to GW-1 based on RIB group import policy again.

NEW QUESTION 8

Exhibit



R1 and R8 are not receiving each other's routes

Referring to the exhibit, what are three configuration commands that would solve this problem? (Choose three.)

- A. Configure loops and advertise-peer-as on routers in AS 64497 and AS 64450.
- B. Configure loops on routers in AS 65412 and advertise-peer-as on routers in AS 64498.
- C. Configure as-override on advertisement from AS 64500 toward AS 64512.
- D. Configure remove-private on advertisements from AS 64497 toward AS 64498
- E. Configure remove-private on advertisements from AS 64500 toward AS 64499

Answer: BDE

Explanation:

The problem in this scenario is that R1 and R8 are not receiving each other's routes because of private AS numbers in the AS path. Private AS numbers are not globally unique and are not advertised to external BGP peers. To solve this problem, you need to do the following:

? Configure loops on routers in AS 65412 and advertise-peer-as on routers in AS 64498. This allows R5 and R6 to advertise their own AS number (65412) instead of their peer's AS number (64498) when sending updates to R7 and R8. This prevents a loop detection issue that would cause R7 and R8 to reject the routes from R5 and R62.

? Configure remove-private on advertisements from AS 64497 toward AS 64498 and from AS 64500 toward AS 64499. This removes any private AS numbers from the AS path before sending updates to external BGP peers. This allows R2 and R3 to receive the routes from R1 and R4, respectively3.

NEW QUESTION 9

You are configuring a BGP signaled Layer 2 VPN across your MPLS enabled core network. Your PE-2 device connects to two sites within the s VPN. In this scenario, which statement is correct?

- A. By default on PE-2, the site's local ID is automatically assigned a value of 0 and must be configured to match the total number of attached sites.
- B. You must create a unique Layer 2 VPN routing instance for each site on the PE-2 device.
- C. You must use separate physical interfaces to connect PE-2 to each site.
- D. By default on PE-2, the remote site IDs are automatically assigned based on the order that you add the interfaces to the site configuration.

Answer: D

Explanation:

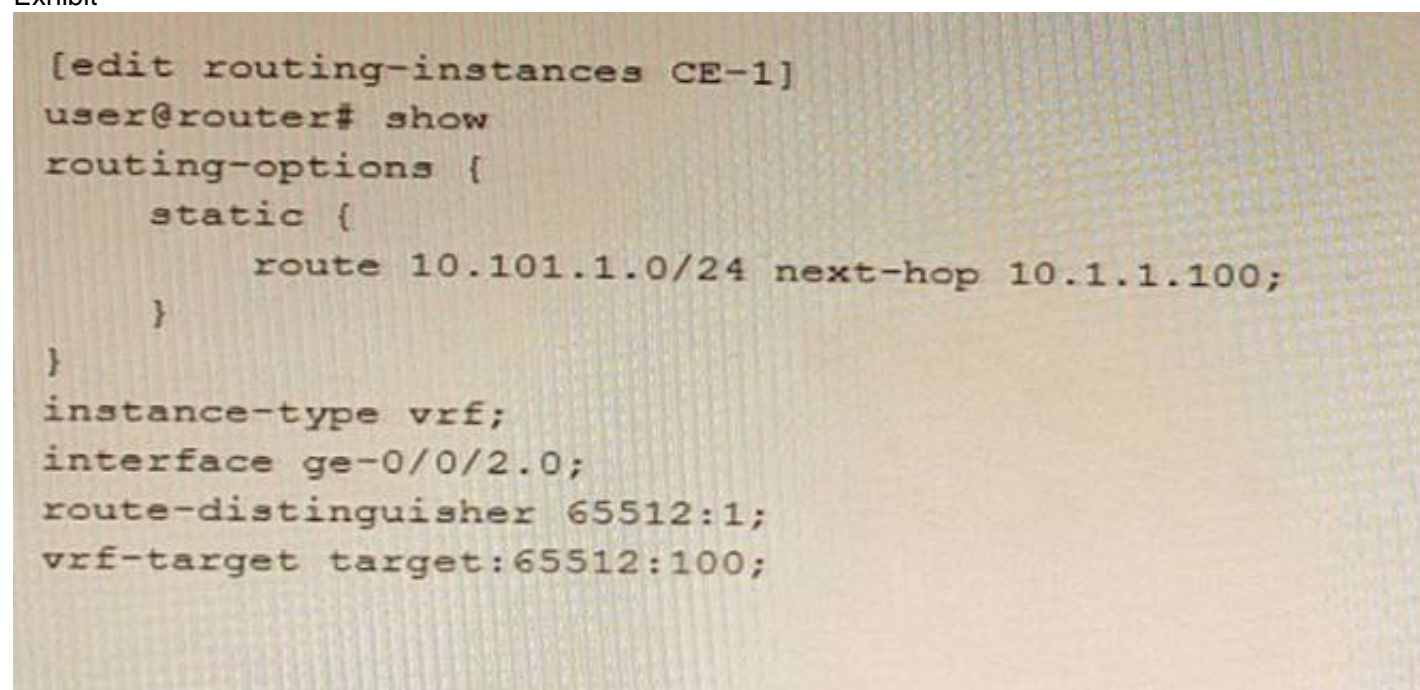
BGP Layer 2 VPNs use BGP to distribute endpoint provisioning information and set up pseudowires between PE devices. BGP uses the Layer 2 VPN (L2VPN) Routing Information Base (RIB) to store endpoint provisioning information, which is updated each time any Layer 2 virtual forwarding instance (VFI) is configured. The prefix and path information is stored in the L2VPN database, which allows BGP to make decisions about the best path.

In BGP Layer 2 VPNs, each site has a unique site ID that identifies it within a VFI. The site ID can be manually configured or automatically assigned by the PE device. By default, the site ID is automatically assigned based on the order that you add the interfaces to the site configuration. The first interface added to a site configuration has a site ID of 1, the second interface added has a site ID of 2, and so on.

Option D is correct because by default on PE-2, the remote site IDs are automatically assigned based on the order that you add the interfaces to the site configuration. Option A is not correct because by default on PE-2, the site's local ID is automatically assigned a value of 0 and does not need to be configured to match the total number of attached sites. Option B is not correct because you do not need to create a unique Layer 2 VPN routing instance for each site on the PE-2 device. You can create one routing instance for all sites within a VFI. Option C is not correct because you do not need to use separate physical interfaces to connect PE-2 to each site. You can use subinterfaces or service instances on a single physical interface.

NEW QUESTION 10

Exhibit



```
[edit routing-instances CE-1]
user@router# show
routing-options {
    static {
        route 10.101.1.0/24 next-hop 10.1.1.100;
    }
}
instance-type vrf;
interface ge-0/0/2.0;
route-distinguisher 65512:1;
vrf-target target:65512:100;
```

Referring to the exhibit, which statement is true?

- A. The 10.101.1.0/24 route will be shared if the vrf-table-label parameter is configured.
- B. The 10.101.1.0/24 route will only be shared if BGP is configured in the routing instance
- C. The 10.101.1.0/24 route will be shared if there are other VRFs that use the same route target community
- D. The 10.101.1.0/24 route will be shared if the auto-export parameter is configured

Answer: D

Explanation:

The auto-export parameter is a routing option that allows a routing instance to share routes with other routing instances or the master routing table. The auto-export parameter automatically exports routes from one routing instance to another based on the route target communities attached to the routes. In this scenario, the 10.101.1.0/24 route will be shared if the auto-export parameter is configured under [edit routing-options] hierarchy level.

NEW QUESTION 10

By default, which statement is correct about OSPF summary LSAs?

- A. All Type 2 and Type 7 LSAs will be summarized into a single Type 5 LSA
- B. The area-range command must be installed on all routers.
- C. Type 3 LSAs are advertised for routes in Type 1 LSAs.
- D. The metric associated with a summary route will be equal to the lowest metric associated with an individual contributing route

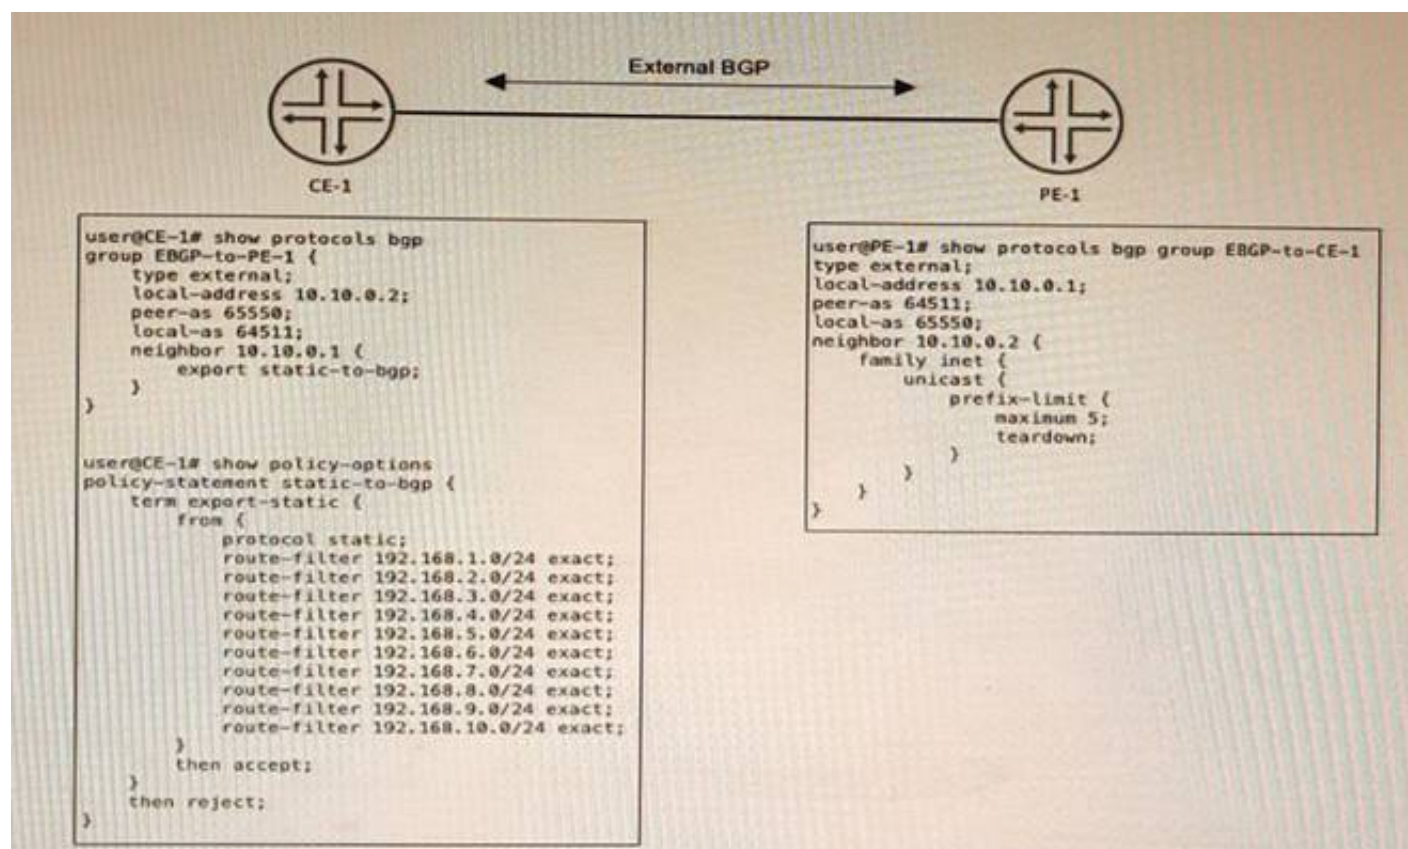
Answer: C

Explanation:

OSPF uses different types of LSAs to describe different aspects of the network topology. Type 1 LSAs are also known as router LSAs, and they describe the links and interfaces of a router within an area. Type 3 LSAs are also known as summary LSAs, and they describe routes to networks outside an area but within the same autonomous system (AS). By default, OSPF will summarize routes from Type 1 LSAs into Type 3 LSAs when advertising them across area boundaries.

NEW QUESTION 11

Exhibit



CE-1 must advertise ten subnets to PE-1 using BGP. Once CE-1 starts advertising the subnets to PE-1, the BGP peering state changes to Active. Referring to the CLI output shown in the exhibit, which statement is correct?

- A. CE-1 is advertising its entire routing table.
- B. CE-1 is configured with an incorrect peer AS
- C. The prefix limit has been reached on PE-1
- D. CE-1 is unreachable

Answer: B

Explanation:

The problem in this scenario is that CE-1 is configured with an incorrect peer AS number for its BGP session with PE-1. The CLI output shows that CE-1 is using AS 65531 as its local AS number and AS 65530 as its peer AS number. However, PE-1 is using AS 65530 as its local AS number and AS 65531 as its peer AS number. This causes a mismatch in the BGP OPEN messages and prevents the BGP session from being established. To solve this problem, CE-1 should configure its peer AS number as 65530 under [edit protocols bgp group external] hierarchy level.

NEW QUESTION 12

Your organization manages a Layer 3 VPN for multiple customers. To support advanced route than one BGP community on advertised VPN routes to remote PE routers.

Which routing-instance configuration parameter would support this requirement?

- A. vrf-export
- B. vrf-import
- C. vrf-target export
- D. vrf-target import

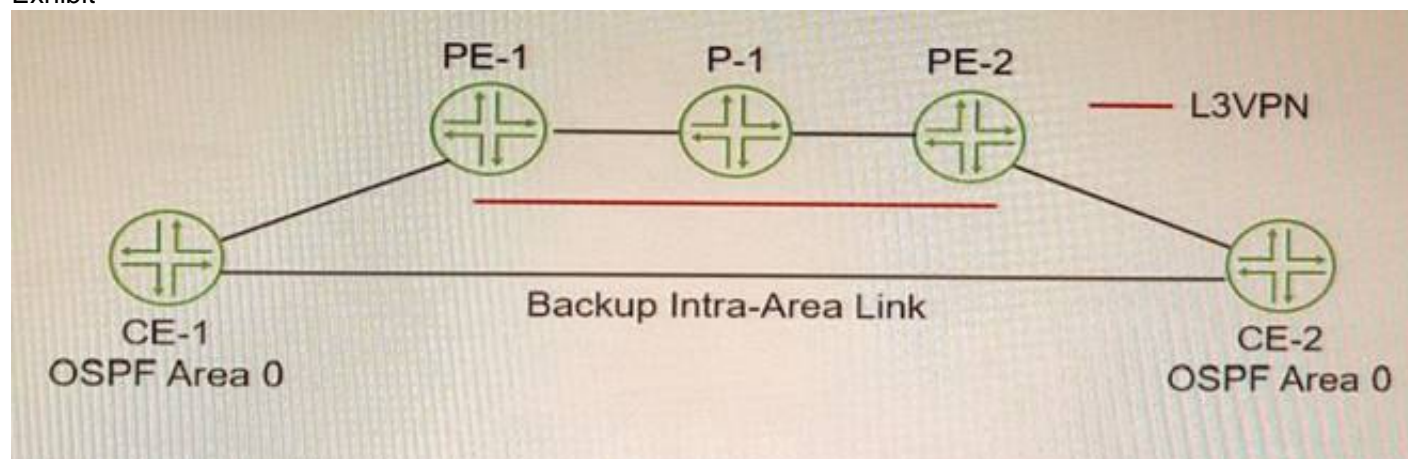
Answer: C

Explanation:

The vrf-target export parameter is used to specify one or more BGP extended community attributes that are attached to VPN routes when they are exported from a VRF routing instance to remote PE routers. This parameter allows you to control which VPN routes are accepted by remote PE routers based on their import policies. You can specify more than one vrf-target export value for a VRF routing instance to support advanced route filtering or route leaking scenarios.

NEW QUESTION 17

Exhibit



You must ensure that the VPN backbone is preferred over the back door intra-area link as long as the VPN is available. Referring to the exhibit, which action will accomplish this task?

- A. Configure an import routing policy on the CE routers that rejects OSPF routes learned on the backup intra-area link.
- B. Enable OSPF traffic-engineering.
- C. Configure the OSPF metric on the backup intra-area link that is higher than the L3VPN link.
- D. Create an OSPF sham link between the PE routers.

Answer: D

Explanation:

A sham link is a logical link between two PE routers that belong to the same OSPF area but are connected through an L3VPN. A sham link makes the PE routers appear as if they are directly connected, and prevents OSPF from preferring an intra-area back door link over the VPN backbone. To create a sham link, you need to configure the local and remote addresses of the PE routers under the [edit protocols ospf area area-id] hierarchy level1.

NEW QUESTION 20

In which two ways does OSPF prevent routing loops in multi-area networks? (Choose two.)

- A. All areas are required to connect as a full mesh.
- B. The LFA algorithm prunes all looped paths within an area.
- C. All areas are required to connect to area 0.
- D. The SPF algorithm prunes looped paths within an area.

Answer: CD

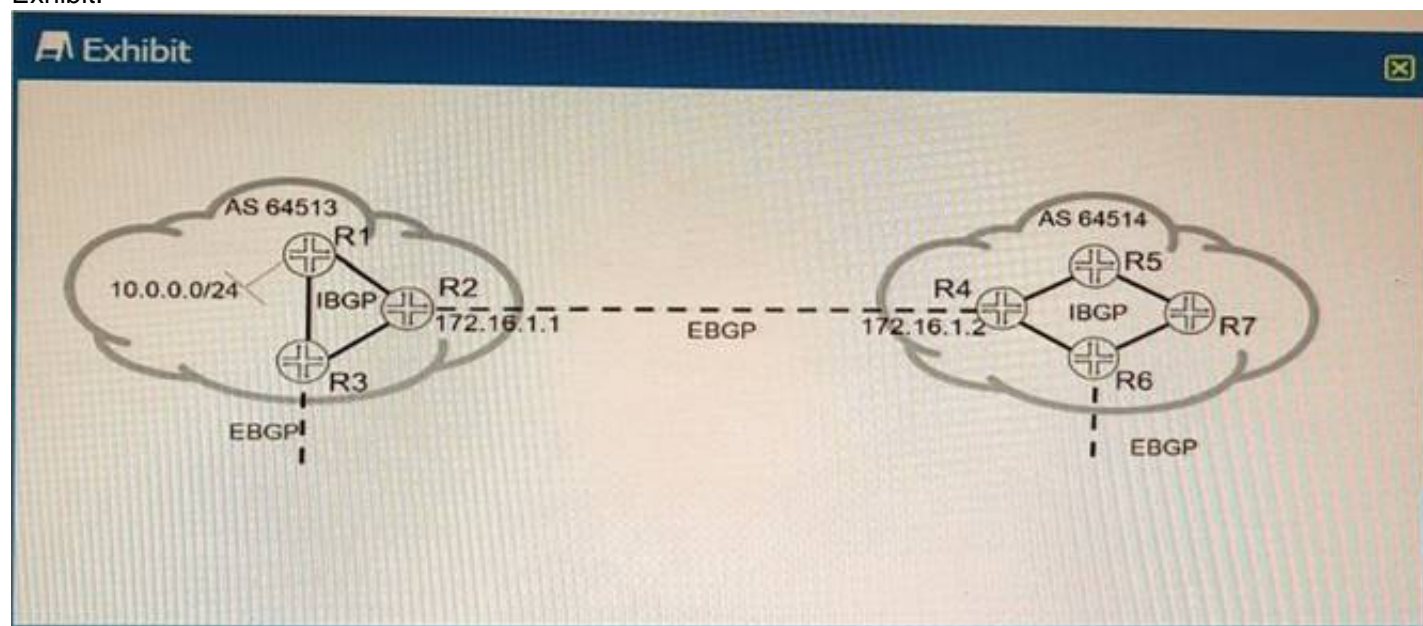
Explanation:

OSPF is an interior gateway protocol that uses link-state routing to exchange routing information among routers within a single autonomous system. OSPF prevents routing loops in multi-area networks by using two methods: area hierarchy and SPF algorithm. Area hierarchy is the concept of dividing a large OSPF network into smaller areas that are connected to a backbone area (area 0). This reduces the amount of routing information that each router has to store and process, and also limits the scope of link-state updates within each area. All areas are required to connect to area 0 either directly or through virtual links2. SPF algorithm is the method that OSPF uses to calculate the shortest path to each destination in the network based on link-state information. The SPF algorithm runs on each router and builds a shortest-path tree that represents the topology of the network from the router's perspective. The SPF algorithm prunes looped paths within an area by choosing only one best path for each destination3.

References: 2: <https://www.juniper.net/documentation/us/en/software/junos/ospf/topics/concept/ospf-area-overview.html> 3: <https://www.juniper.net/documentation/us/en/software/junos/ospf/topics/concept/ospf-spf-algorithm-overview.html>

NEW QUESTION 25

Exhibit.



Referring to the exhibit; the 10.0.0.0/24 EBGP route is received on R5; however, the route is being hidden. What are two solutions that will solve this problem? (Choose two.)

- A. On R4, create a policy to change the BGP next hop to itself and apply it to IBGP as an export policy
- B. Add the external interface prefix to the IGP routing tables
- C. Add the internal interface prefix to the BGP routing tables.
- D. On R4, create a policy to change the BGP next hop to 172.16.1.1 and apply it to IBGP as an export policy

Answer: AB

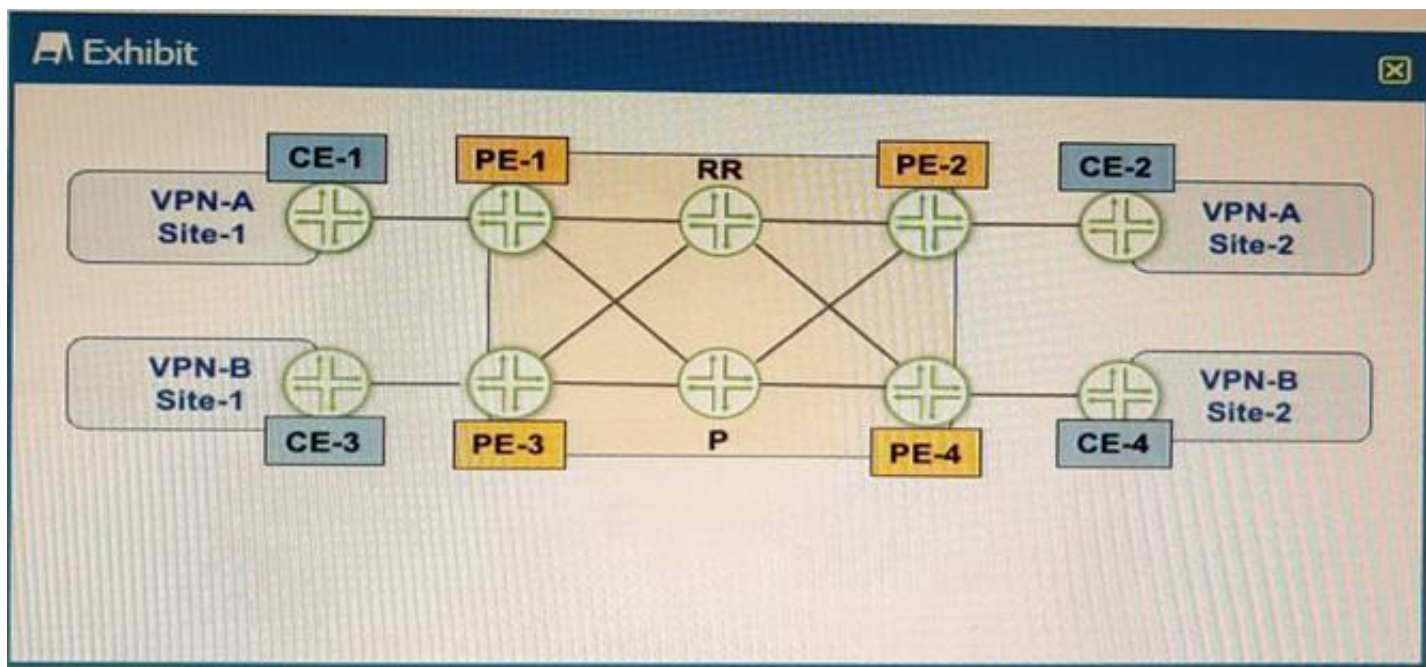
Explanation:

the default behavior for iBGP is to propagate EBGP-learned prefixes without changing the next-hop. This can cause issues if the next-hop is not reachable via the IGP. One solution is to use the next-hop self command on R4, which will change the next-hop attribute to its own loopback address. This way, R5 can reach the next-hop via the IGP and install the route in its routing table.

Another solution is to add the external interface prefix (120.0.4.16/30) to the IGP routing tables of R4 and R5. This will also make the next-hop reachable via the IGP and allow R5 to use the route. According to2, this is a possible workaround for a pure IP network, but it may not work well for an MPLS network.

NEW QUESTION 29

Exhibit



Referring to the exhibit, PE-1 and PE-2 are getting route updates for VPN-B when neither of them service that VPN. Which two actions would optimize this process? (Choose two.)

- A. Configure the family route-target statement on the PEs.
- B. Configure the family route-target statement on the RR
- C. Configure the resolution rib bgp . l3vpn . 0 resolution-ribs ine
- D. 0 Statement on the PEs.
- E. Configure the resolution rib bgp.l3vpn.0 resolution-ribs ine
- F. 0 Statement on the RR

Answer: BD

Explanation:

BGP route target filtering is a technique that reduces the number of routers that receive VPN routes and route updates, helping to limit the amount of overhead associated with running a VPN. BGP route target filtering is based on the exchange of the route-target address family, which contains information about the VPN membership of each PE device. Based on this information, a PE device can decide whether to accept or reject VPN routes from another PE device. BGP route target filtering can be configured on PE devices or on route reflectors (RRs). Configuring BGP route target filtering on RRs is more efficient and scalable, as it reduces the number of BGP sessions and updates between PE devices. To configure BGP route target filtering on RRs, the following steps are required:


? Configure the family route-target statement under the BGP group or neighbor configuration on the RRs. This enables the exchange of the route-target address family between the RRs and their clients (PE devices).

? Configure the resolution rib bgp.l3vpn.0 resolution-ribs inet.0 statement under the routing-options configuration on the RRs. This enables the RRs to resolve next hops for VPN routes using the inet.0 routing table.

? Configure an export policy for BGP route target filtering under the routing-options configuration on the RRs. This policy controls which route targets are advertised to each PE device based on their VPN membership.

NEW QUESTION 30

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: AB

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:

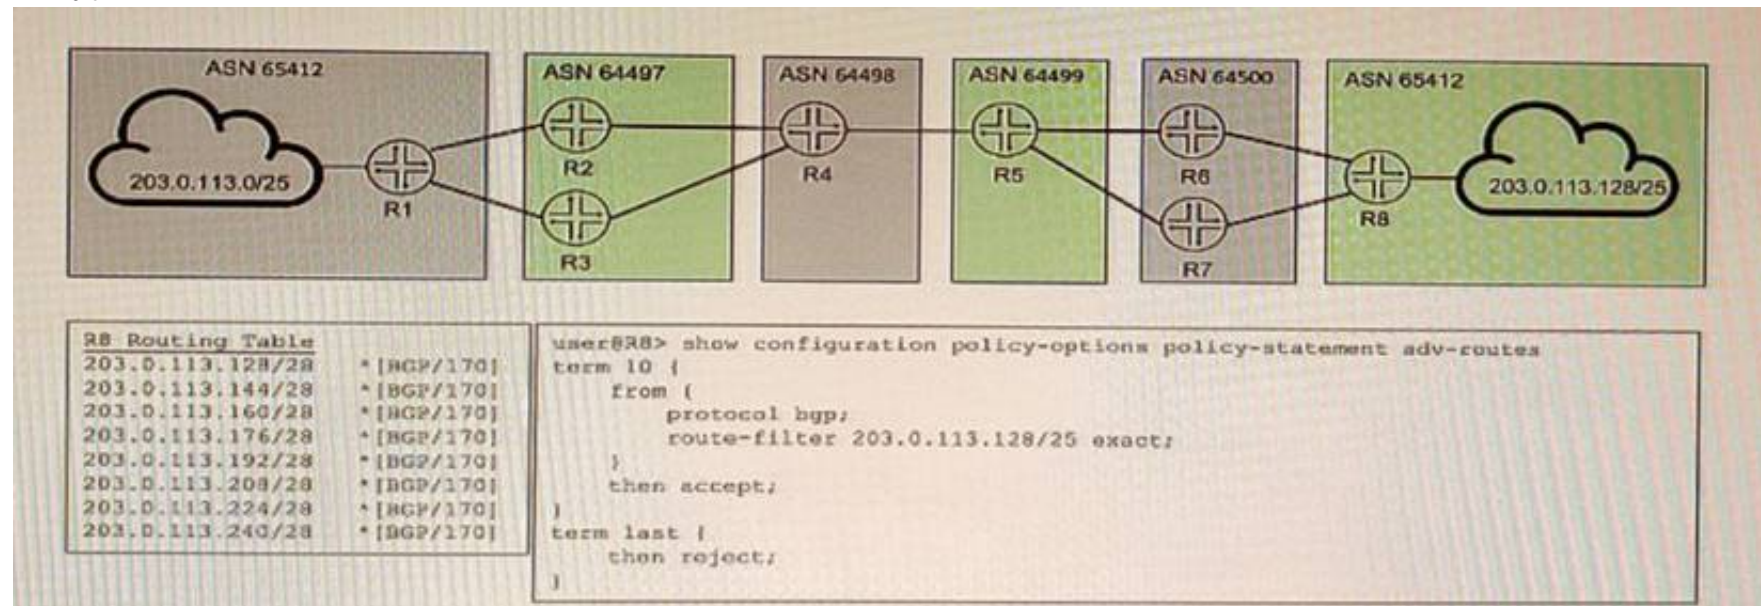
? By default, the next-hop value for these routes is not changed by ISP-A before being sent to R3. This is because the default behavior of EBGP is to preserve the next-hop attribute of the routes received from another EBGP neighbor. The next-hop attribute indicates the IP address of the router that should 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>

NEW QUESTION 35

Exhibit



You are attempting to summarize routes from the 203.0.113.128/25 IP block on R8 to AS 64500. You implement the export policy shown in the exhibit and all routes from the routing table stop being advertised.

In this scenario, which two steps would you take to summarize the route in BGP? (Choose two.)

- A. Remove the from protocol bgp command from the export policy.
- B. Add the set protocols bgp family inet unicast add-path command to allow additional routes to the RIB table
- C. -
- D. Add the set routing-options static route 203.0.113.123/25 discard command.
- E. Replace exact in the export policy with orlonger.

Answer: CD

Explanation:

To summarize routes from the 203.0.113.128/25 IP block on R8 to AS 64500, you need to do the following:

? Add the set routing-options static route 203.0.113.128/25 discard command. This creates a static route for the summary prefix and discards any traffic destined to it. This is necessary because BGP can only advertise routes that are present in the routing table.

? Replace exact in the export policy with orlonger. This allows R8 to match and advertise any route that is equal or more specific than the summary prefix. The exact term only matches routes that are exactly equal to the summary prefix, which is not present in the routing table.

NEW QUESTION 40

Exhibit


```

user@router> show l2vpn connections
Layer-2 VPN connections:
Legend for connection status (St)
EI -- encapsulation invalid          NC -- interface encapsulation not
CCC/TCC/VPLS                        WE -- interface and instance encaps not same
EM -- encapsulation mismatch        NP -- interface hardware not present
VC-Dn -- Virtual circuit down       -> -- only outbound connection is up
CM -- control-word mismatch         <- -- only inbound connection is up
CN -- circuit not provisioned        Up -- operational
OR -- out of range                  Dn -- down
OL -- no outgoing label             CF -- call admission control failure
LD -- local site signaled down       SC -- local and remote site ID collision
RD -- remote site signaled down      LM -- local site ID not minimum designated
LN -- local site not designated      RM -- remote site ID not minimum designated
RN -- remote site not designated     IL -- no incoming label
XX -- unknown connection status      MI -- Mesh-Group ID not available
MM -- MTU mismatch                  ST -- Standby connection
BK -- Backup connection              PB -- Profile busy
PF -- Profile parse failure           SN -- Static Neighbor
RS -- remote site standby            RB -- Remote site not best-site
LB -- Local site not best-site       HS -- Hot-standby Connection
VM -- VLAN ID mismatch
Legend for interface status
Up -- operational
Dn -- down
Instance: vpn-A
Edge protection: Not-Primary
Local site: CE1-2 (2)
connection-site Type St      Time last up      # Up trans
1               rmt  Up      Apr 11 14:35:27 2020      1
Remote PE: 172.17.20.1, Negotiated control-word: Yes (Null)
Incoming label: 21, Outgoing label: 22
Local interface: ge-0/0/6.610, Status: Up, Encapsulation: VLAN
Flow Label Transmit: No, Flow Label Receive: No

```

Which two statements about the output shown in the exhibit are correct? (Choose two.)

- A. The PE is attached to a single local site.
- B. The connection has not flapped since it was initiated.
- C. There has been a VLAN ID mismatch.
- D. The PE router has the capability to pop flow labels

Answer: AD


Explanation:

According to 1 and 2, BGP Layer 2 VPNs use BGP to distribute endpoint provisioning information and set up pseudowires between PE devices. BGP uses the Layer 2 VPN (L2VPN) Routing Information Base (RIB) to store endpoint provisioning information, which is updated each time any Layer 2 virtual forwarding instance (VFI) is configured. The prefix and path information is stored in the L2VPN database, which allows BGP to make decisions about the best path. In the output shown in the exhibit, we can see some information about the L2VPN RIB and the pseudowire state. Based on this information, we can infer the following statements:

- ? The PE is attached to a single local site. This is correct because the output shows only one local site ID (1) under the L2VPN RIB section. A local site ID is a unique identifier for a site within a VPLS domain. If there were multiple local sites attached to the PE, we would see multiple local site IDs with different prefixes.
- ? The connection has not flapped since it was initiated. This is correct because the output shows that the uptime of the pseudowire is equal to its total uptime (1w6d). This means that the pseudowire has been up for one week and six days without any interruption or flap.
- ? There has been a VLAN ID mismatch. This is not correct because the output shows that the remote and local VLAN IDs are both 0 under the pseudowire state section. A VLAN ID mismatch occurs when the remote and local VLAN IDs are different, which can cause traffic loss or misdelivery. If there was a VLAN ID mismatch, we would see different values for the remote and local VLAN IDs.
- ? The PE router has the capability to pop flow labels. This is correct because the output shows that the flow label pop bit is set under the pseudowire state section. The flow label pop bit indicates that the PE router can pop (remove) the MPLS flow label from the packet before forwarding it to the CE device. The flow label is an optional MPLS label that can be used for load balancing or traffic engineering purposes.

NEW QUESTION 41

Exhibit



Referring to the exhibit, a working L3VPN exists that connects VPN-A sites CoS is configured correctly to match on the MPLS EXP bits of the LSP, but when traffic is sent from Site-1 to Site-2, PE-2 is not classifying the traffic correctly

What should you do to solve the problem?

- A. Configure the explicit-null statement on PE-1.
- B. Configure the explicit-null statement on PE-2
- C. Configure VPN prefix mapping for the PE-1_to_PE-2 LSP
- D. Set a static CoS value for the PE-1_to_PE-2 LSP


Answer: A

Explanation:

The explicit-null statement enables the PE router to send an MPLS label with a value of 0 (explicit null) instead of an IP header for packets destined to the VPN customer sites. This allows the penultimate hop router (the router before the egress PE router) to preserve the EXP bits of the MPLS label and pass them to the egress PE router. The egress PE router can then use these EXP bits to classify the traffic according to the CoS policy². In this example, PE-1 should configure the explicit-null statement under [edit protocols mpls label-switched-path PE-1_to_PE-2] hierarchy level.

NEW QUESTION 45

Exhibit



Referring to the exhibit, PIM-SM is configured on all routers, and Anycast-RP with Anycast- PIM is used for the discovery mechanism on RP1 and RP2. The interface metric values are shown for the OSPF area.

In this scenario, which two statements are correct about which RP is used? (Choose two.)

- A. Source2 will use RP2 and Receiver1 will use RP2 for group 224.2.2.2.
- B. Source2 will use RP1 and Receiver2 will use RP1 for group 224.2.2.2.
- C. Source1 will use RP1 and Receiver1 will use RP1 for group 224.1.1.1.
- D. Source1 will use RP1 and Receiver1 will use RP2 for group 224.1.1.1.

Answer: AC

Explanation:

A sham link is a logical link between two PE routers that belong to the same OSPF area but are connected through an L3VPN. A sham link makes the PE routers appear as if they are directly connected, and prevents OSPF from preferring an intra-area back door link over the VPN backbone. A sham link creates an OSPF multihop neighborhood between the PE routers using TCP port 646. The PEs exchange Type 1 OSPF LSAs instead of Type 3 OSPF LSAs for the L3VPN routes, which allows OSPF to use the correct metric for route selection¹.

NEW QUESTION 50

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

JN0-664 Practice Exam Features:

- * JN0-664 Questions and Answers Updated Frequently
- * JN0-664 Practice Questions Verified by Expert Senior Certified Staff
- * JN0-664 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * JN0-664 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The JN0-664 Practice Test Here](#)