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-unica3t 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

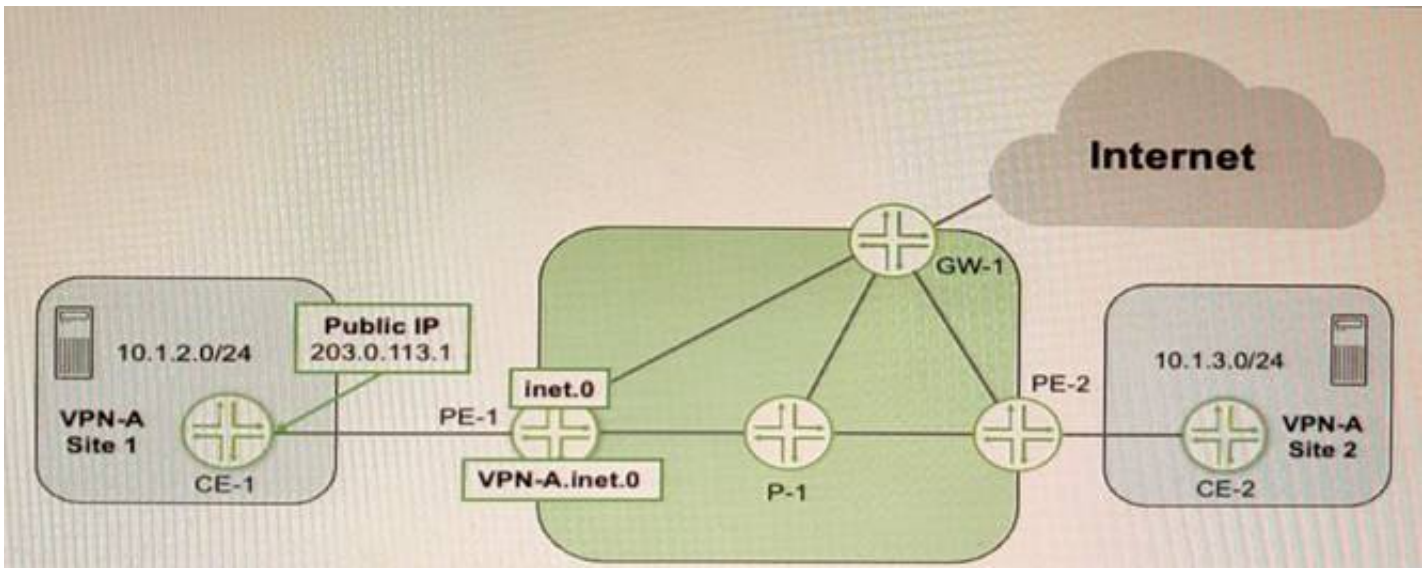

Referring to the exhibit, CE-1 is providing NAT services for the hosts at Site 1 and you must provide Internet access for those hosts
Which two statements are correct in this scenario? (Choose two.)

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
Which three mechanisms are used by Junos platforms to evaluate incoming traffic for CoS purposes? (Choose three)

A. rewrite rules
B. behavior aggregate classifiers
C. traffic shapers
D. fixed classifiers
E. multifield classifiers

Answer: BDE

Explanation:
Junos platforms use different mechanisms to evaluate incoming traffic for CoS purposes, such as:
? Behavior aggregate classifiers: These classifiers use a single field in a packet header to classify traffic into different forwarding classes and loss priorities based on predefined or user-defined values.
? Fixed classifiers: These classifiers use a fixed field in a packet header to classify traffic into different forwarding classes and loss priorities based on predefined values.
? Multifield classifiers: These classifiers use multiple fields in a packet header to classify traffic into different forwarding classes and loss priorities based on user-defined values and filters.
Rewrite rules and traffic shapers are not used to evaluate incoming traffic for CoS purposes, but rather to modify or shape outgoing traffic based on CoS policies.

NEW QUESTION 4
Exhibit

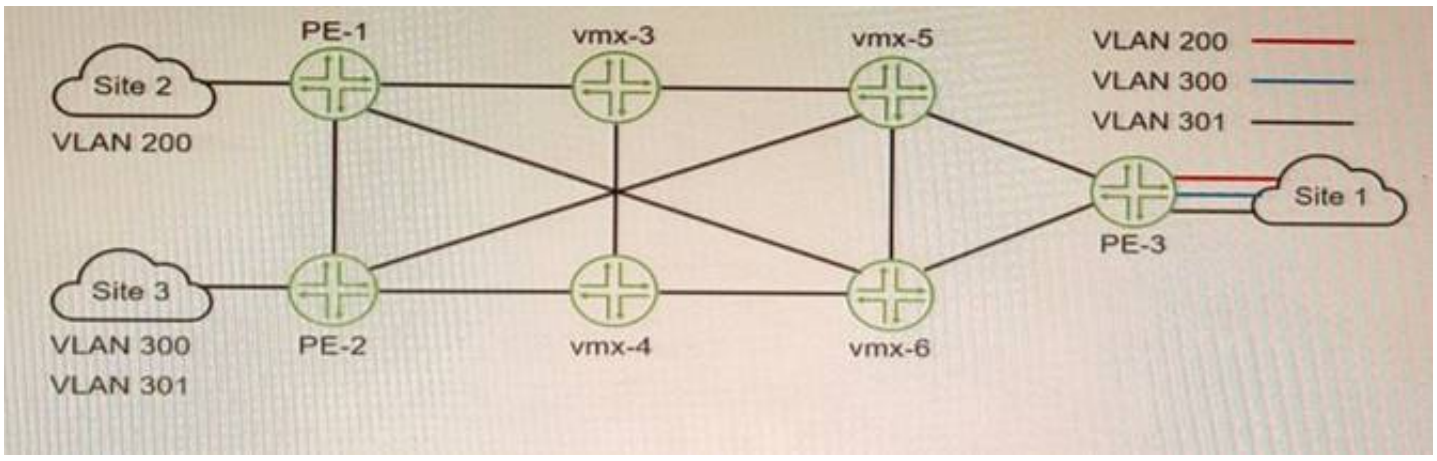

You want Site 1 to access three VLANs that are located in Site 2 and Site 3 The customer- facing interface on the PE-1 router is configured for Ethernet-VLAN encapsulation.
What is the minimum number of L2VPN routing instances to be configured to accomplish this task?

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 5
Which two statements are correct about VPLS tunnels? (Choose two.)

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

Answer: BC

Explanation:
VPLS is a Layer 2 VPN technology that allows multiple sites to connect over a shared IP/MPLS network as if they were on the same LAN. VPLS tunnels can be signaled using either Label Distribution Protocol (LDP) or Border Gateway Protocol (BGP). LDP-signaled VPLS tunnels use auto-discovery to provision sites, meaning that PE routers can automatically discover other PE routers that belong to the same VPLS instance

NEW QUESTION 6
Exhibit

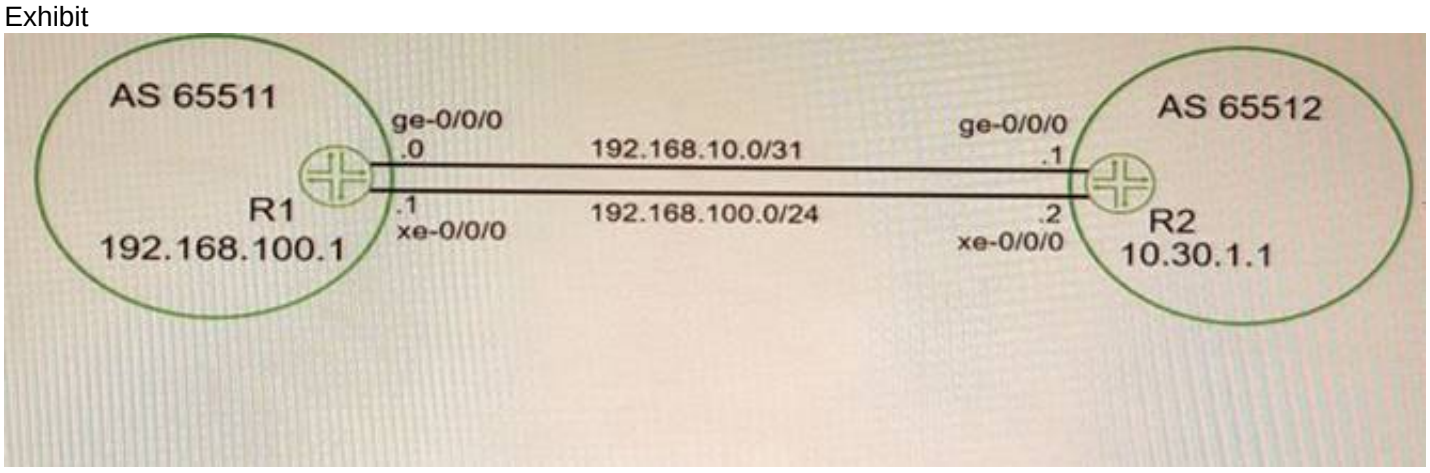

You want to use both links between R1 and R2 Because of the bandwidth difference between the two links, you must ensure that the links are used as much as possible.
Which action will accomplish this goal?

A. Define a policy to tag routes with the appropriate bandwidth community.
B. Disable multipath.
C. Ensure that the metric-out parameter on the Gigabit Ethernet interface is higher than the 10 Gigibit Ethernet interface.
D. Enable per-prefix load balancing.

Answer: D

Explanation:
VPLS is a Layer 2 VPN technology that allows multiple sites to connect over a shared IP/MPLS network as if they were on the same LAN. VPLS tunnels can be signaled using either Label Distribution Protocol (LDP) or Border Gateway Protocol (BGP). In this question, we have two links between R1 and R2 with different bandwidths (10 Gbps and 1 Gbps). We want to use both links as much as possible for VPLS traffic. To achieve this, we need to enable per-prefix load balancing on both routers. Per-prefix load balancing is a feature that allows a router to distribute traffic across multiple equal-cost or unequal- cost paths based on the destination prefix of each packet. This improves the utilization of multiple links and provides better load sharing than per-flow load balancing, which distributes traffic based on a hash of source and destination addresses4. Per-prefix load balancing can be enabled globally or per interface using the load-balance per-packet command.
Reference: 4: https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/137544-technote-mpls-00.html

NEW QUESTION 7
Exhibit

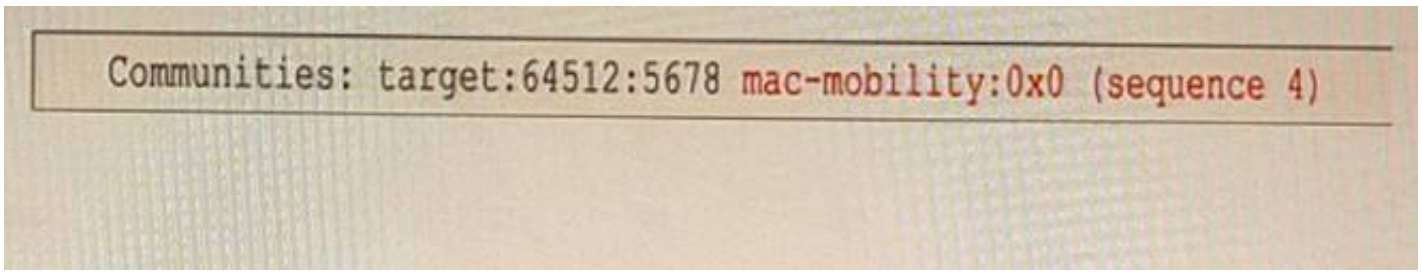

You have MAC addresses moving in your EVPN environment
Referring to the exhibit, which two statements are correct about the sequence number? (Choose two)

A. It identifies MAC addresses that should be discarded.
B. It resolves conflicting MAC address ownership claims.
C. It helps the local PE to identify the latest advertisement.
D. It is advertised using a Type 2 message

Answer: BC

Explanation:
The sequence number is a field in the MAC mobility extended community that is used to resolve conflicting MAC address ownership claims and to help the local PE to identify the latest advertisement. The sequence number is incremented by one for every MAC address mobility event, such as when a host moves from one Ethernet segment to another segment in the EVPN network. The PE device that receives multiple MAC advertisements for the same MAC address chooses the one with the highest sequence number as the most recent and valid advertisement.

NEW QUESTION 8
Exhibit

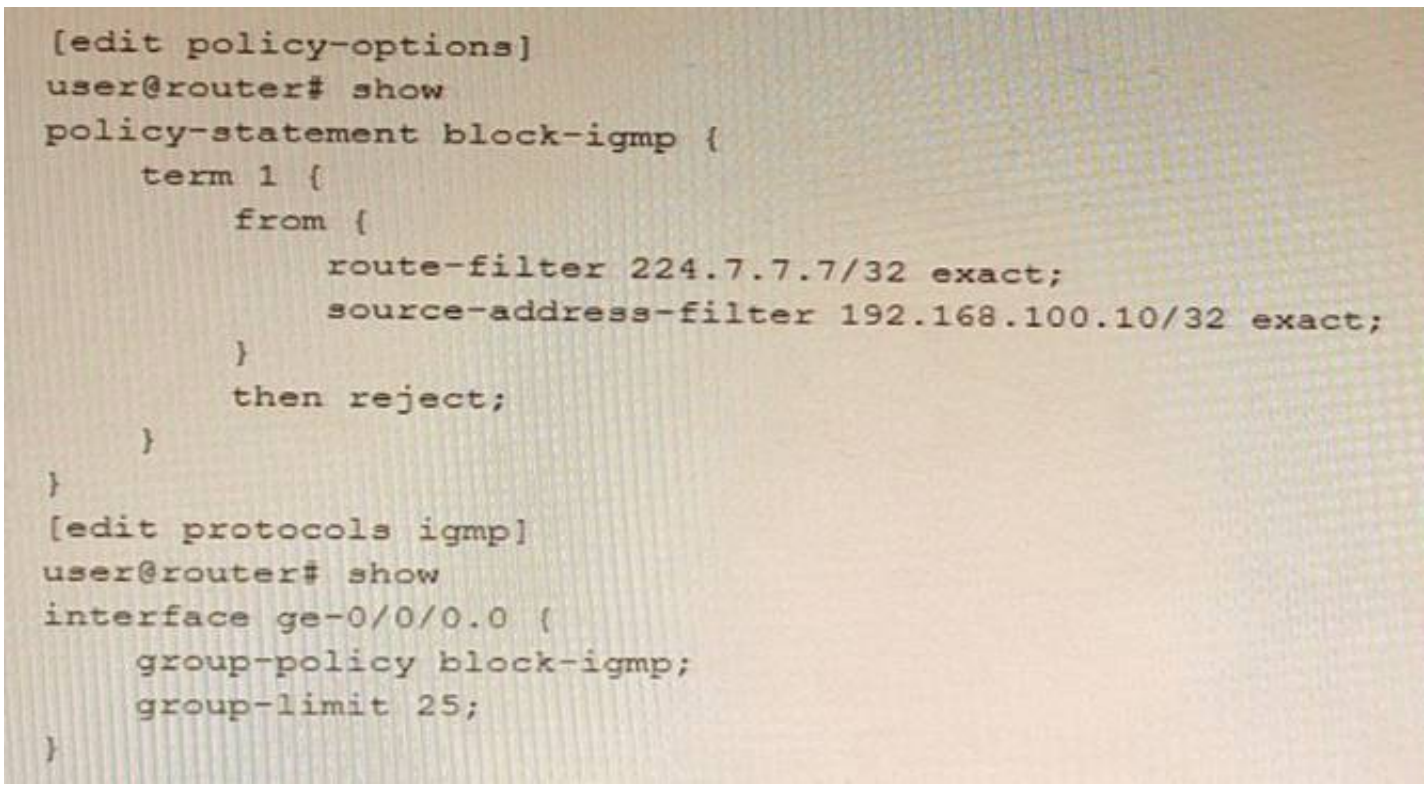

Based on the configuration contents shown in the exhibit, which statement is true?

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 range3. 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/xe-16/irg-xe-16-book/bgp-policy-framework.html

NEW QUESTION 9
Exhibit

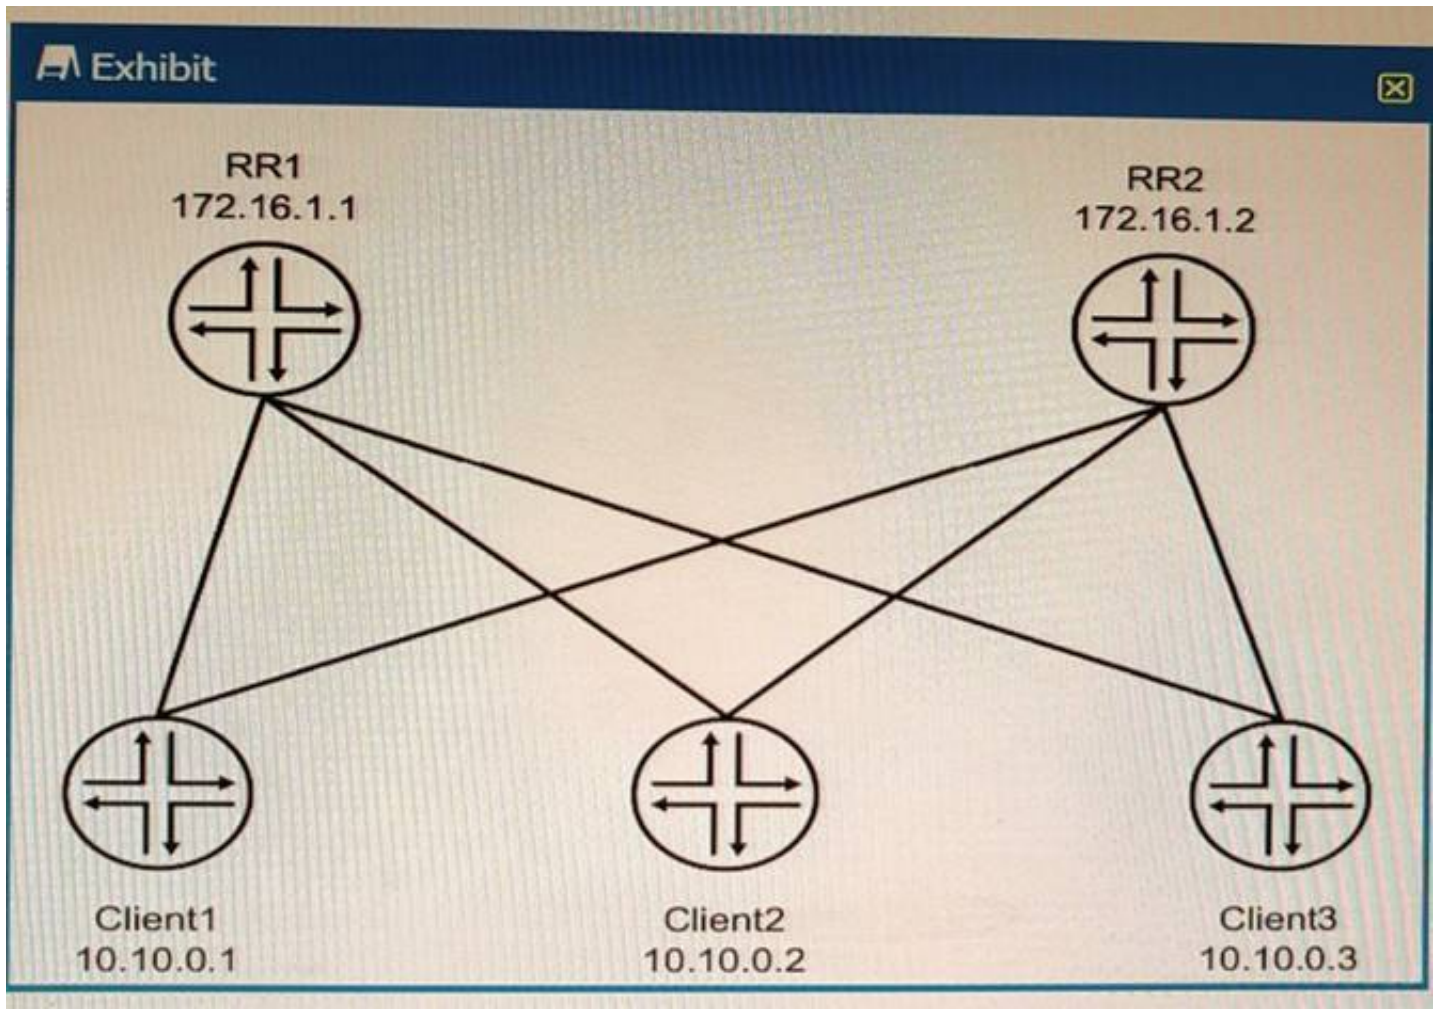

The environment is using BGP All devices are in the same AS with reachability redundancy Referring to the exhibit, which statement is correct?

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 10
In IS-IS, which two statements are correct about the designated intermediate system (DIS) on a multi-access network segment? (Choose two)

A. A router with a priority of 10 wins the DIS election over a router with a priority of 1.
B. A router with a priority of 1 wins the DIS election over a router with a priority of 10.
C. On the multi-access network, each router forms an adjacency to every other router on the segment
D. On the multi-access network, each router only forms an adjacency to the DIS.

Answer: AD

Explanation:
In IS-IS, a designated intermediate system (DIS) is a router that is elected on a multi-access network segment (such as Ethernet) to perform some functions on behalf of other routers on the same segment. A DIS is responsible for sending network link-state advertisements (LSPs), which describe all the routers attached to the network. These LSPs are flooded throughout a single area. A DIS also generates pseudonode LSPs, which represent the multi-access network as a single node in the link-state database. A DIS election is based on the priority value configured on each router's interface connected to the multi-access network. The priority value ranges from 0 to 127, with higher values indicating higher priority. The router with the highest priority becomes the DIS for the area (Level 1, Level 2, or both). If routers have the same priority, then the router with the highest MAC address is elected as the DIS. By default, routers have a priority value of 64. On a multi-access network, each router only forms an adjacency to the DIS, not to every other router on the segment. This reduces the amount of hello packets and LSP

NEW QUESTION 10
Exhibit

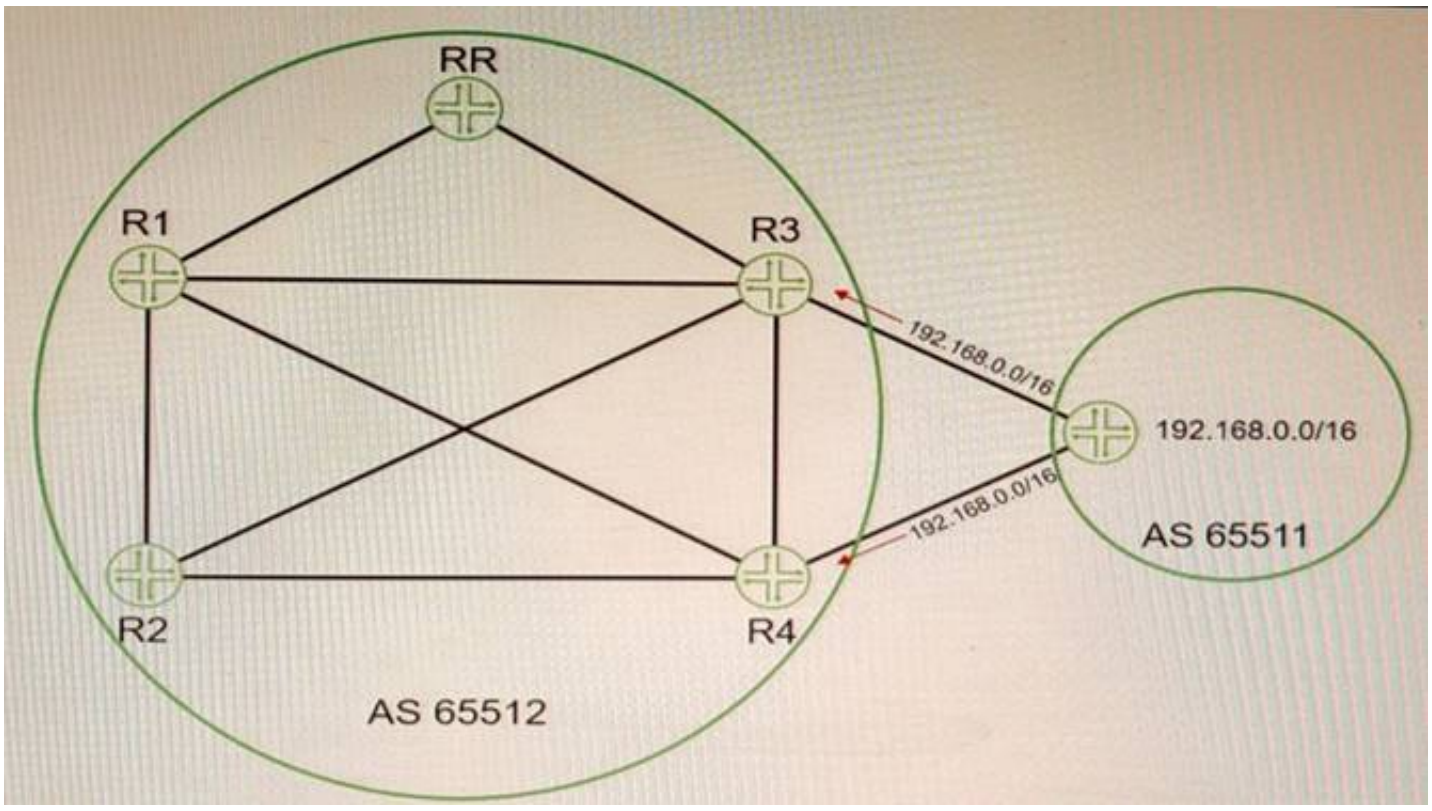

Referring to the exhibit, you are receiving the 192.168 0 0/16 route on both R3 and R4 from your EBGP neighbor You must ensure that R1 and R2 receive both BGP routes from the route reflector
In this scenario, which BGP feature should you configure to accomplish this behavior?

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 network3. 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/xe-16/irg-xe-16-book/bgp-additional-paths.html

NEW QUESTION 13
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 ,o,ns 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 18
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 19
An interface is configured with a behavior aggregate classifier and a multifield classifier How will the packet be processed when received on this interface?

A. The packet will be discarded.
B. The packet will be processed by the BA classifier first, then the MF classifier.
C. The packet will be forwarded with no classification changes.
D. The packet will be processed by the MF classifier first, then the BA classifier.

Answer: C

Explanation:
behavior aggregate (BA) classifiers and multifield (MF) classifiers are two types of classifiers that are used to assign packets to a forwarding class and a loss priority based on different criteria. The forwarding class determines the output queue for a packet. The loss priority is used by a scheduler to control packet discard during periods of congestion.
A BA classifier maps packets to a forwarding class and a loss priority based on a fixed- length field in the packet header, such as DSCP, IP precedence, MPLS EXP, or IEEE 802.1p CoS bits. A BA classifier is computationally efficient and suitable for core devices that handle high traffic volumes. A BA classifier is useful if the traffic comes from a trusted source and the CoS value in the packet header is trusted.
An MF classifier maps packets to a forwarding class and a loss priority based on multiple fields in the packet header, such as source address, destination address, protocol type, port number, or VLAN ID. An MF classifier is more flexible and granular than a BA classifier and can match packets based on complex filter rules. An MF classifier is suitable for edge devices that need to classify traffic from untrusted sources or rewrite packet headers.
You can configure both a BA classifier and an MF classifier on an interface. If you do this, the BA classification is performed first and then the MF classification. If the two classification results conflict, the MF classification result overrides the BA classification result.
Based on this information, we can infer the following statements:
? The packet will be discarded. This is not correct because the packet will not be discarded by the classifiers unless it matches a filter rule that specifies discard as an action. The classifiers only assign packets to a forwarding class and a loss priority based on their match criteria.
? The packet will be processed by the BA classifier first, then the MF classifier. This is correct because if both a BA classifier and an MF classifier are configured on an interface, the BA classification is performed first and then the MF classification. If they conflict, the MF classification result overrides the BA classification result.
? The packet will be forwarded with no classification changes. This is not correct because the packet will be classified by both the BA classifier and the MF classifier if they are configured on an interface. The final classification result will determine which output queue and which discard policy will be applied to the packet.
? The packet will be processed by the MF classifier first, then the BA classifier. This is not correct because if both a BA classifier and an MF classifier are configured on an interface, the BA classification is performed first and then the MF classification. If they conflict, the MF classification result overrides the BA classification result.

NEW QUESTION 24
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 27
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 31
Exhibit

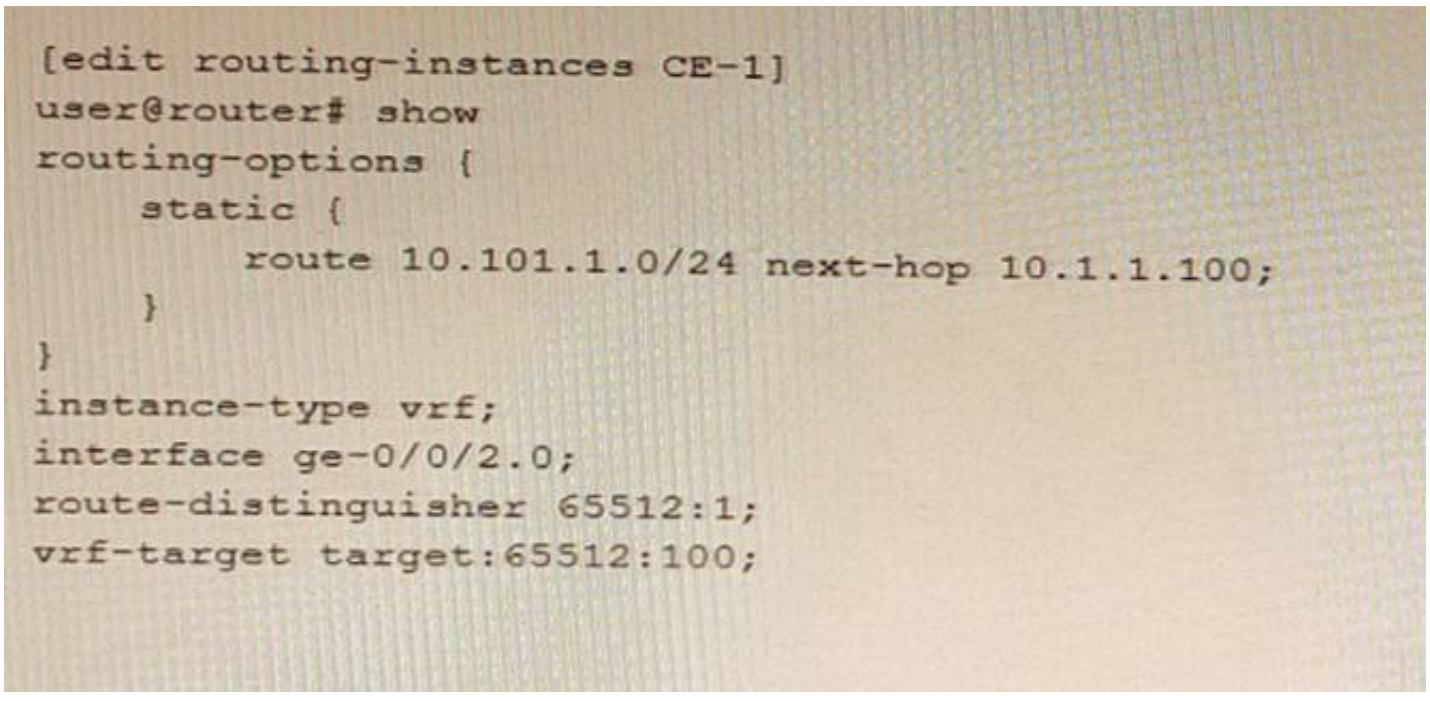

Referring to the exhibit, which statement is true?

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 36
By default, which statement is correct about OSPF summary LSAs?

A. All Type 2 and Type 7 LSAs will be summanzed 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 41
When using OSPFv3 for an IPv4 environment, which statement is correct?

A. OSPFv3 only supports IPv4.
B. OSPFv3 supports both IPv6 and IPv4, but not in the same routing instance.

C. OSPFv3 is not backward compatible with IPv4
D. OSPFv3 supports IPv4 only on interfaces with family inet6 defined

Answer: C

Explanation:
OSPFv3 is an extension of OSPFv2 that supports IPv6 routing and addressing. OSPFv3 is not backward compatible with IPv4 because it uses a different packet format and a different link-state advertisement (LSA) structure than OSPFv2. OSPFv3 also uses IPv6 link-local addresses as router IDs and neighbor addresses, instead of IPv4 addresses. To use OSPFv3 for an IPv4 environment, you need to enable the IPv4 unicast address family under [edit protocols ospf3] hierarchy level and configure IPv4 addresses on the interfaces.

NEW QUESTION 46
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 47
After a recent power outage, your manager asks you to investigate ways to automatically reduce the impact caused by suboptimal routing in your OSPF and OSPFv3 network after devices reboot.
Which three configuration statements accomplish this task? (Choose three.)

A. set protocols ospf overload timeout 900
B. set protocols ospf3 realm ipv4-unicast overload timeout 900
C. set protocols ospf overload
D. set protocols oapf3 overload timeout 900
E. set protocols ospf3 overload

Answer: AE

Explanation:
To reduce the impact of suboptimal routing in OSPF and OSPFv3 after devices reboot, you can use the overload feature to prevent a router from being used as a transit router for a specified period of time. This allows the router to stabilize its routing table before forwarding traffic for other routers. To enable the overload feature, you need to do the following:
? For OSPF, configure the overload statement under [edit protocols ospf] hierarchy level. You can also specify a timeout value in seconds to indicate how long the router should remain in overload state after it boots up. For example, set protocols ospf overload timeout 900 means that the router will be in overload state for 15 minutes after it boots up.
? For OSPFv3, configure the overload statement under [edit protocols ospf3] hierarchy level. You can also specify a realm (ipv4-unicast or ipv6-unicast) and a timeout value in seconds to indicate how long the router should remain in overload state after it boots up for each realm. For example, set protocols ospf3 realm ipv4- unicast overload timeout 900 means that the router will be in overload state for 15 minutes after it boots up for IPv4 unicast routing.

NEW QUESTION 51
Exhibit

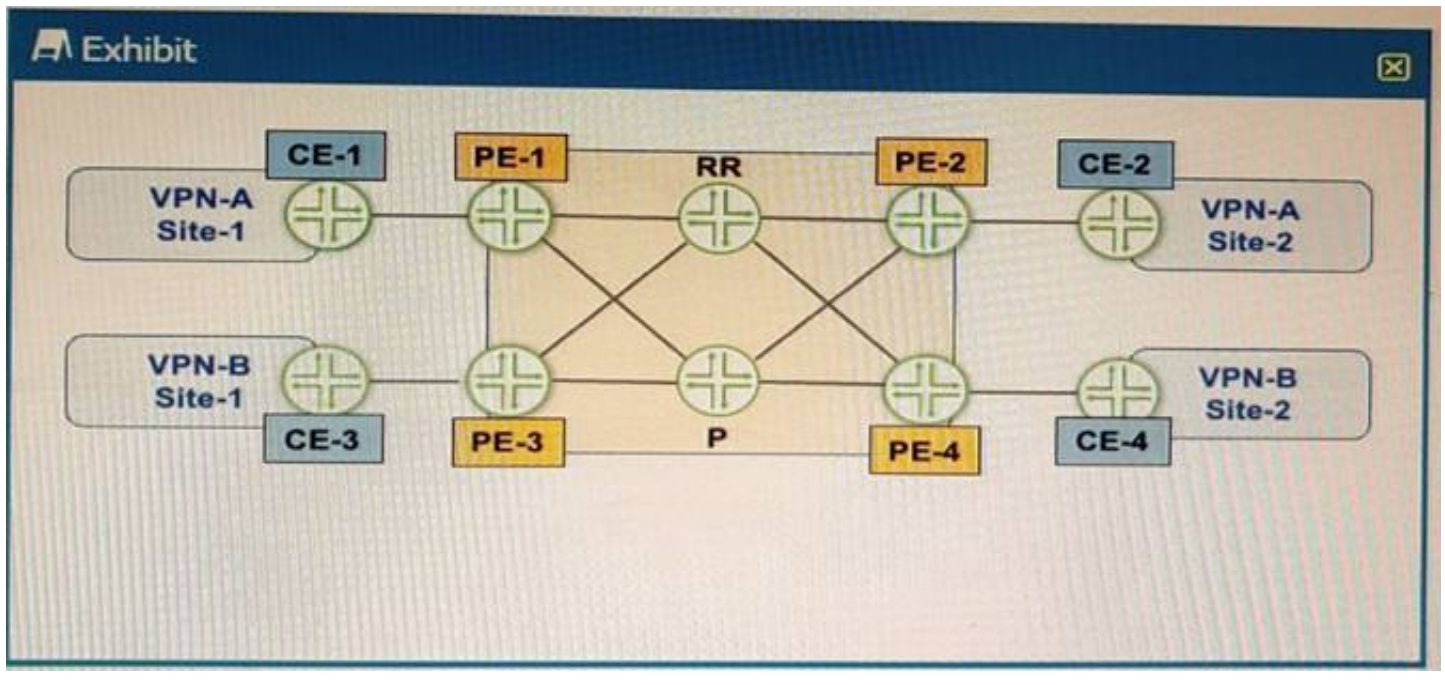

Referring to the exhibit, PE-1 and PE-2 are getting route updates for VPN-B when neither of them service that VPN
Which two actions would optimize this process? (Choose two.)

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 . 13vpn . 0 resolution-ribs ine

D. 0 Statement on the PEs.
E. Configure the resolution rib bgp.l3vpn.O 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 56
Exhibit

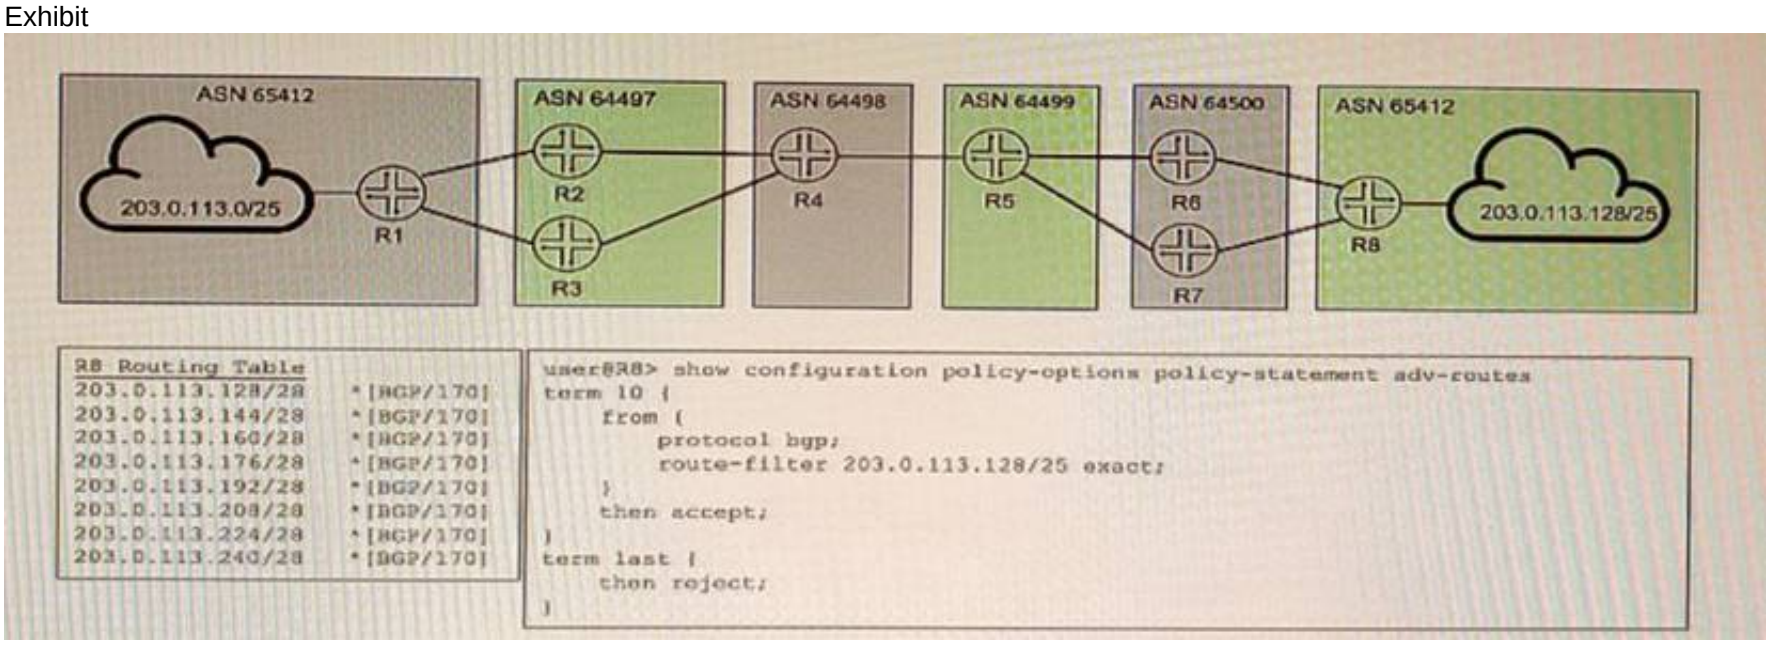

You are attempting to summarize routes from the 203.0.113.128/25 IP block on R8 to AS 64500. You implement the export policy shown in the exhibit and all routes from the routing table stop being advertised.
In this scenario, which two steps would you take to summarize the route in BGP? (Choose two.)

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 57
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 61
Exhibit

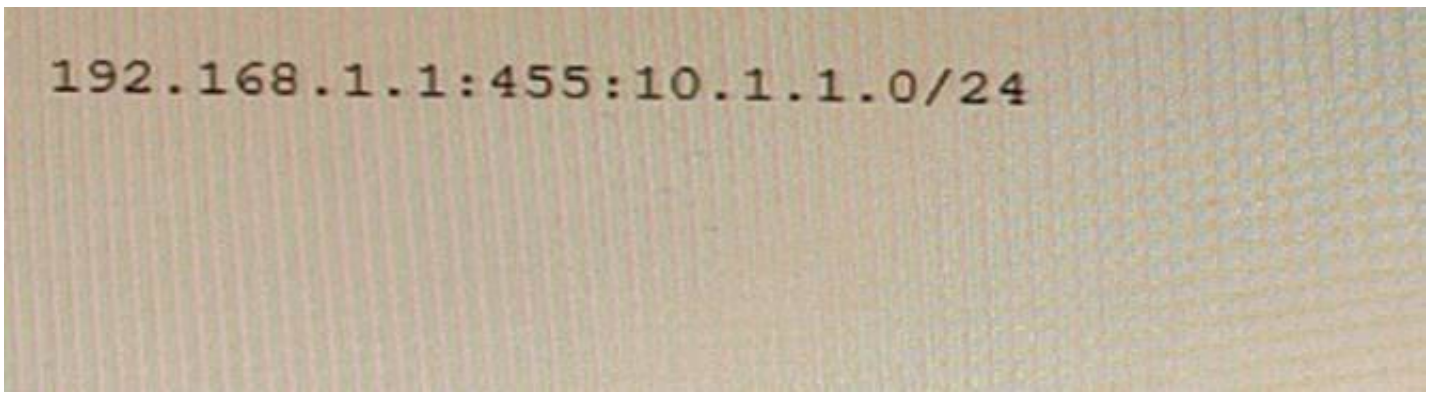

You are examining an L3VPN route that includes the information shown in the exhibit Which statement is correct in this scenario?

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 66
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 68
Which two statements are correct about IS-IS interfaces? (Choose two.)

A. If a broadcast interface is in both L1 and L2, one combined hello message is sent for both levels.
B. If a point-to-point interface is in both L1 and L2, separate hello messages are sent for each level.
C. If a point-to-point interface is in both L1 and L2, one combined hello message is sent for both levels.
D. If a broadcast interface is in both L1 and L2, separate hello messages are sent for each level

Answer: BD

Explanation:
IS-IS supports two levels of routing: Level 1 (intra-area) and Level 2 (interarea). An IS-IS router can be either Level 1 only, Level 2 only, or both Level 1 and Level 2. A router that is both Level 1 and Level 2 is called a Level 1-2 router. A Level 1-2 router sends separate hello messages for each level on both point-to-point and broadcast interfaces1. A point-to-point interface provides a connection between a single source and a single destination. A broadcast interface behaves as if the router is connected to a LAN.

NEW QUESTION 73
Exhibit

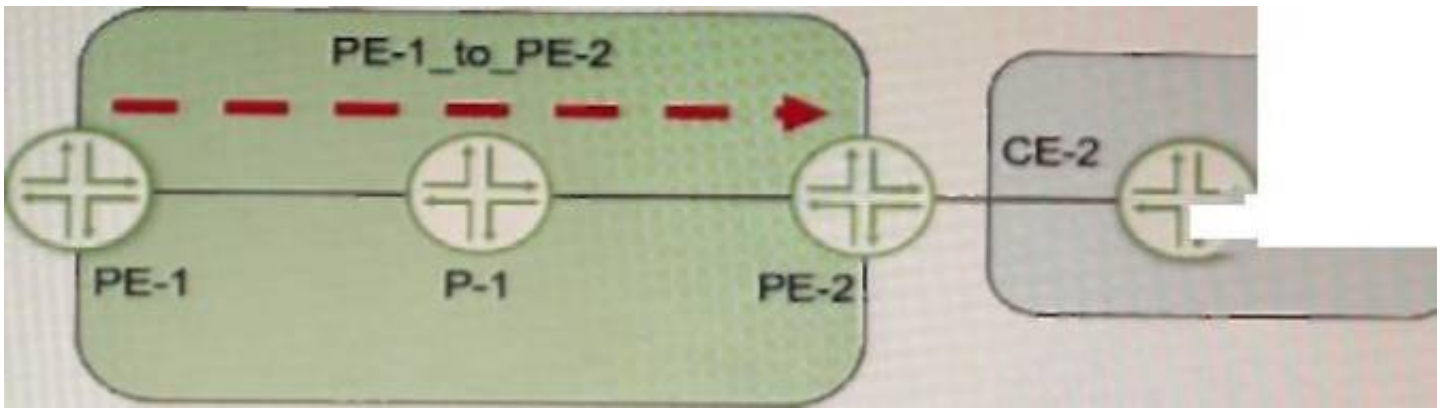

Referring to the exhibit, a working L3VPN exists that connects VPN-A sites CoS is configured correctly to match on the MPLS EXP bits of the LSP, but when traffic is sent from Site-1 to Site-2, PE-2 is not classifying the traffic correctly
What should you do to solve the problem?

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 74
Exhibit

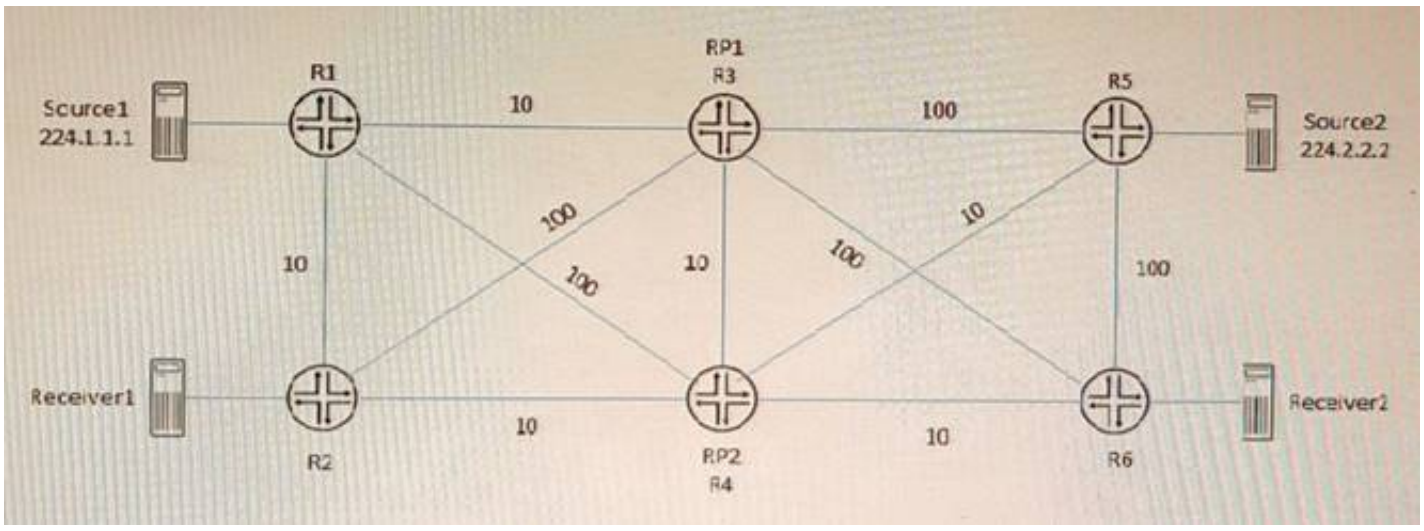

Referring to the exhibit, PIM-SM is configured on all routers, and Anycast-RP with Anycast- PIM is used for the discovery mechanism on RP1 and RP2. The interface metric values are shown for the OSPF area.
In this scenario, which two statements are correct about which RP is used? (Choose two.)

A. Source2 will use RP2 and Received 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 neighborship 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 77
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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 FirstTry

* JN0-664 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year