

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

You are asked to protect your company's customers from amplification attacks. In this scenario, what is Juniper's recommended protection method?

- A. ASN prepending
- B. BGP FlowSpec
- C. destination-based Remote Triggered Black Hole
- D. unicast Reverse Path Forwarding

Answer: C


Explanation:

amplification attacks are a type of distributed denial-of-service (DDoS) attack that exploit the characteristics of certain protocols to amplify the traffic sent to a victim. For example, an attacker can send a small DNS query with a spoofed source IP address to a DNS server, which will reply with a much larger response to the victim. This way, the attacker can generate a large amount of traffic with minimal resources.

One of the methods to protect against amplification attacks is destination-based Remote Triggered Black Hole (RTBH) filtering. This technique allows a network operator to drop traffic destined to a specific IP address or prefix at the edge of the network, thus preventing it from reaching the victim and consuming bandwidth and resources. RTBH filtering can be implemented using BGP to propagate a special route with a next hop of 192.0.2.1 (a reserved address) to the edge routers. Any traffic matching this route will be discarded by the edge routers.

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@RI show configuration interpolated-profile { interpolate {
fill-level [ 50 75 drop—probability [ > }
class-of-service drop-profiles
];
20 60 ];
```

Which two statements are correct about the class-of-service configuration shown in the exhibit? (Choose two.)

- A. The drop probability jumps immediately from 20% to 60% when the queue level reaches 75% full.
- B. The drop probability gradually increases from 20% to 60% as the queue level increases from 50% full to 75% full
- C. To use this drop profile, you reference it in a scheduler.
- D. To use this drop profile, you apply it directly to an interface.

Answer: BC

Explanation:

class-of-service (CoS) is a feature that allows you to prioritize and manage network traffic based on various criteria, such as application type, user group, or packet loss priority. CoS uses different components to classify, mark, queue, schedule, shape, and drop traffic according to the configured policies. One of the components of CoS is drop profiles, which define how packets are dropped when a queue is congested. Drop profiles use random early detection (RED) algorithm to drop packets randomly before the queue is full, which helps to avoid global synchronization and improve network performance. Drop profiles can be discrete or interpolated. A discrete drop profile maps a specific fill level of a queue to a specific drop probability. An interpolated drop profile maps a range of fill levels of a queue to a range of drop probabilities and interpolates the values in between.

In the exhibit, we can see that the class-of-service configuration shows an interpolated drop profile with two fill levels (50 and 75) and two drop probabilities (20 and 60). Based on this configuration, we can infer the following statements:

? The drop probability jumps immediately from 20% to 60% when the queue level reaches 75% full. This is not correct because the drop profile is interpolated, not discrete. This means that the drop probability gradually increases from 20% to 60% as the queue level increases from 50% full to 75% full. The drop probability for any fill level between 50% and 75% can be calculated by using linear interpolation formula.

? The drop probability gradually increases from 20% to 60% as the queue level increases from 50% full to 75% full. This is correct because the drop profile is interpolated and uses linear interpolation formula to calculate the drop probability for any fill level between 50% and 75%. For example, if the fill level is 60%, the drop probability is 28%, which is calculated by using the formula: $(60 - 50) / (75 - 50) * (60 - 20) + 20 = 28$.

? To use this drop profile, you reference it in a scheduler. This is correct because a scheduler is a component of CoS that determines how packets are dequeued from different queues and transmitted on an interface. A scheduler can reference a drop profile by using the random-detect statement under the [edit class-of-service schedulers] hierarchy level. For example: scheduler test { transmit-rate percent 10; buffer-size percent 10; random-detect test-profile; }

? To use this drop profile, you apply it directly to an interface. This is not correct because a drop profile cannot be applied directly to an interface. A drop profile can only be referenced by a scheduler, which can be applied to an interface by using the scheduler-map statement under the [edit class-of-service interfaces] hierarchy level. For example: interfaces ge-0/0/0 { unit 0 { scheduler-map test-map; } }

NEW QUESTION 5

Exhibit

```
[edit policy-options]
user@router# show
policy-statement block-igmp {
  term 1 {
    from {
      route-filter 224.7.7.7/32 exact;
      source-address-filter 192.168.100.10/32 exact;
    }
    then reject;
  }
}
[edit protocols igmp]
user@router# show
interface ge-0/0/0.0 {
  group-policy block-igmp;
  group-limit 25;
}
```

Based on the configuration contents shown in the exhibit, which statement is true?

- A. Joins for group 224.7.7.7 are rejected if the source address is 192.168.100.10
- B. Joins for any group are accepted if the group count value is less than 25.
- C. Joins for group 224.7.7.7 are always rejected, regardless of the group count.
- D. Joins for group 224.7.7.7 are accepted if the group count is less than 25

Answer: D

Explanation:

BGP policy framework is a set of tools that allows you to control the flow of routing information and apply routing policies based on various criteria. BGP policy framework consists of several components, such as route maps, prefix lists, community lists, AS path lists, and route filters. Route maps are used to define routing policies by matching certain conditions and applying certain actions. Prefix lists are used to filter routes based on their prefixes. Community lists are used to filter routes based on their community attributes. AS path lists are used to filter routes based on their AS path attributes. Route filters are used to filter routes based on their prefix length or range. In this question, we have a route map named ISP-A that has two clauses: clause 10 and clause 20. Clause 10 matches any route with a prefix length between 8 and 24 bits and sets the local preference to 200. Clause 20 matches any route with a prefix of 224.7.7.7/32 and rejects it. The route map is applied inbound on the BGP neighborship with ISP-A. Based on this configuration, the correct statement is that joins for group 224.7.7.7 are always rejected, regardless of the group count. This is because clause 20 explicitly denies any route with a prefix of 224.7.7.7/32, which corresponds to the multicast group 224.7.7.7.

Reference: 3: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xr-16/irg-xr-16-book/bgp-policy-framework.html

NEW QUESTION 6

Which two statements about IS-IS are correct? (Choose two.)

- A. PSNPs are flooded periodically.
- B. PSNPs contain only descriptions of LSPs.
- C. CSNPs are flooded periodically

D. CSNPs contain only descriptions of LSPs.

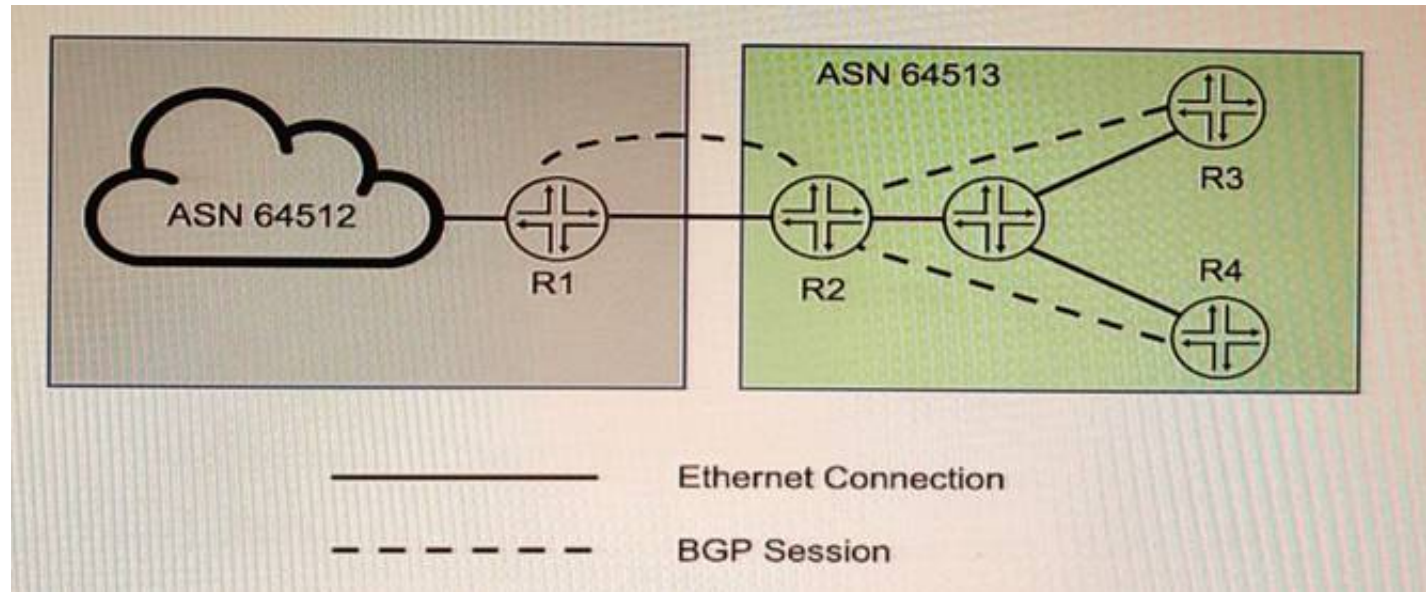
Answer: BC

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 packets to synchronize link-state databases among routers: Link State Packets (LSPs) and Partial Sequence Number Packets (PSNPs). LSPs contain information about the state and cost of links in the network, and are flooded periodically throughout the network. PSNPs are used to acknowledge receipt of LSPs and request retransmission of missing or corrupted LSPs. PSNPs contain only descriptions of LSPs, such as their sequence numbers and checksums³. IS-IS also uses another type of packet called Complete Sequence Number Packets (CSNPs), which are used to summarize the entire link-state database at regular intervals or when a new adjacency is formed. CSNPs are flooded periodically throughout the network and contain only descriptions of LSPs⁴. Therefore, PSNPs contain only descriptions of LSPs and CSNPs are flooded periodically. References: 3: <https://www.juniper.net/documentation/us/en/software/junos/routing-policy/topics/concept/routing-policy-is-is-partial-sequence-number-packet-psnp.html> 4: <https://www.juniper.net/documentation/us/en/software/junos/routing-policy/topics/concept/routing-policy-is-is-complete-sequence-number-packet-csnp.html>

NEW QUESTION 7

Exhibit



You want to implement the BGP Generalized TTL Security Mechanism (GTSM) on the network. Which three statements are correct in this scenario? (Choose three)

- A. You can implement BGP GTSM between R2, R3, and R4
- B. BGP GTSM requires a firewall filter to discard packets with incorrect TTL.
- C. You can implement BGP GTSM between R2 and R1.
- D. BGP GTSM requires a TTL of 1 to be configured between neighbors.
- E. BGP GTSM requires a TTL of 255 to be configured between neighbors.

Answer: ADE

Explanation:


BGP GTSM is a technique that protects a BGP session by comparing the TTL value in the IP header of incoming BGP packets against a valid TTL range. If the TTL value is within the valid TTL range, the packet is accepted. If not, the packet is discarded. The valid TTL range is from 255 – the configured hop count + 1 to 255. When GTSM is configured, the BGP packets sent by the device have a TTL of 255. GTSM provides best protection for directly connected EBGP sessions, but not for multihop EBGP or IBGP sessions because the TTL of packets might be modified by intermediate devices.

In the exhibit, we can see that R2, R3, and R4 are in the same AS (AS 20) and R1 is in a different AS (AS 10). Based on this information, we can infer the following statements:

- ? You can implement BGP GTSM between R2, R3, and R4. This is not correct because R2, R3, and R4 are IBGP peers and GTSM does not provide effective protection for IBGP sessions. The TTL of packets between IBGP peers might be changed by intermediate devices or routing protocols.
- ? BGP GTSM requires a firewall filter to discard packets with incorrect TTL. This is not correct because BGP GTSM does not require a firewall filter to discard packets with incorrect TTL. BGP GTSM uses TCP option 19 to negotiate GTSM capability between peers and uses TCP option 20 to carry the expected TTL value in each packet. The receiver checks the expected TTL value against the actual TTL value and discards packets with incorrect TTL values.
- ? You can implement BGP GTSM between R2 and R1. This is correct because R2 and R1 are EBGP peers and GTSM provides effective protection for directly connected EBGP sessions. The TTL of packets between directly connected EBGP peers is not changed by intermediate devices or routing protocols.
- ? BGP GTSM requires a TTL of 1 to be configured between neighbors. This is not correct because BGP GTSM requires a TTL of 255 to be configured between neighbors. The sender sets the TTL of packets to 255 and the receiver expects the TTL of packets to be 255 minus the configured hop count.
- ? BGP GTSM requires a TTL of 255 to be configured between neighbors. This is correct because BGP GTSM requires a TTL of 255 to be configured between neighbors. The sender sets the TTL of packets to 255 and the receiver expects the TTL of packets to be 255 minus the configured hop count.

NEW QUESTION 8

Exhibit



Referring to the exhibit, you are receiving the 192.168.0.0/16 route on both R3 and R4 from your EBGP neighbor. You must ensure that R1 and R2 receive both BGP routes from the route reflector.

In this scenario, which BGP feature should you configure to accomplish this behavior?

- A. add-path
- B. multihop
- C. multipath
- D. route-target

Answer: A

Explanation:

BGP add-path is a feature that allows the advertisement of multiple paths through the same peering session for the same prefix without the new paths implicitly replacing any previous paths. This behavior promotes path diversity and reduces multi-exit discriminator (MED) oscillations. BGP add-path is implemented by adding a path identifier to each path in the NLRI. The path identifier can be considered as something similar to a route distinguisher in VPNs, except that a path ID can apply to any address family. Path IDs are unique to a peering session and are generated for each network. In this question, we have a route reflector (RR) that receives two routes for the same prefix (192.168.0.0/16) from an EBGP neighbor. By default, the RR will only advertise its best path to its clients (R1 and R2). However, we want R1 and R2 to receive both routes from the RR. To achieve this, we need to configure BGP add-path on the RR and enable it to send multiple paths for the same prefix to its clients.

Reference: 3: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xr-16/irg-xe-16-book/bgp-additional-paths.html

NEW QUESTION 9

Which two EVPN route types are used to advertise a multihomed Ethernet segment? (Choose two)

- A. Type 1
- B. Type 3
- C. Type 4
- D. Type 2

Answer: AC

Explanation:

EVPN is a solution that provides Ethernet multipoint services over MPLS networks. EVPN uses BGP to distribute endpoint provisioning information and set up pseudowires between PE devices. EVPN uses different route types to convey different information in the control plane. The following are the main EVPN route types:

? Type 1 - Ethernet Auto-Discovery Route: This route type is used for network-wide messaging and discovery of other PE devices that are part of the same EVPN instance. It also carries information about the redundancy mode and load balancing algorithm of the PE devices.

? Type 2 - MAC/IP Advertisement Route: This route type is used for MAC and IP address learning and advertisement between PE devices. It also carries information about the Ethernet segment identifier (ESI) and the label for forwarding traffic to the MAC or IP address.

? Type 3 - Inclusive Multicast Ethernet Tag Route: This route type is used for broadcast, unknown unicast, and multicast (BUM) traffic forwarding. It also carries information about the multicast group and the label for forwarding BUM traffic.

? Type 4 - Ethernet Segment Route: This route type is used for multihoming scenarios, where a CE device is connected to more than one PE device. It also carries information about the ESI and the designated forwarder (DF) election process.

NEW QUESTION 10

Which origin code is preferred by BGP?

- A. Internal
- B. External
- C. Incomplete
- D. Null

Answer: C

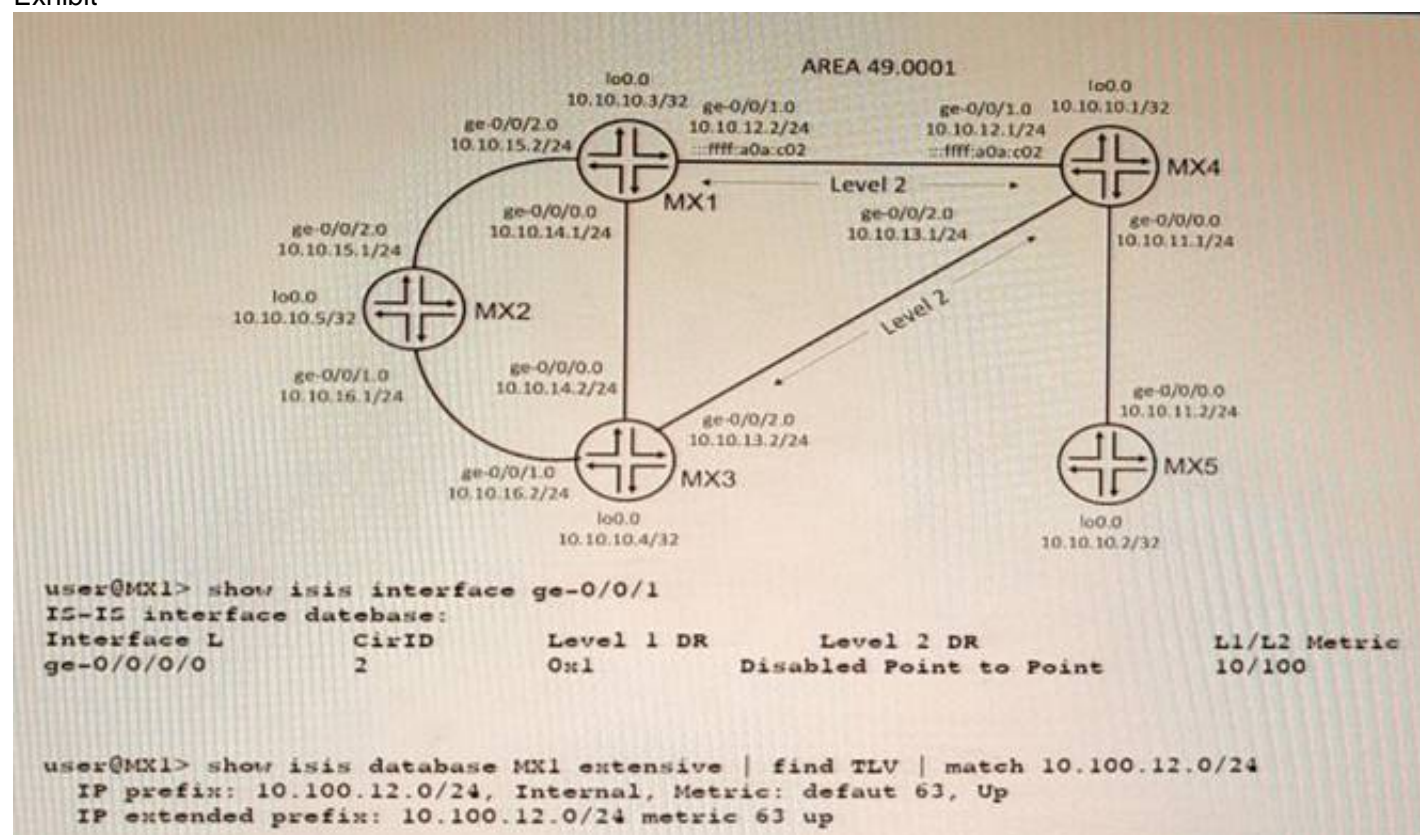
Explanation:

BGP uses several attributes to select the best path for a destination prefix. One of these attributes is origin, which indicates how BGP learned about a route. The origin attribute can have one of three values: IGP, EGP, or Incomplete. IGP means that the route was originated by a network or aggregate statement within BGP.

or by redistribution from an IGP into BGP. EGP means that the route was learned from an external BGP peer (this value is obsolete since BGP version 4). Incomplete means that the route was learned by some other means, such as redistribution from a static route into BGP. BGP prefers routes with lower origin values, so Incomplete is preferred over EGP, which is preferred over IGP.

NEW QUESTION 10

Exhibit



A network is using IS-IS for routing.
 In this scenario, why are there two TLVs shown in the exhibit?

- A. There are both narrow and wide metric devices in the topology
- B. The interface specified a metric of 100 for L2.
- C. Wide metrics have specifically been requested
- D. Both IPv4 and IPv6 are being used in the topology

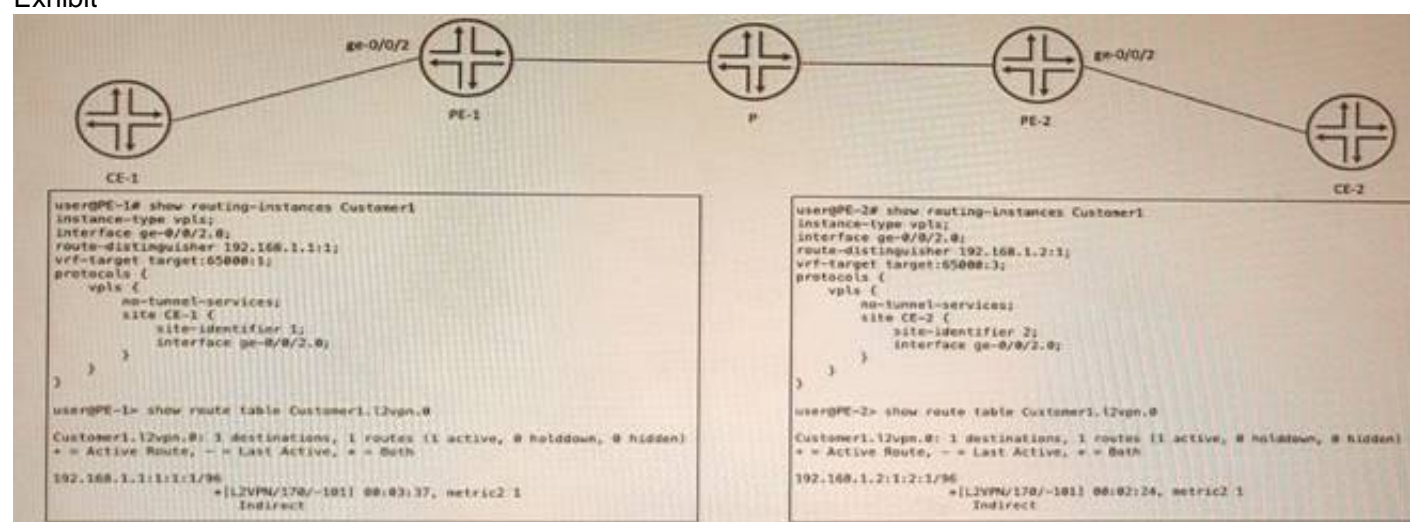
Answer: A

Explanation:

TLVs are tuples of (Type, Length, Value) that can be advertised in IS-IS packets. TLVs can carry different kinds of information in the Link State Packets (LSPs). IS-IS supports both narrow and wide metrics for link costs. Narrow metrics use a single octet to encode the link cost, while wide metrics use three octets. Narrow metrics have a maximum value of 63, while wide metrics have a maximum value of 16777215. If there are both narrow and wide metric devices in the topology, IS-IS will advertise two TLVs for each link: one with the narrow metric and one with the wide metric. This allows backward compatibility with older devices that only support narrow metrics.

NEW QUESTION 14

Exhibit



CE-1 and CE-2 are part of a VPLS called Customer1. No connectivity exists between CE-1 and CE-2. In the process of troubleshooting, you notice PE-1 is not learning any routes for this VPLS from PE-2, and PE-2 is not learning any routes for this VPLS from PE-1.

- A. The route target must match on PE-1 and PE-2.
- B. The route distinguisher must match on PE-1 and PE-2.
- C. The instance type should be changed to l2vpn.
- D. The no-tunnel-services statement should be deleted on both PEs.

Answer: A

Explanation:

VPLS is a technology that provides Layer 2 VPN services over an MPLS network. VPLS uses BGP as its control protocol to exchange VPN membership information between PE routers. The route target is a BGP extended community attribute that identifies which VPN a route belongs to. The route target must match on PE routers that participate in the same VPLS instance, otherwise they will not accept or advertise routes for that VPLS.

NEW QUESTION 19

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

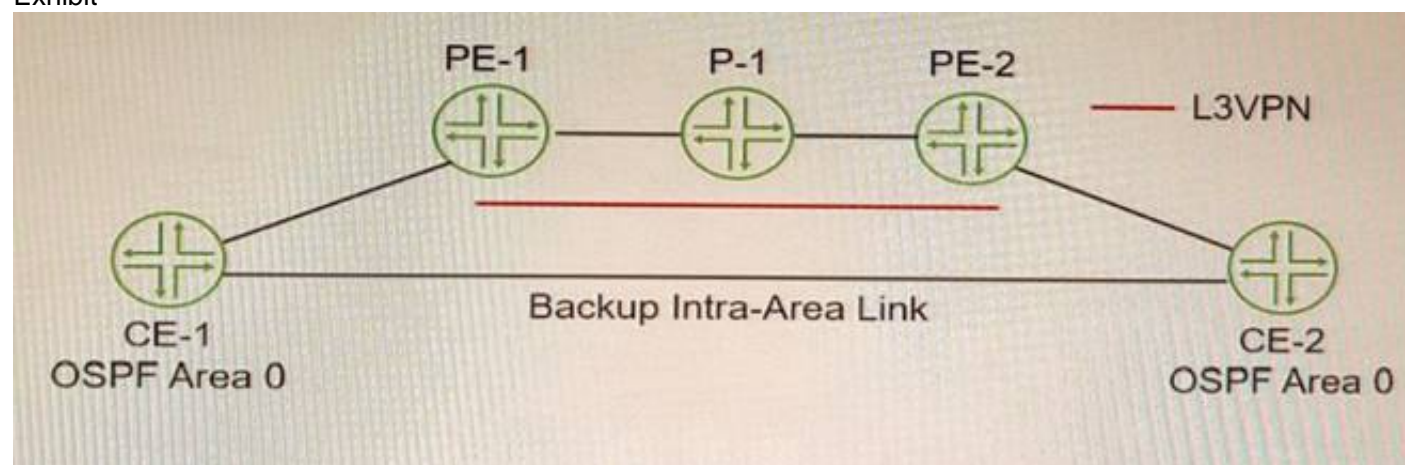
Answer: C

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 23

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 28

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 31

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: AB

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.

NEW QUESTION 33

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: AB

Explanation:

Transmission scheduling is a CoS feature that allows you to allocate bandwidth among different queues on an interface. Each queue has a configured bandwidth percentage that determines how much of the available bandwidth it can use. If a queue exceeds its allocated bandwidth, it is considered to have negative bandwidth credit and its excess traffic will be discarded by default. If a queue does not use all of its allocated bandwidth, it is considered to have positive bandwidth credit and its unused bandwidth can be shared by other queues.

NEW QUESTION 34

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:

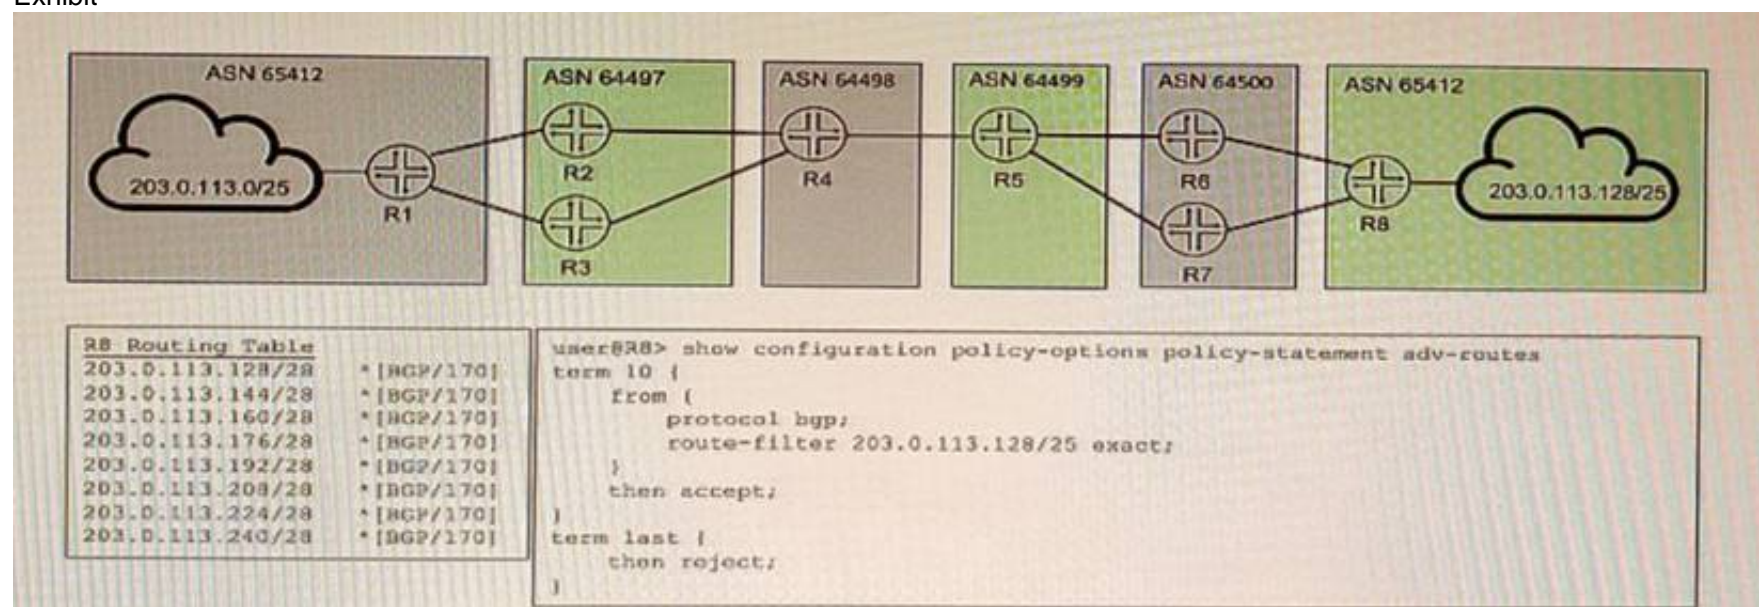
? By default, the next-hop value for these routes is not changed by ISP-A before being sent to R3. This is because the default behavior of EBGP is to preserve the next-hop attribute of the routes received from another EBGP neighbor. The next-hop attribute indicates the IP address of the router that should 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 38

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 42

Which two statements are correct about the customer interface in an LDP-signaled pseudowire? (Choose two)

- A. When the encapsulation is vlan-ccc or extended-vlan-ccc, the configured VLAN tag is not included in the control plane LDP advertisement
- B. When the encapsulation is ethernet-ccc, only frames without a VLAN tag are accepted in the data plane
- C. When the encapsulation is vLan-ccc or extended-vlan-ccc, the configured VLAN tag is included in the control plane LDP advertisement
- D. When the encapsulation is ethemet-ccc, tagged and untagged frames are both accepted in the data plane.

Answer: CD

Explanation:

The customer interface in an LDP-signaled pseudowire is the interface on the PE router that connects to the CE device. An LDP-signaled pseudowire is a type of Layer 2 circuit that uses LDP to establish a point-to-point connection between two PE routers over an MPLS network. The customer interface can have different encapsulation types depending on the type of traffic that is carried over the pseudowire. The encapsulation types are ethernet-ccc, vlan-ccc, extended-vlan-ccc, atm-ccc, frame-relay-ccc, ppp-ccc, cisco-hdlc-ccc, and tcc-ccc. Depending on the encapsulation type, the customer interface can accept or reject tagged or untagged frames in the data plane, and include or exclude VLAN tags in the control plane LDP advertisement. The following table summarizes the behavior of different encapsulation types:

NEW QUESTION 47

Which two statements are correct about a sham link? (Choose two.)

- A. It creates an OSPF multihop neighborhood between two PE routers.
- B. It creates a BGP multihop neighborhood between two PE routers.
- C. The PEs exchange Type 1 OSPF LSAs instead of Type 3 OSPF LSAs for the L3VPN routes
- D. The PEs exchange Type 3 OSPF LSAs instead of Type 1 OSPF LSAs for the L3VPN routes.

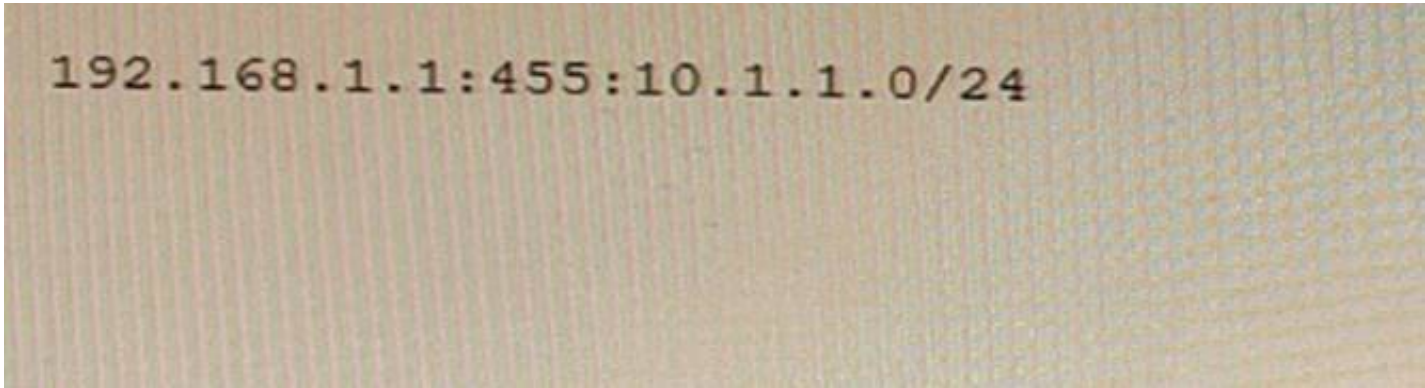
Answer: AC

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 selection.

NEW QUESTION 48

Exhibit



192.168.1.1:455:10.1.1.0/24

You are examining an L3VPN route that includes the information shown in the exhibit Which statement is correct in this scenario?

- A. The information shows a Type 1 route distinguisher.
- B. The information shows a Type 0 route distinguisher
- C. The information shows a Type 2 route distinguisher.
- D. The information shows a route target

Answer: B

Explanation:

The information shows a Type 0 route distinguisher, which is one of the three types of route distinguishers defined by RFC 4364. A route distinguisher is a 64-bit value that is prepended to an IPv4 address to create a VPN-IPv4 address, which is unique within a VPN routing and forwarding (VRF) table. A Type 0 route distinguisher has two fields: an administrator subfield (2 bytes) and an assigned number subfield (6 bytes). The administrator subfield can be an AS number or an IP address, and the assigned number subfield can be any value assigned by the administrator. In this example, the administrator subfield is 65530 (an AS number) and the assigned number subfield is 1.

NEW QUESTION 51

Which statement is correct about IS-IS when it performs the Dijkstra algorithm?

- A. The local router moves its own local tuples into the candidate database
- B. When a new neighbor ID in the tree database matches a router ID in the LSDB, the neighbor ID is moved to the candidate database
- C. Tuples with the lowest cost are moved from the tree database to the LSDB.
- D. The algorithm will stop processing once the tree database is empty.

Answer: A

Explanation:

IS-IS is a link-state routing protocol that uses the Dijkstra algorithm to compute the shortest paths between nodes in a network. The Dijkstra algorithm maintains three data structures: a tree database, a candidate database, and a link-state database (LSDB). The tree database contains the nodes that have been visited and their shortest distances from the source node. The candidate database contains the nodes that have not been visited yet and their tentative distances from the source node. The LSDB contains the topology information of the network, such as the links and their costs.

The Dijkstra algorithm works as follows:

- ? The local router moves its own local tuples into the tree database. A tuple consists of a node ID, a distance, and a parent node ID. The local router's tuple has a distance of zero and no parent node.
- ? The local router moves its neighbors' tuples into the candidate database. The neighbors' tuples have distances equal to the costs of the links to them and parent node IDs equal to the local router's node ID.
- ? The local router selects the tuple with the lowest distance from the candidate database and moves it to the tree database. This tuple becomes the current node.
- ? The local router updates the distances of the current node's neighbors in the candidate database by adding the current node's distance to the link costs. If a shorter distance is found, the parent node ID is also updated.
- ? The algorithm repeats steps 3 and 4 until either the destination node is reached or the candidate database is empty.

NEW QUESTION 56

You are responding to an RFP for a new MPLS VPN implementation. The solution must use LDP for signaling and support Layer 2 connectivity without using BGP The solution must be scalable and support multiple VPN connections over a single MPLS LSP The customer wants to maintain all routing for their Private network In this scenario, which solution do you propose?

- A. circuit cross-connect
- B. BGP Layer 2 VPN
- C. LDP Layer 2 circuit
- D. translational cross-connect

Answer: C

Explanation:

AToM (Any Transport over MPLS) is a framework that supports various Layer 2 transport types over an MPLS network core. One of the transport types supported by AToM is LDP Layer 2 circuit, which is a point-to-point Layer 2 connection that uses LDP for signaling and MPLS for forwarding. LDP Layer 2 circuit can support Layer 2 connectivity without using BGP and can be scalable and efficient by using a single MPLS LSP for multiple VPN connections. The customer can maintain all routing for their private network by using their own CE switches.

NEW QUESTION 60

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

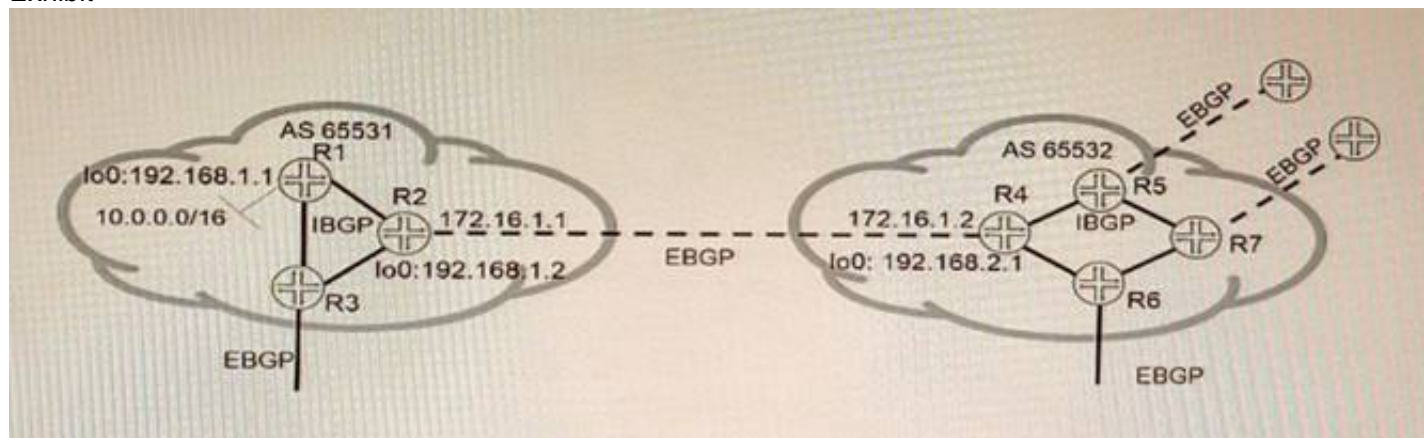
Explanation:

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 65

Exhibit



Referring to the exhibit, which three statements are correct about route 10 0 0.0/16 when using the default BGP advertisement rules'? (Choose three.)

- A. R1 will prepend AS 65531 when advertising 10 0 0.0/16 to R2.
- B. R1 will advertise 10.0.0.0/16 to R2 with 192.168.1.1 as the next hop.
- C. R2 will advertise 10.0.0.0/16 to R3 with 192.168.1.1 as the next hop
- D. R4 will advertise 10 0 0.0/16 to R6 with 172.16.1.1 as the next hop
- E. R2 will advertise 10.0.0.0/16 to R4 with 172.16.1.1 as the next hop

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 68

.....

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](#)