

Juniper

Exam Questions JN0-664

Service Provider - Professional (JNCIP-SP)



NEW QUESTION 1

You are configuring a BGP signaled Layer 2 VPN across your MPLS enabled core network. In this scenario, which statement is correct?

- A. You must assign a unique site number to each attached site's configuration.
- B. This type of VPN only supports Ethernet interfaces when connecting to CE devices.
- C. This type of VPN requires the support of the inet-vpn NLRI on all core BGP devices
- D. You must use the same route-distinguisher value on both PE devices.

Answer: C

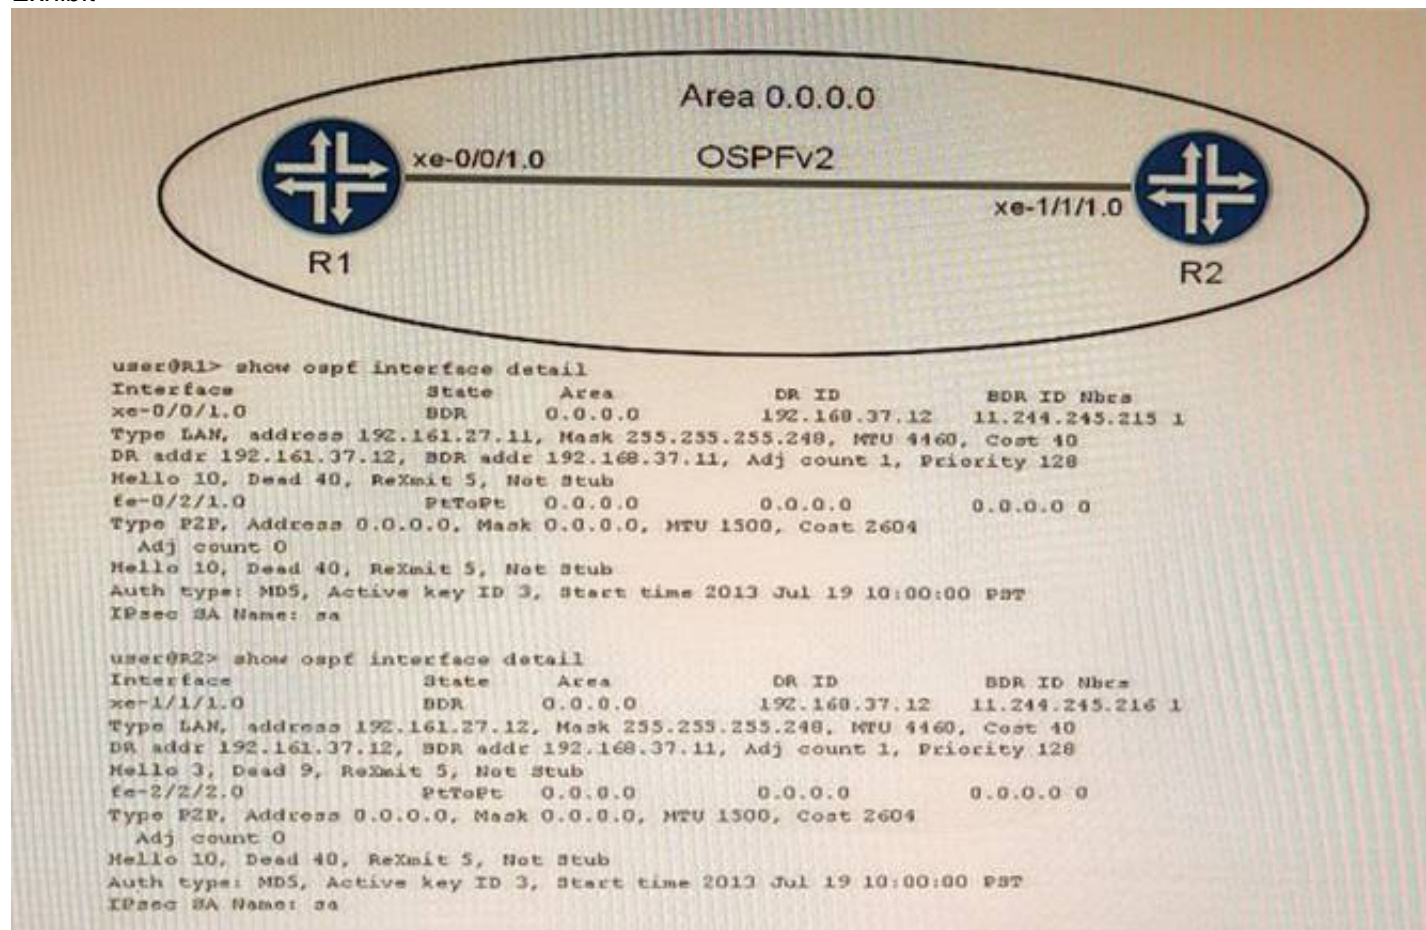
Explanation:

BGP signaled Layer 2 VPN is a type of VPN that uses BGP to distribute VPN labels and information for Layer 2 connectivity between sites over an MPLS network. BGP signaled Layer 2 VPN requires the support of the l2vpn NLRI on all core BGP devices¹. The l2vpn NLRI is a new address family that carries Layer 2 VPN information such as the VPN identifier, the attachment circuit identifier, and the route distinguisher. The l2vpn NLRI is used for both auto-discovery and signaling of Layer 2 VPNs². In this scenario, we are configuring a BGP signaled Layer 2 VPN across an MPLS enabled core network.

Therefore, we need to ensure that all core BGP devices support the l2vpn NLRI. References: 1: <https://www.juniper.net/documentation/us/en/software/junos/vpn-l2/topics/concept/vpn-layer-2-overview.html> 2: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l2_vpns/configuration/xr-16/mp-l2-vpns-xr-16-book/vpls-bgp-signaling-l2vpn-inter-as-option-a.html

NEW QUESTION 2

Exhibit



Which two statements are true about the OSPF adjacency displayed in the exhibit? (Choose two.)

- A. There is a mismatch in the hello interval parameter between routers R1 and R2
- B. There is a mismatch in the dead interval parameter between routers R1 and R2.
- C. There is a mismatch in the OSPF hold timer parameter between routers R1 and R2.
- D. There is a mismatch in the poll interval parameter between routers R1 and R2.

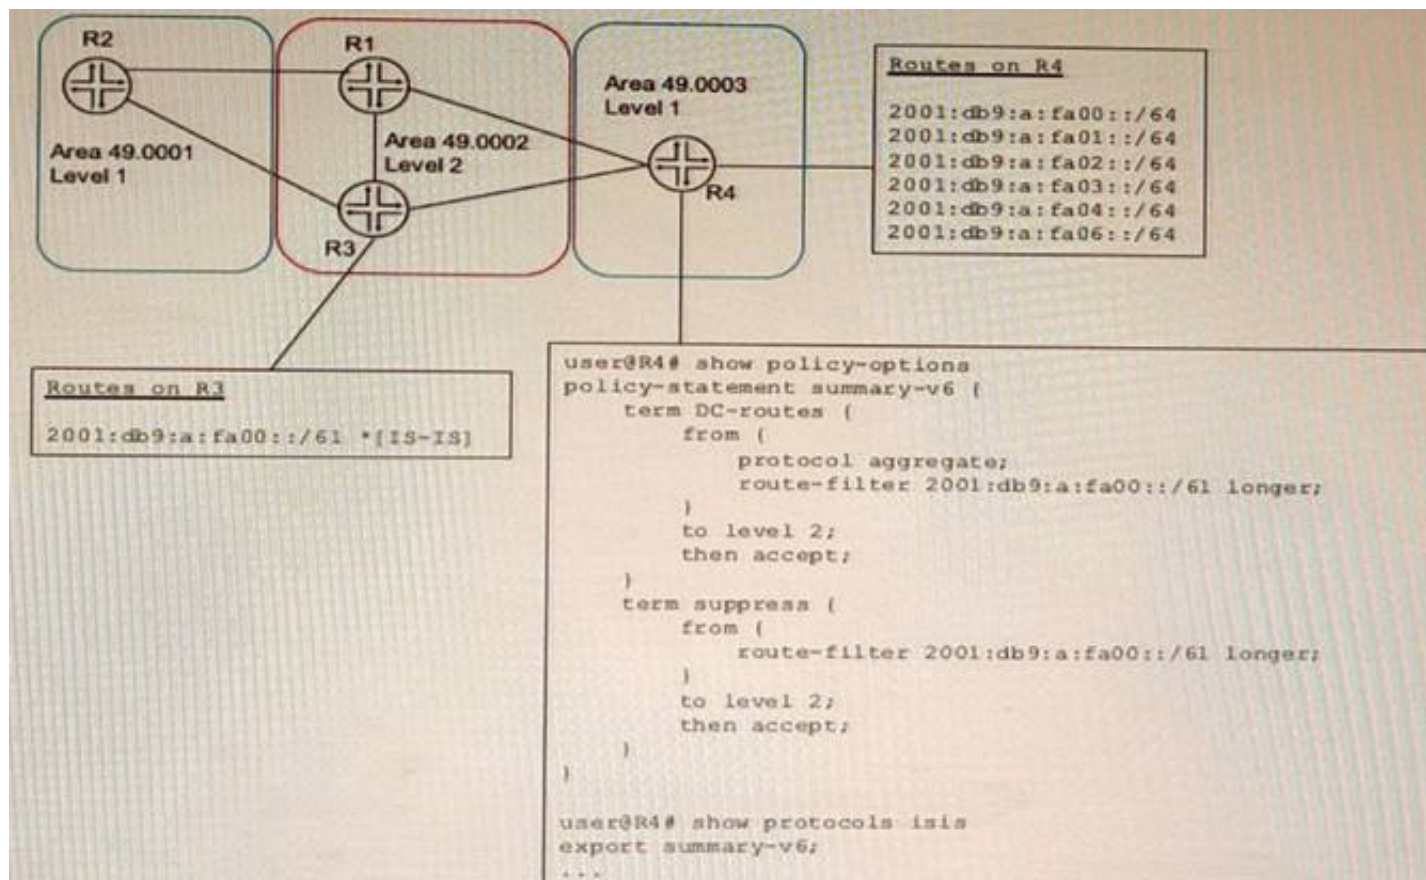
Answer: AB

Explanation:

The hello interval is the time interval between two consecutive hello packets sent by an OSPF router on an interface. The dead interval is the time interval after which a neighbor is declared down if no hello packets are received from it. These parameters must match between two OSPF routers for them to form an adjacency. In the exhibit, router R1 has a hello interval of 10 seconds and a dead interval of 40 seconds, while router R2 has a hello interval of 30 seconds and a dead interval of 120 seconds. This causes a mismatch and prevents them from becoming neighbors²³.

NEW QUESTION 3

Exhibit



A network designer would like to create a summary route as shown in the exhibit, but the configuration is not working.

Which three configuration changes will create a summary route? (Choose three.)

- A. set policy-options policy-statement leak-v6 term DC-routes then reject
- B. delete policy-options policy-statement leak-v6 term DC-routes from route-filter 2001: db9:a: fa00 : :/61 longer
- C. set policy—options policy-statement leak-v€ term DC—routes from route-filter 2001:db9:a:faOO::/61 exact
- D. delete protocols isis export summary-v6
- E. set protocols isis import summary-v6

Answer: BCD

Explanation:

To create a summary route for IS-IS, you need to configure a policy statement that matches the prefixes to be summarized and sets the next-hop to discard. You also need to configure a summary-address statement under the IS-IS protocol hierarchy that references the policy statement. In this case, the policy statement leak-v6 is trying to match the prefix 2001:db9:a:fa00::/61 exactly, but this prefix is not advertised by any router in the network. Therefore, no summary route is created. To fix this, you need to delete the longer keyword from the route-filter term and change the prefix length to /61 exact. This will match any prefix that falls within the /61 range. You also need to delete the export statement under protocols isis, because this will export all routes that match the policy statement to other IS-IS routers, which is not desired for a summary route.

NEW QUESTION 4

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:

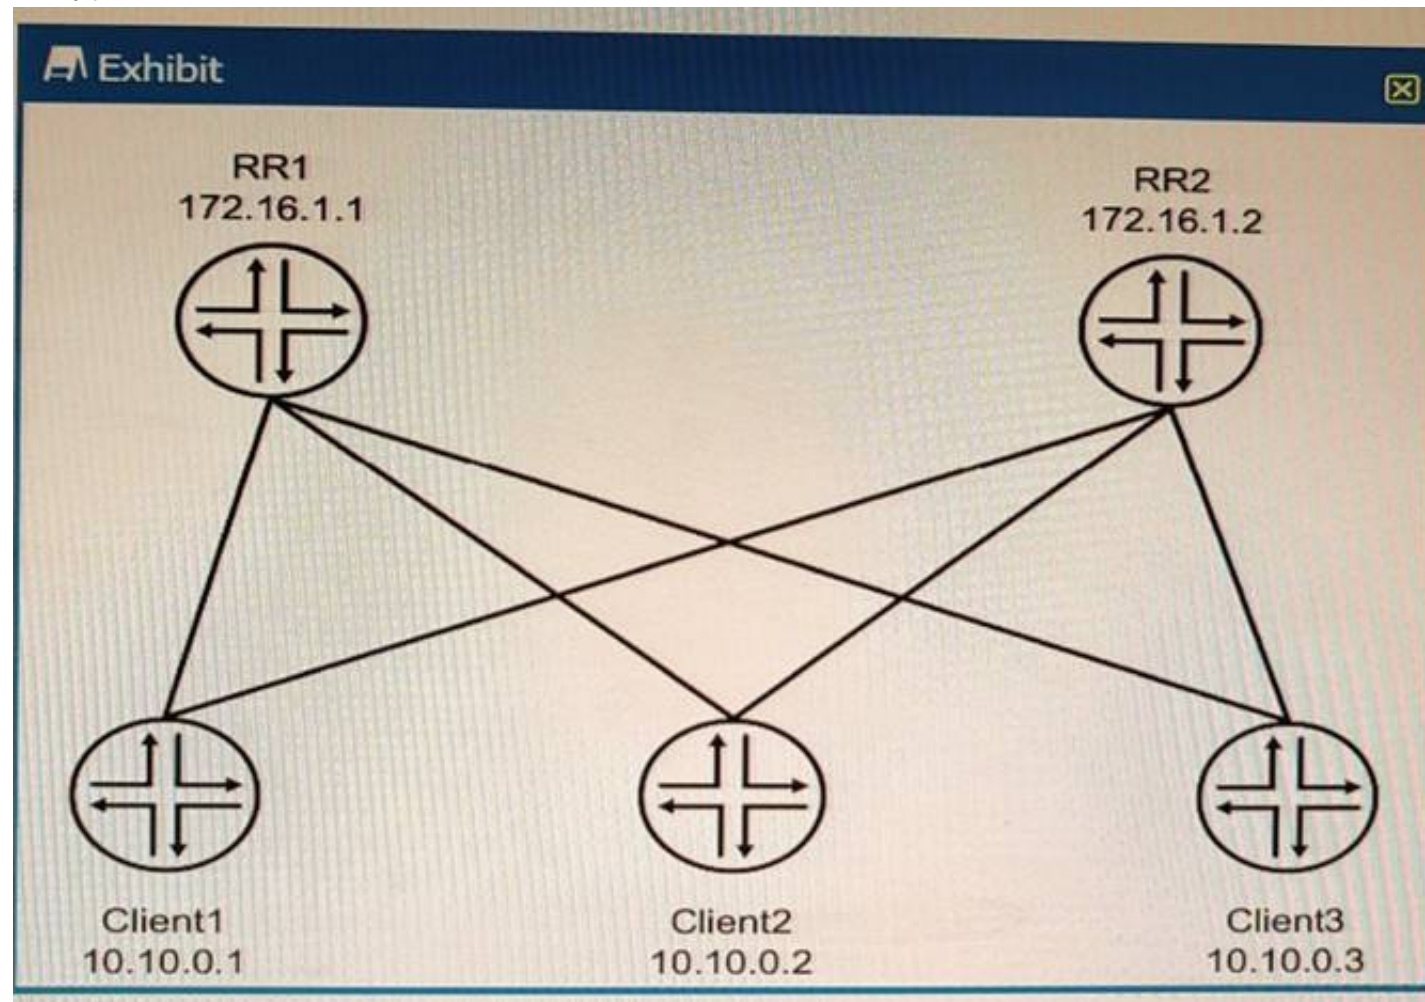
BGP policy framework is a set of tools that allows you to control the flow of routing information and apply routing policies based on various criteria. BGP policy framework consists of several components, such as route maps, prefix lists, community lists, AS path lists, and route filters. Route maps 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 neighborhood 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 5

Exhibit



The environment is using BGP All devices are in the same AS with reachability redundancy Referring to the exhibit, which statement is correct?

- A. RR1 is peered to Client2 and RR2
- B. RR2 is in an OpenConfirm State until RR1 becomes unreachable.
- C. Client1 is peered to Client2 and Client3.
- D. Peering is dynamically discovered between all devices.

Answer: A

Explanation:

BGP route reflectors are BGP routers that are allowed to ignore the IBGP loop avoidance rule and advertise IBGP learned routes to other IBGP peers under specific conditions. BGP route reflectors can reduce the number of IBGP sessions and updates in a network by eliminating the need for a full mesh of IBGP peers. BGP route reflectors can have three types of peerings:

? EBGP neighbor: A BGP router that belongs to a different autonomous system (AS) than the route reflector.

? IBGP client neighbor: An IBGP router that receives reflected routes from the route reflector. A client does not need to peer with other clients or non-clients.

? IBGP non-client neighbor: An IBGP router that does not receive reflected routes from the route reflector. A non-client needs to peer with other non-clients and the route reflector.

In the exhibit, we can see that RR1 and RR2 are route reflectors in the same AS with reachability redundancy. They have two types of peerings: EBGP neighbors (R1 and R4) and IBGP client neighbors (Client1, Client2, and Client3). RR1 and RR2 are also peered with each other as IBGP non-client neighbors.

NEW QUESTION 6

Exhibit


```

user@R4> show pim rps
Instance: PIM.master
address-family INET
RP address      Type      Mode      Holdtime Timeout Groups Group prefixes
10.1.255.2      bootstrap sparse    150       118      0 224.1.1.0/24
10.1.255.3      bootstrap sparse    150       118      2 224.1.1.0/28
user@R4> show route 10.1.255.2
inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.1.255.2/32    *{IS-IS/18} 00:32:27, metric 10
                  > to 10.1.1.2 via ge-0/0/0.0
inet.2: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
0.0.0.0/0        *{Static/5} 00:13:55
                  > to 10.1.1.6 via ge-0/0/1.0
user@R4> show route 10.1.255.3

inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.1.255.3/32    *{IS-IS/18} 00:32:43, metric 10
                  > to 10.1.1.6 via ge-0/0/1.0
inet.2: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
0.0.0.0/0        *{Static/5} 00:14:25
                  > to 10.1.1.6 via ge-0/0/1.0
[edit]
user@R2# show protocols pim
rp {
  bootstrap {
    family inet {
      priority 200;
    }
  }
  local {
    address 10.1.255.2;
    group-ranges {
      224.1.1.0/24;
    }
  }
}
interface all;
[edit]
user@R3# show protocols pim
rp {
  bootstrap {
    family inet {
      priority 210;
    }
  }
  local {
    address 10.1.255.3;
    group-ranges {
      224.1.1.0/28;
    }
  }
}
interface all;

```

R4 is directly connected to both RPs (R2 and R3) R4 is currently sending all joins upstream to R3 but you want all joins to go to R2 instead Referring to the exhibit, which configuration change will solve this issue?

- A. Change the bootstrap priority on R2 to be higher than R3
- B. Change the default route in inet.2 on R4 from R3 as the next hop to R2
- C. Change the local address on R2 to be higher than R3.
- D. Change the group-range to be more specific on R2 than R3.

Answer: A

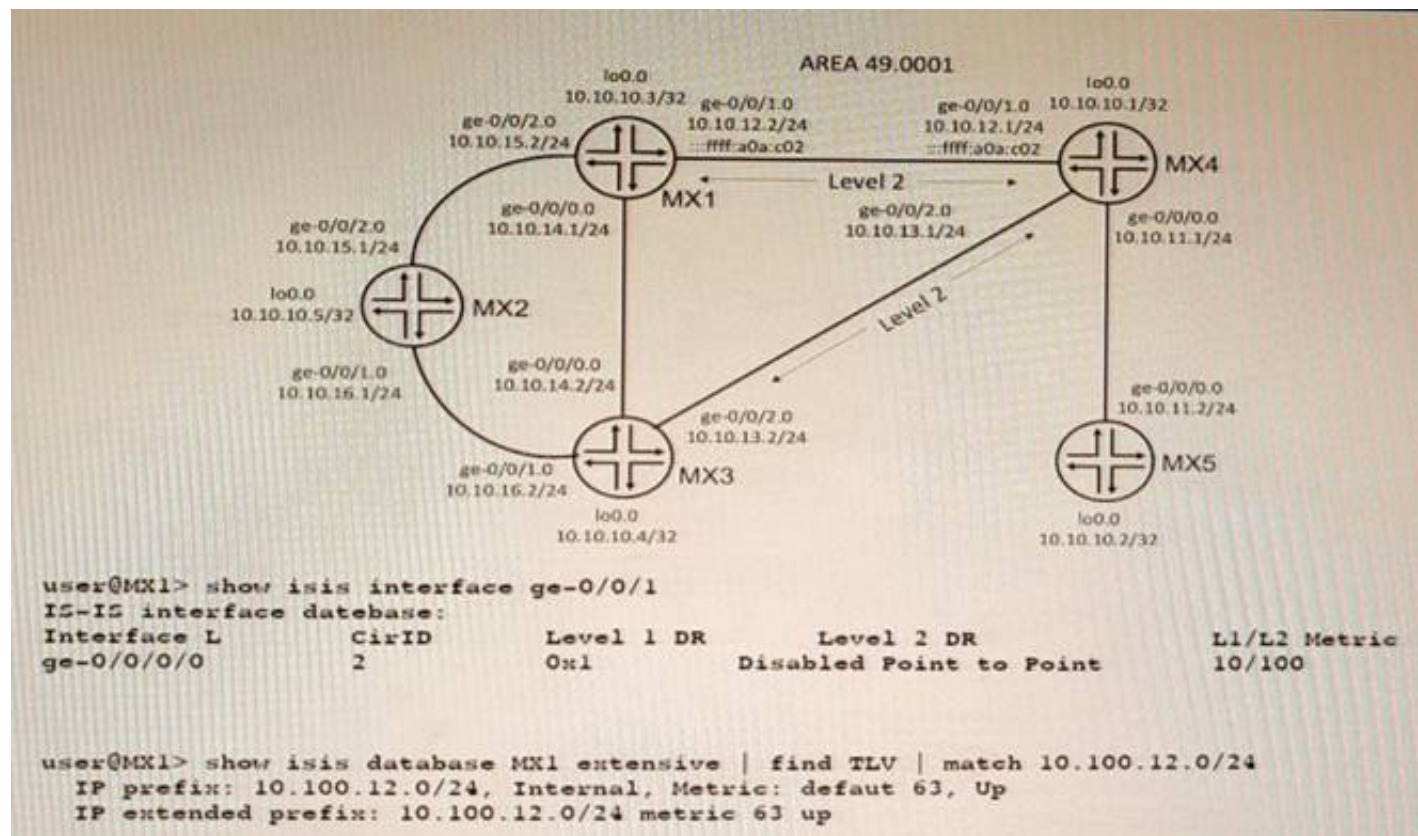
Explanation:

PIM Bootstrap Router (BSR) is a mechanism that allows PIM routers to discover and announce rendezvous point (RP) information for multicast groups. BSR uses two roles: candidate BSR and candidate RP. Candidate BSR is the router that collects information from all available RPs in the network and advertises it throughout the network. Candidate RP is the router that wants to become the RP and registers itself with the BSR. There can be only one active BSR in the network, which is elected based on the highest priority or highest IP address if the priority is the same. The BSR priority can be configured manually or assigned automatically. The default priority is 0 and the highest priority is 2515. In this question, R4 is directly connected to both RPs (R2 and R3) and is currently sending all joins upstream to R3 but we want all joins to go to R2 instead. To achieve this, we need to change the BSR priority on R2 to be higher than R3 so that R2 becomes the active BSR and advertises its RP information to R4.

Reference: 1: <https://study-ccnp.com/multicast-rendezvous-points-explained/>

NEW QUESTION 7

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

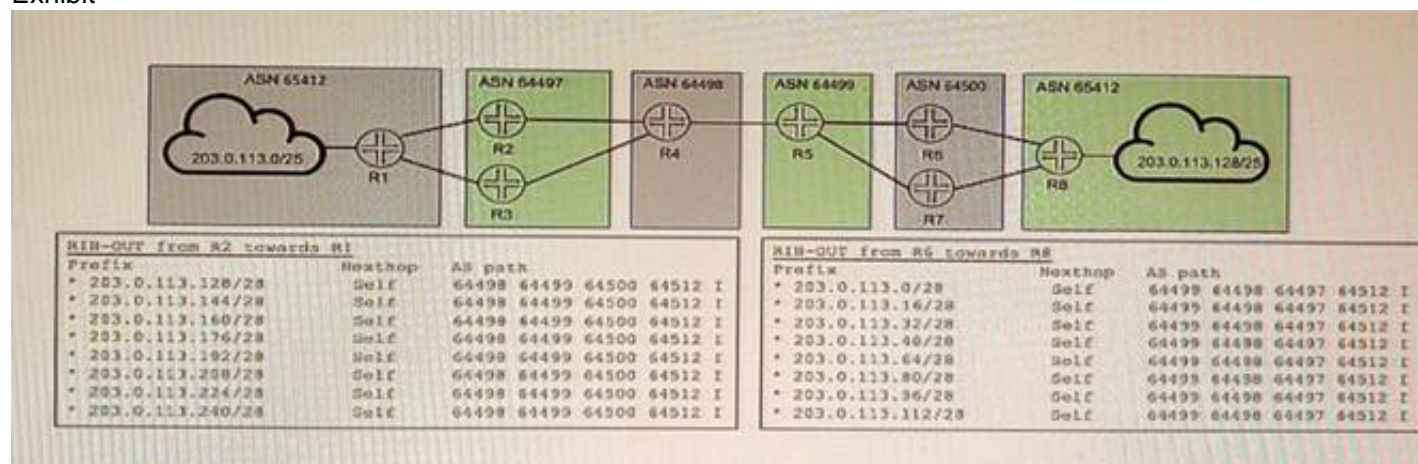
Answer: A

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 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 R6.

? 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, respectively.

NEW QUESTION 9

What is the correct order of packet flow through configurable components in the Junos OS CoS features?

- A. Multifield Classifier -> Behavior Aggregate Classifier -> Input Policer -> Forwarding Policy Options -> Fabric Scheduler -> Output Policer -> Rewrite Marker ->

Scheduler/Shaper/RED

B. Behavior Aggregate Classifier -> Multifield Classifier -> Input Policer -> Forwarding Policy Options -> Fabric Scheduler -> Output Policer ->

Scheduler/Shaper/RED -> Rewrite Marker

C. Behavior Aggregate Classifier -> Input Policer -> Multifield Classifier -> Forwarding Policy Options -> Fabric Scheduler -> Output Policer ->

Scheduler/Shaper/RED -> Rewrite Marker

D. Behavior Aggregate Classifier -> Multifield Classifier -> Input Policer -> Forwarding Policy Options -> Fabric Scheduler -> Scheduler/Shaper/RED -> Output Policer -> Rewrite Marker

Answer: C

Explanation:

The correct order of packet flow through configurable components in the Junos OS CoS features is as follows:

? Behavior Aggregate Classifier: This component uses a single field in a packet header to classify traffic into different forwarding classes and loss priorities based on predefined or user-defined values.

? Input Policer: This component applies rate-limiting and marking actions to incoming traffic based on the forwarding class and loss priority assigned by the classifier.

? Multifield Classifier: This component uses multiple fields in a packet header to classify traffic into different forwarding classes and loss priorities based on user-defined values and filters.

? Forwarding Policy Options: This component applies actions such as load balancing, filtering, or routing to traffic based on the forwarding class and loss priority assigned by the classifier.

? Fabric Scheduler: This component schedules traffic across the switch fabric based on the forwarding class and loss priority assigned by the classifier.

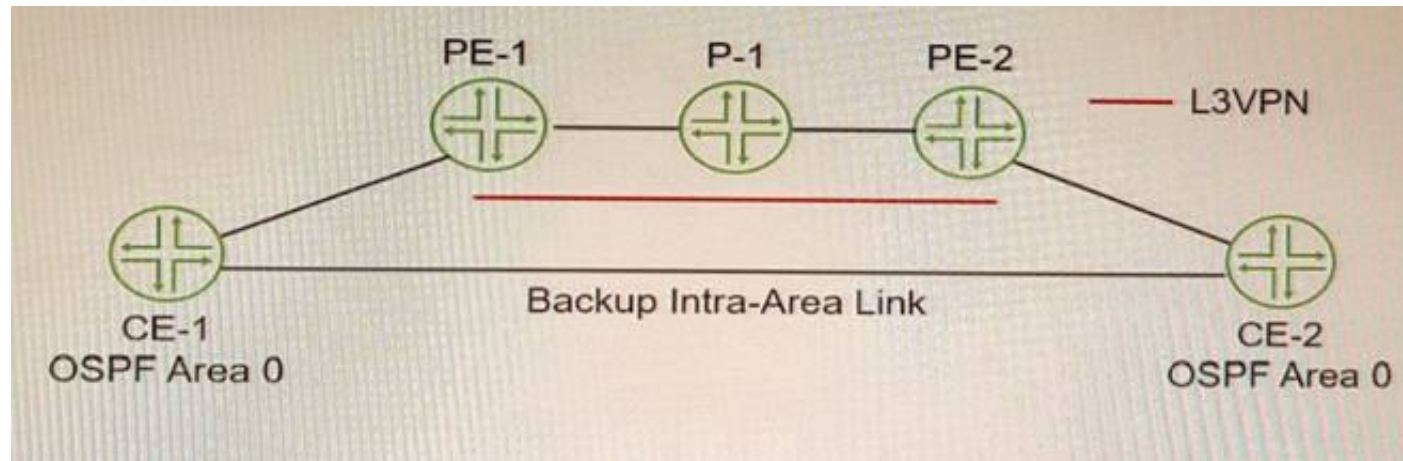
? Output Policer: This component applies rate-limiting and marking actions to outgoing traffic based on the forwarding class and loss priority assigned by the classifier.

? Scheduler/Shaper/RED: This component schedules, shapes, and drops traffic at the egress interface based on the forwarding class and loss priority assigned by the classifier.

? Rewrite Marker: This component rewrites the code-point bits of packets leaving an interface based on the forwarding class and loss priority assigned by the classifier.

NEW QUESTION 10

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 10

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 12

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

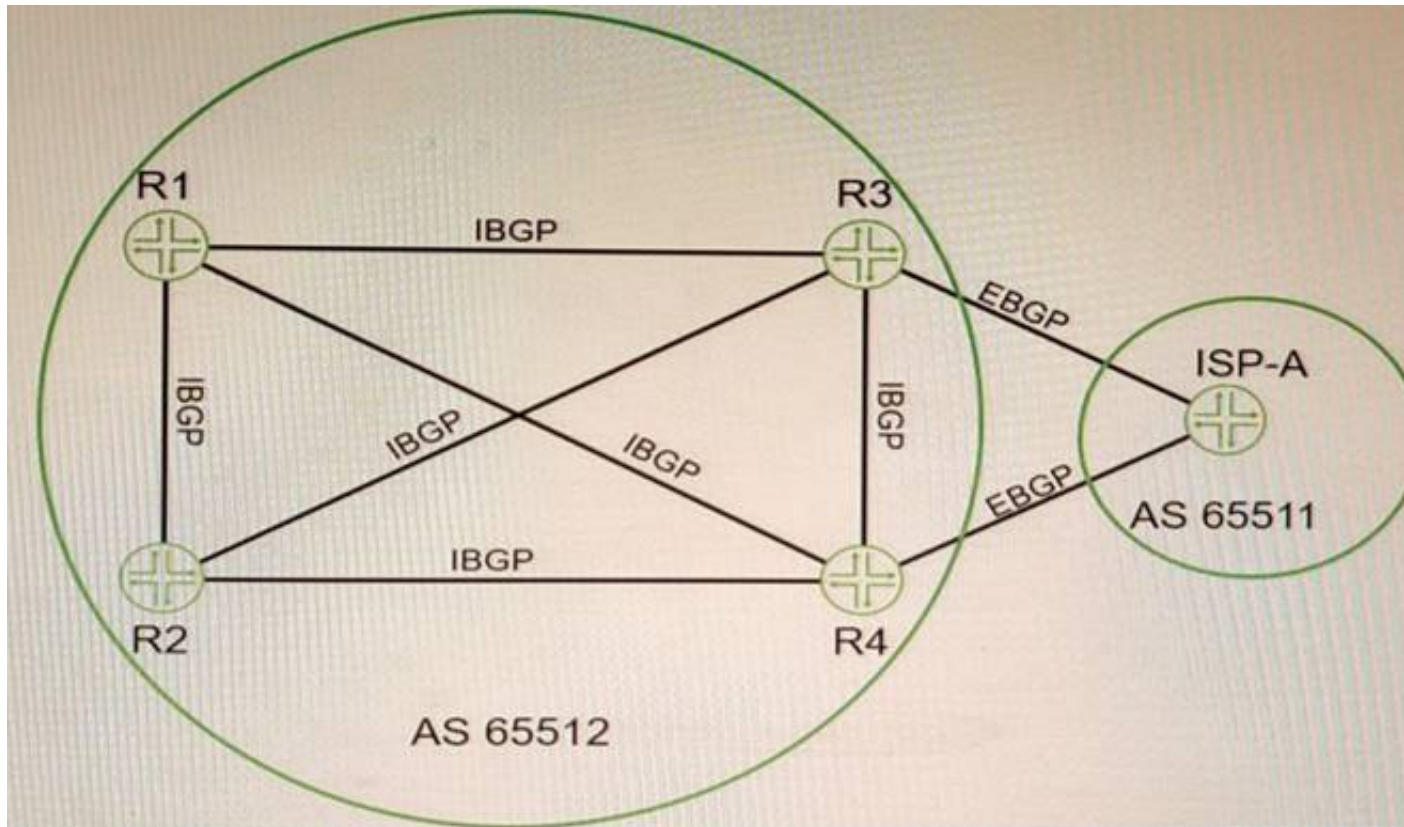
Answer: AB

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 13

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 15

You are a network architect for a service provider and want to offer Layer 2 services to your customers. You want to use EVPN for Layer 2 services in your existing MPLS network.

Which two statements are correct in this scenario? (Choose two.)

- A. Segment routing must be configured on all PE routers.

- B. VXLAN must be configured on all PE routers.
- C. EVPN uses Type 2 routes to advertise MAC address and IP address pairs learned using ARP snooping
- D. EVPN uses Type 3 routes to join a multicast tree to flood traffic.

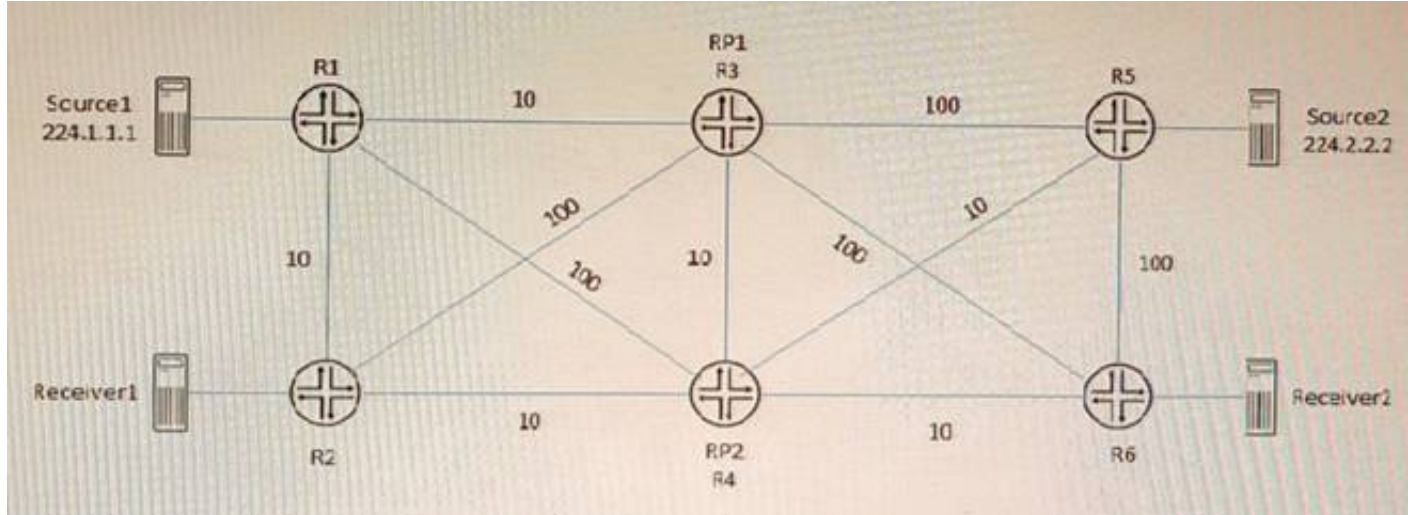
Answer: CD

Explanation:

EVPN is a technology that connects L2 network segments separated by an L3 network using a virtual Layer 2 network overlay over the Layer 3 network. EVPN uses BGP as its control protocol to exchange different types of routes for different purposes. Type 2 routes are used to advertise MAC address and IP address pairs learned using ARP snooping from the local CE devices. Type 3 routes are used to join a multicast tree to flood traffic such as broadcast, unknown unicast, and multicast (BUM) traffic.

NEW QUESTION 19

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

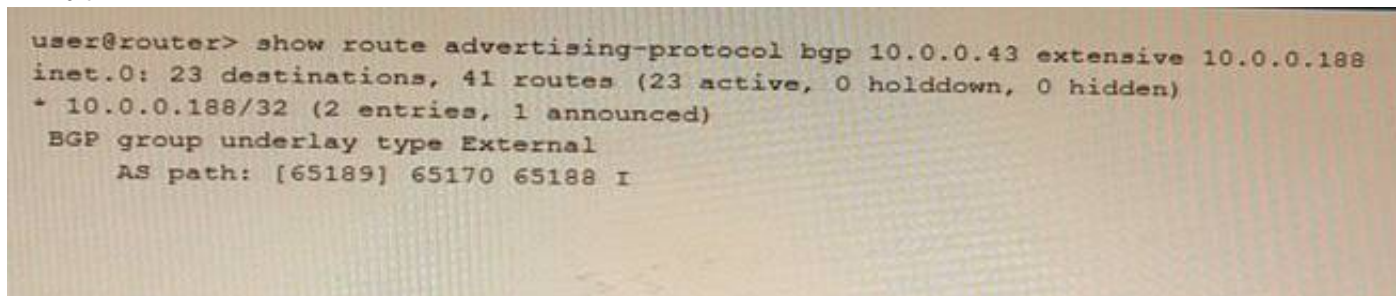
Answer: AC

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 selection.

NEW QUESTION 24

Exhibit



Referring to the exhibit, what do the brackets [] in the AS path identify?

- A. They identify the local AS number associated with the AS path if configured on the router, or if AS path prepending is configured
- B. They identify an AS set, which are groups of AS numbers in which the order does not matter
- C. They identify that the autonomous system number is incomplete and awaiting more information from the BGP protocol.
- D. They identify that a BGP confederation is being used to ensure that there are no routing loops.

Answer: B

Explanation:

The brackets [] in the AS path identify an AS set, which are groups of AS numbers in which the order does not matter. An AS set is used when BGP aggregates routes from different ASs into a single prefix. For example, if BGP aggregates routes 10.0.0.0/16 and 10.1.0.0/16 from AS 100 and AS 200, respectively, into a single prefix 10.0.0.0/15, then the AS path for this prefix will be [100 200]. An AS set reduces the length of the AS path and prevents routing loops.

NEW QUESTION 28

.....

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](#)