Number: JN0-664
Passing Score: 800
Time Limit: 120
File Version: 3.0

Exam Code: JN0-664

Exam Name: Service Provider Routing and Switching, Professional

Exam A

QUESTION 1
Exhibit

```
[edit routing-instances CE-1]
user@R1# show
protocols {
    bgp {
        group CE-1 {
            type external;
            peer-as 65555;
            neighbor 10.1.1.100;
        }
    }
}
instance-type vrf;
interface ge-0/0/2.0;
route-distinguisher 65512:1;
vrf-target target:65512:100;
[edit routing-instances CE-2]
user@R2# show
protocols {
    bgp {
        group CE-2 {
            type external;
            peer-as 64444;
            neighbor 10.1.5.100;
        }
    }
}
instance-type vrf;
interface ge-0/0/3.0;
route-distinguisher 65512:1;
vrf-target target:65512:100;
```

Referring to the exhibit, which statement is correct?

A. The vrf-target configuration will allow routes to be shared between CE-1 and CE-2.
B. The vrf-target configuration will stop routes from being shared between CE-1 and CE-2.
C. The route-distinguisher configuration will allow overlapping routes to be shared between CE-1 and CE-2.

D. The route-diatinguisher configuration will stop routes from being shared between CE-1 and CE-2.

Correct Answer: C
Section:
Explanation:
The route distinguisher (RD) is a BGP attribute that is used to create unique VPN IPv4 prefixes for each VPN in an MPLS network. The RD is a 64-bit value that consists of two parts: an administrator field and an assigned number field. The administrator field can be an AS number or an IP address, and the assigned number field can be any arbitrary value chosen by the administrator. The RD is prepended to the IPv4 prefix to create a VPN IPv4 prefix that can be advertised across the MPLS network without causing any overlap or conflict with other VPNs. In this question, we have two PE routers (PE-1 and PE-2) that are connected to two CE devices (CE-1 and CE-2) respectively. PE-1 and PE-2 are configured with VRFs named Customer-A and Customer-B respectively.

QUESTION 2
Exhibit

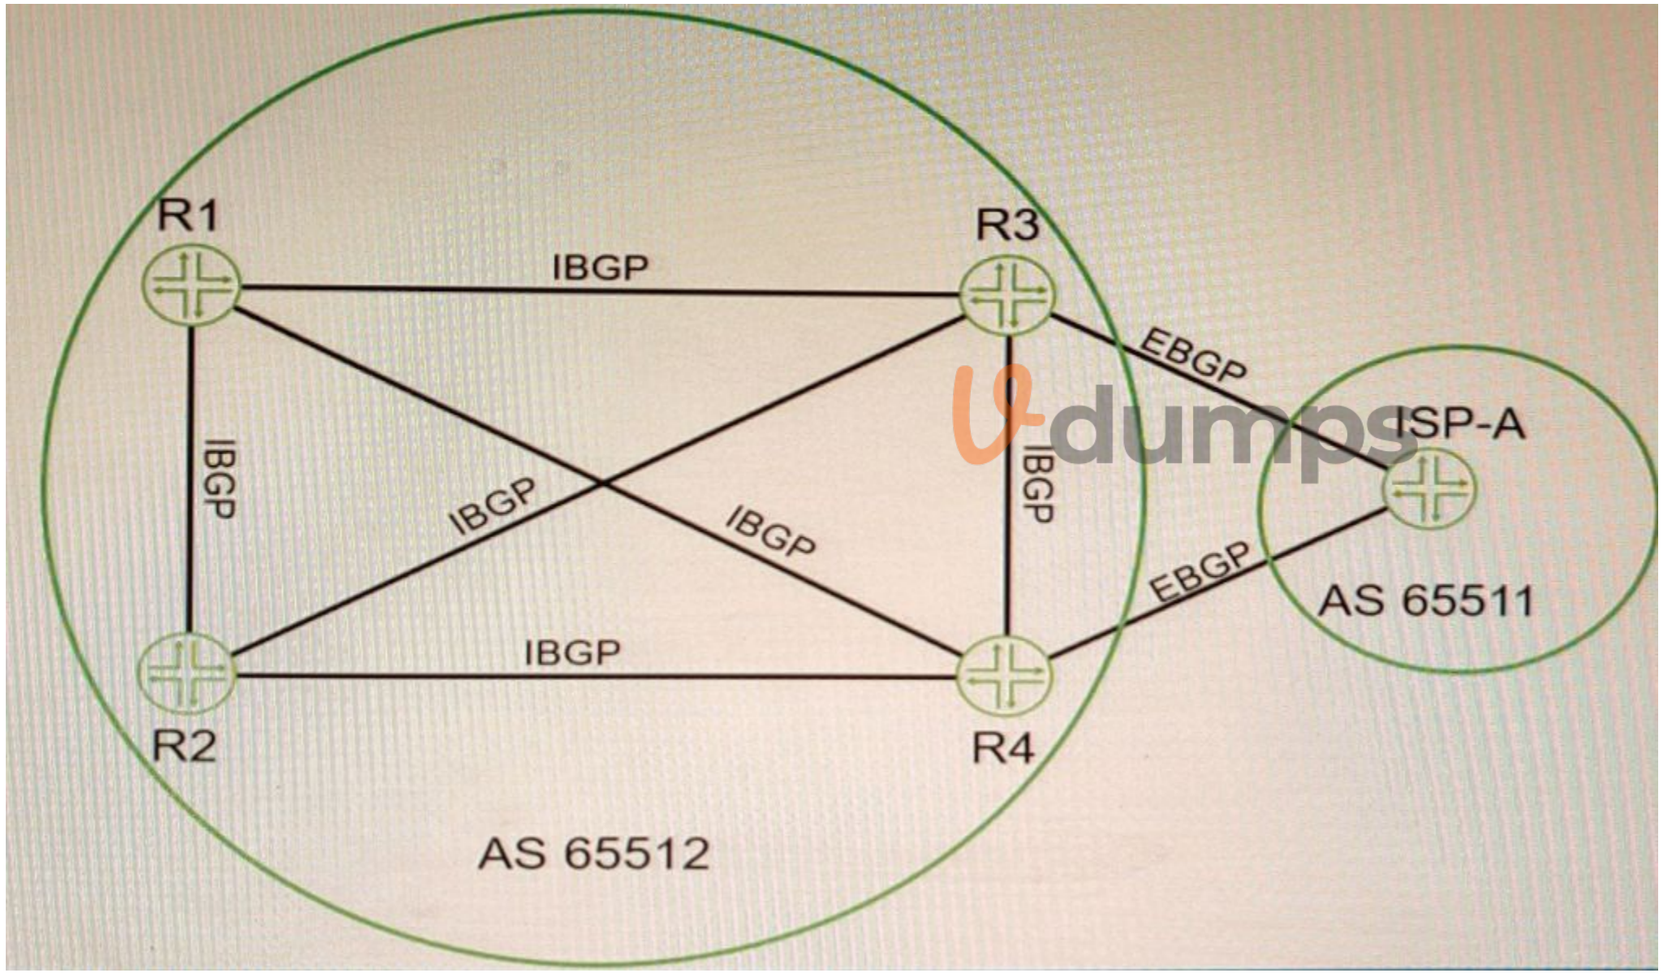

Click the Exhibit button-Referring to the exhibit, which two statements are correct about BGP routes on R3 that are learned from the ISP-A neighbor? (Choose two.)

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

Correct Answer: A, B
Section:
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.

QUESTION 3
Exhibit

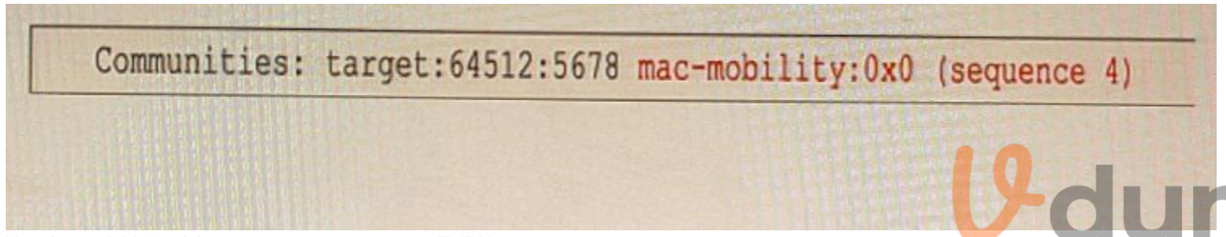

You have MAC addresses moving in your EVPN environment
Referring to the exhibit, which two statements are correct about the sequence number? (Choose two)

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

Correct Answer: B, C
Section:
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.

QUESTION 4
Exhibit.

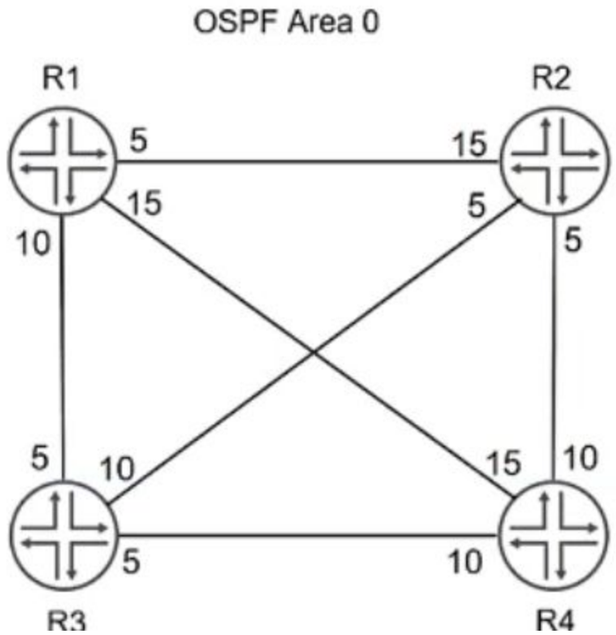

Referring to the exhibit, which path would traffic passing through R1 take to get to R4?

A. R1 -> R3 -> R4
B. R1 -> R2 -> R3 -> R4
C. R1 -> R2 -> R4
D. R1 -> R4

Correct Answer: C
Section:
Explanation:
The OSPF cost is carried in the LSAs that are exchanged within an OSPF area. When a router calculates the cost to a destination it uses the cost of the exit interface of each router in the path to the destination.

QUESTION 5
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.

Correct Answer: C, D
Section:
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.

QUESTION 6
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 ospf3 realm ipv4-unicast overload timeout 900
B. set protocols ospf overload
C. set protocols ospf overload timeout 900
D. set protocols ospf3 overload
E. set protocols ospf3 overload timeout 900

Correct Answer: A, C, E
Section:
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.

QUESTION 7
Exhibit

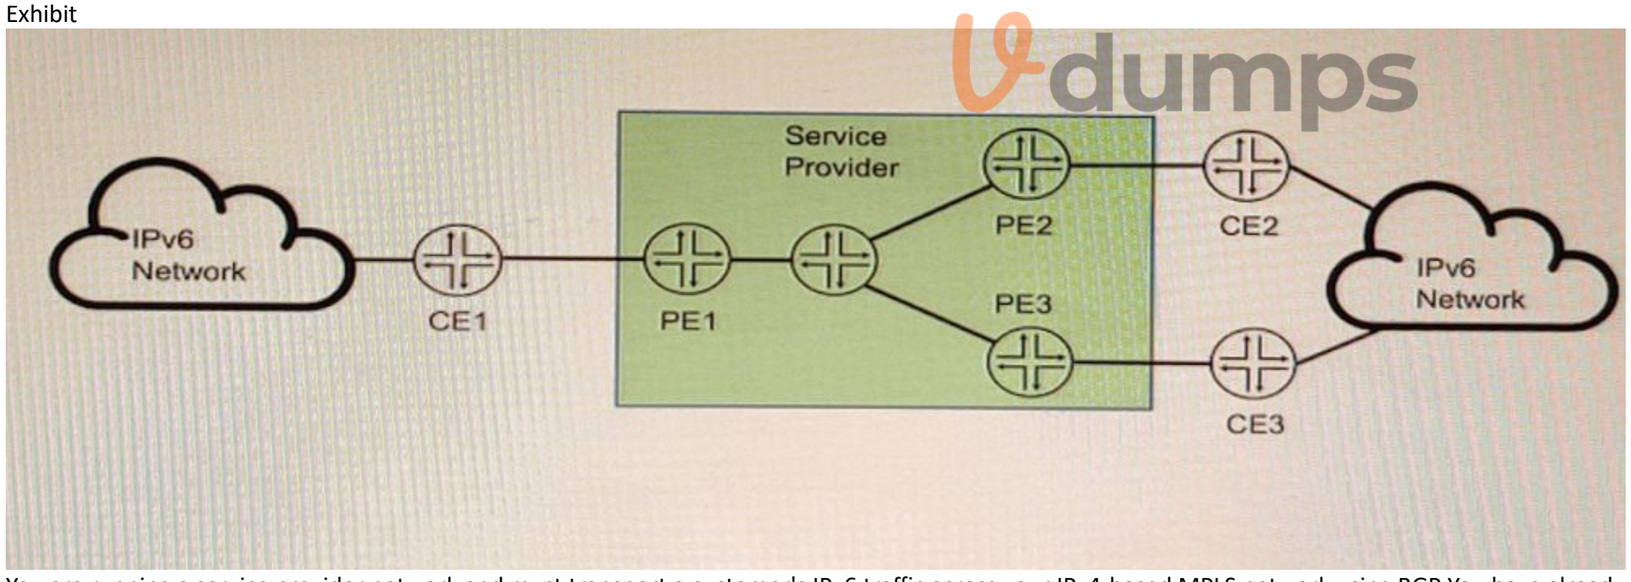

You are running a service provider network and must transport a customer's IPv6 traffic across your IPv4-based MPLS network using BGP You have already configured mpis ipv6-tunneling on your PE routers.
Which two statements are correct about the BGP configuration in this scenario? (Choose two.)

A. You must configure family inet6 labcled-unicast between PE routers.
B. You must configure family inet6 unicaat between PE and CE routers.
C. You must configure family inet6 add-path between PE and CE routers.
D. You must configure family inet6 unicast between PE routers

Correct Answer: A, B
Section:
Explanation:
To transport IPv6 traffic over an IPv4-based MPLS network using BGP, you need to configure two address families: family inet6 labeled-unicast and family inet6 unicast. The former is used to exchange IPv6 routes with MPLS labels between PE routers, and the latter is used to exchange IPv6 routes without labels between PE and CE routers. The mpis ipv6-tunneling command enables the PE routers to encapsulate the IPv6 packets with an MPLS label stack and an IPv4 header before sending them over the MPLS network.

QUESTION 8
You want to ensure that L1 IS-IS routers have only the most specific routes available from L2 IS-IS routers. Which action accomplishes this task?

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

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

QUESTION 9
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

Correct Answer: C
Section:
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.

QUESTION 10
Exhibit

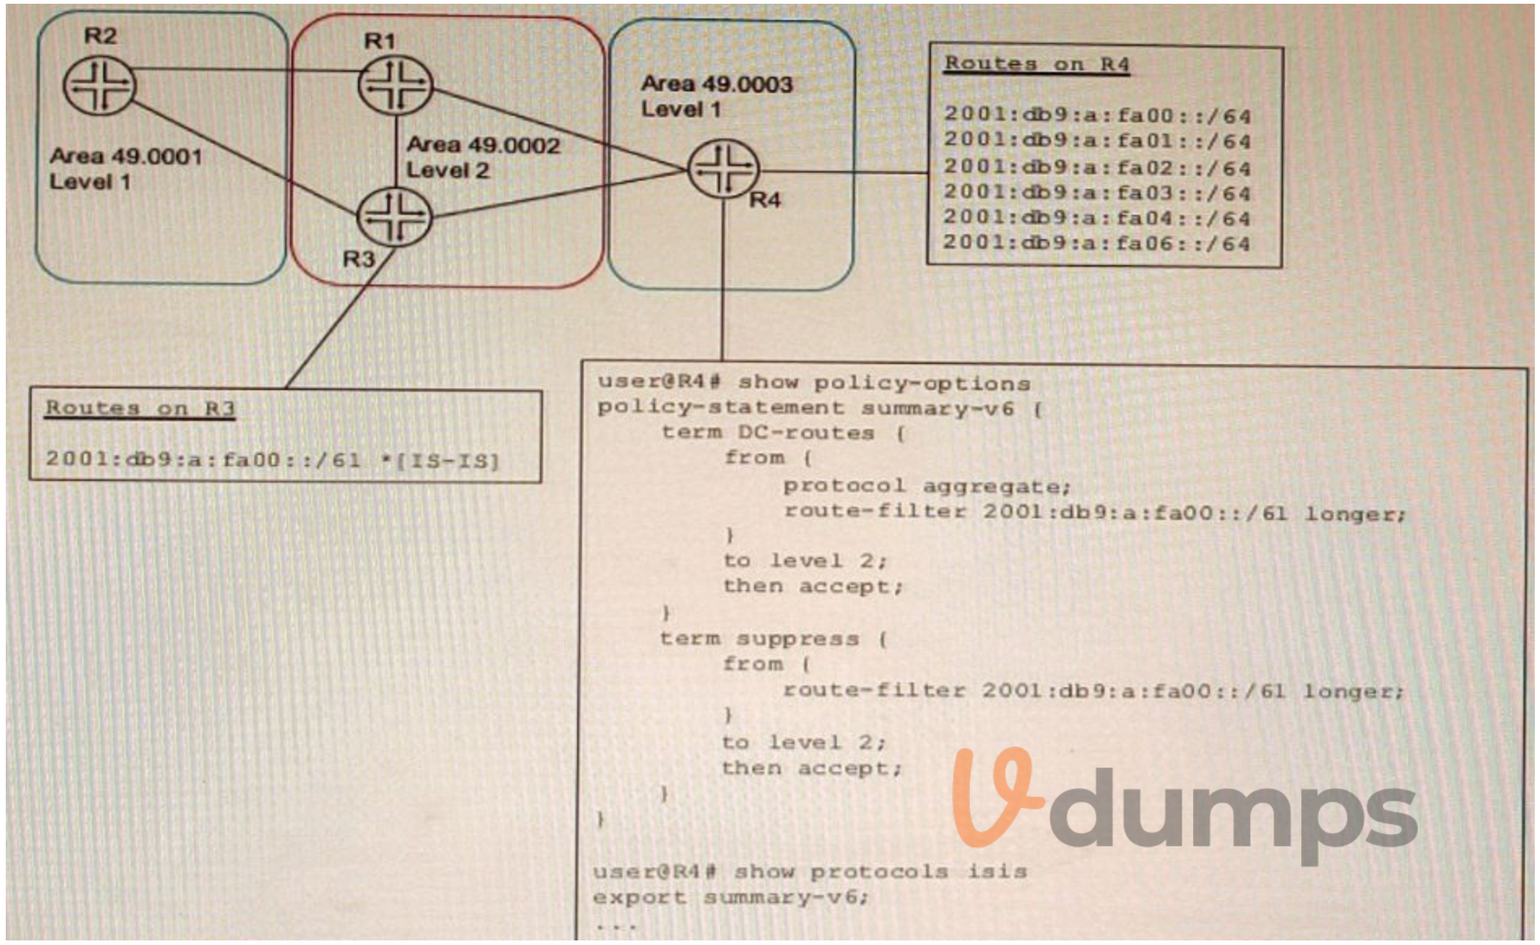

A network designer would like to create a summary route as shown in the exhibit, but the configuration is not working.
Which three configuration changes will create a summary route? (Choose three.)

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 : :/6l 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

Correct Answer: B, C, D
Section:
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.

QUESTION 11
Exhibit

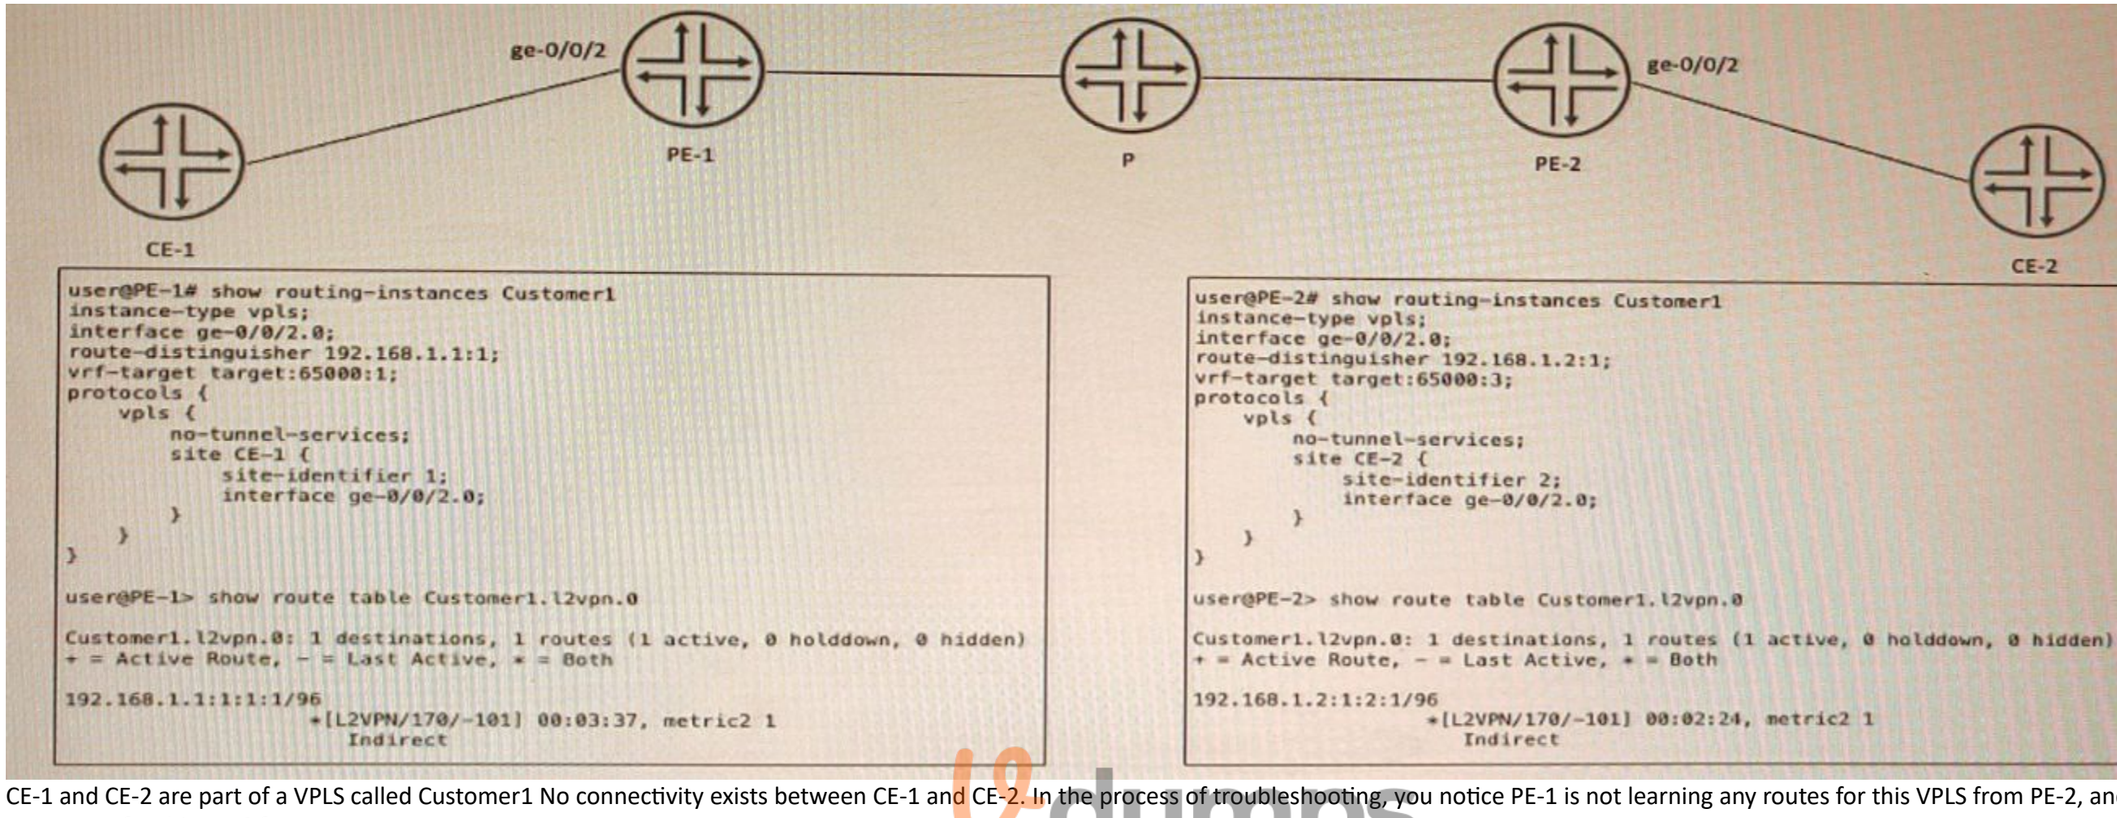

CE-1 and CE-2 are part of a VPLS called Customer1 No connectivity exists between CE-1 and CE-2. In the process of troubleshooting, you notice PE-1 is not learning any routes for this VPLS from PE-2, and PE-2 is not learning any routes for this VPLS from PE-1.

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.

Correct Answer: A
Section:
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.

QUESTION 12
Exhibit

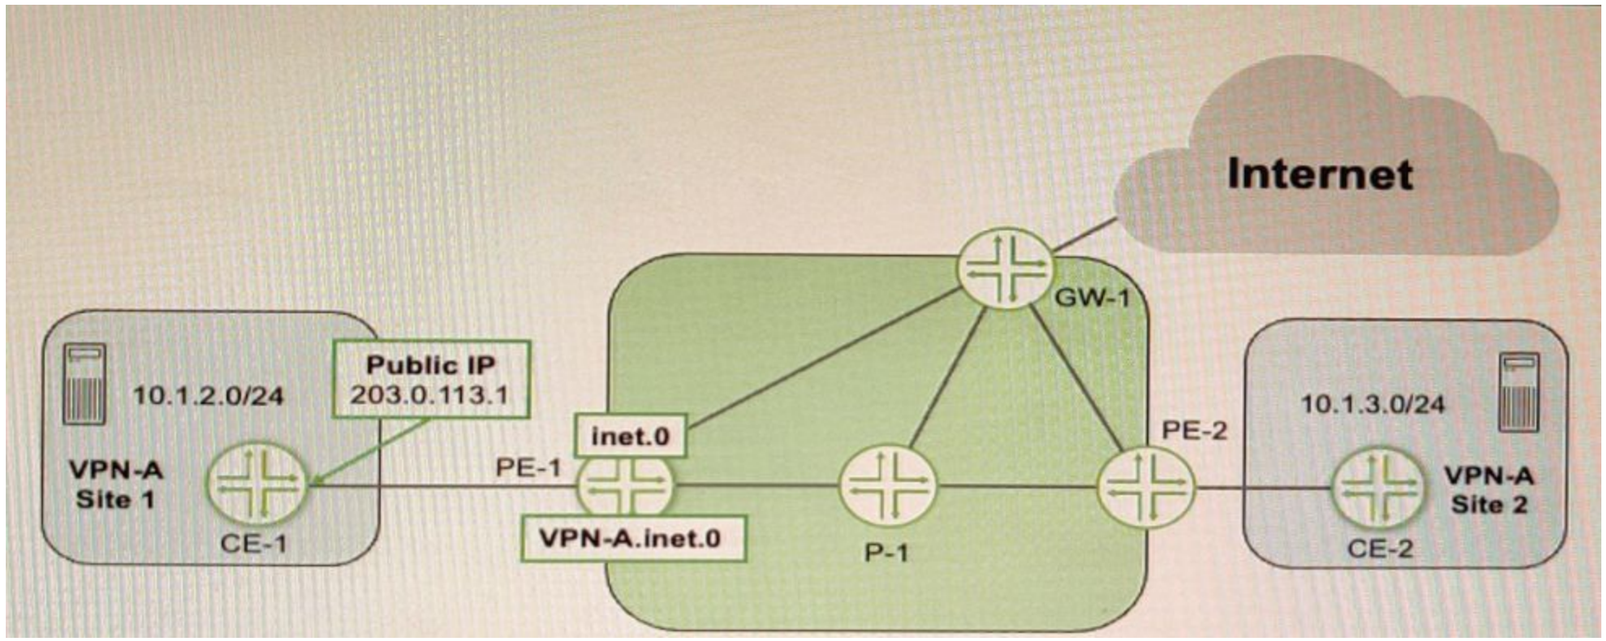

Referring to the exhibit, CE-1 is providing NAT services for the hosts at Site 1 and you must provide Internet access for those hosts
Which two statements are correct in this scenario? (Choose two.)

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.

Correct Answer: B, C
Section:

QUESTION 13
Exhibit

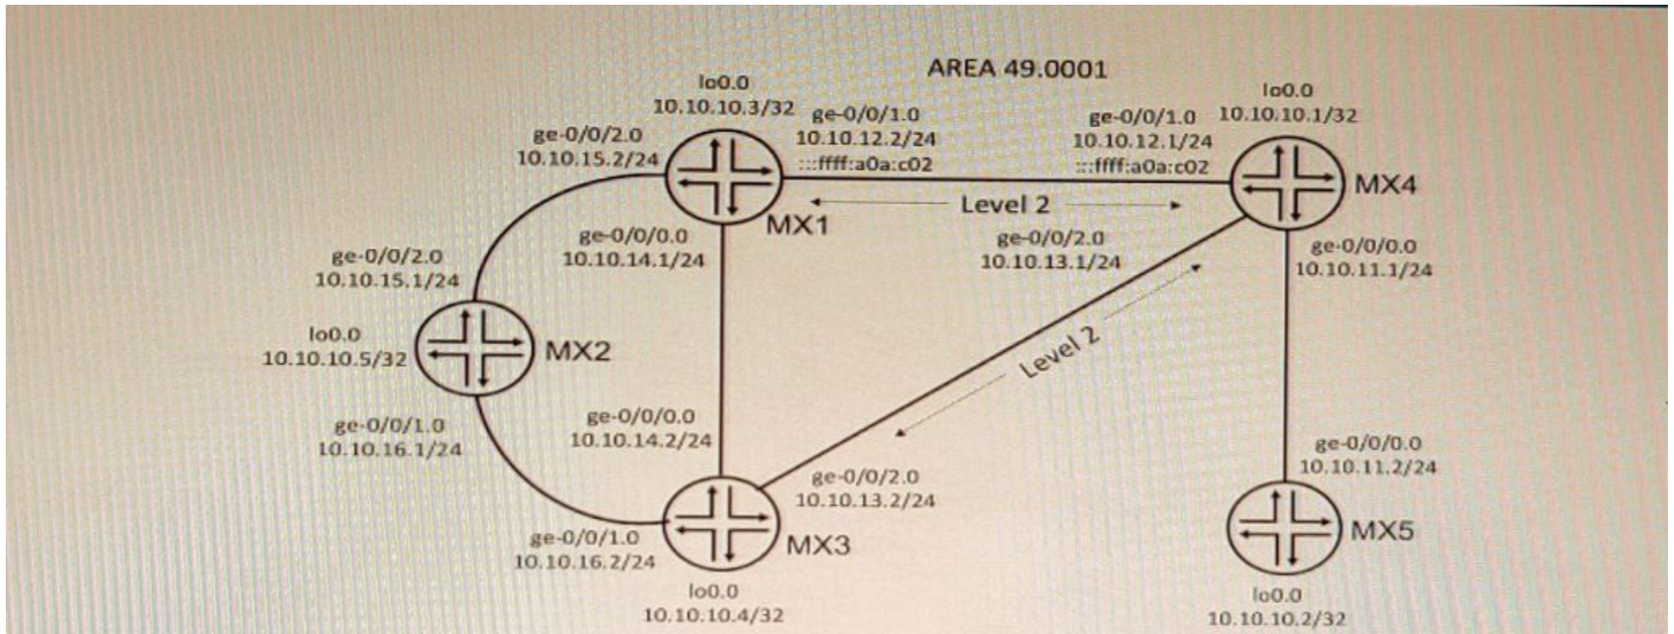


```
user@MX1> show isis interface ge-0/0/1
IS-IS interface database:
Interface L         CirID        Level 1 DR        Level 2 DR                L1/L2 Metric
ge-0/0/0/0          2            0x1               Disabled Point to Point   10/100

user@MX1> show isis database MX1 extensive | find TLV | match 10.100.12.0/24
  IP prefix: 10.100.12.0/24, Internal, Metric: defaut 63, Up
  IP extended prefix: 10.100.12.0/24 metric 63 up
```

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

Correct Answer: A
Section:
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 metrics12.

QUESTION 14
Exhibit

You are asked to exchange routes between R1 and R4 as shown in the exhibit. These two routers use the same AS number Which two steps will accomplish this task? (Choose two.)

A. Configure the BGP group with the advertise-peer-as parameter on R1 and R4.
B. Configure the BGP group with the as-override parameter on R2 and R3
C. Configure the BGP group with the advertise-peer-as parameter on R2 and R3.
D. Configure the BGP group with the as-override parameter on R1 and R4

Correct Answer: B, C
Section:

QUESTION 15
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

Correct Answer: C
Section:
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 .

QUESTION 16
Which statement is true regarding BGP FlowSpec?

A. It uses a remote triggered black hole to protect a network from a denial-of-service attack.
B. It uses dynamically created routing policies to protect a network from denial-of-service attacks
C. It is used to protect a network from denial-of-service attacks dynamically
D. It verifies that the source IP of the incoming packet has a resolvable route in the routing table

Correct Answer: B
Section:
Explanation:
BGP FlowSpec is a feature that extends the Border Gateway Protocol (BGP) to enable routers to exchange traffic flow specifications, allowing for more precise control of network traffic. The BGP FlowSpec feature enables routers to advertise and receive information about specific flows in the network, such as those originating from a particular source or destined for a particular destination.Routers can then use this information to construct traffic filters that allow or deny packets of a certain type, rate limit flows, or perform other actions1.BGP FlowSpec can also help in filtering traffic and taking action against distributed denial of service (DDoS) attacks by dropping the DDoS traffic or diverting it to an analyzer2.BGP FlowSpec rules are internally converted to equivalent Cisco Common Classification Policy Language (C3PL) representing corresponding match and action parameters2. Therefore, BGP FlowSpec uses dynamically created routing policies to protect a network from denial-of-service attacks.

QUESTION 17
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

Correct Answer: B, D, E
Section:
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.

QUESTION 18
Which origin code is preferred by BGP?

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

Correct Answer: C
Section:
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.

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.

Correct Answer: C
Section:
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.

QUESTION 20
Exhibit

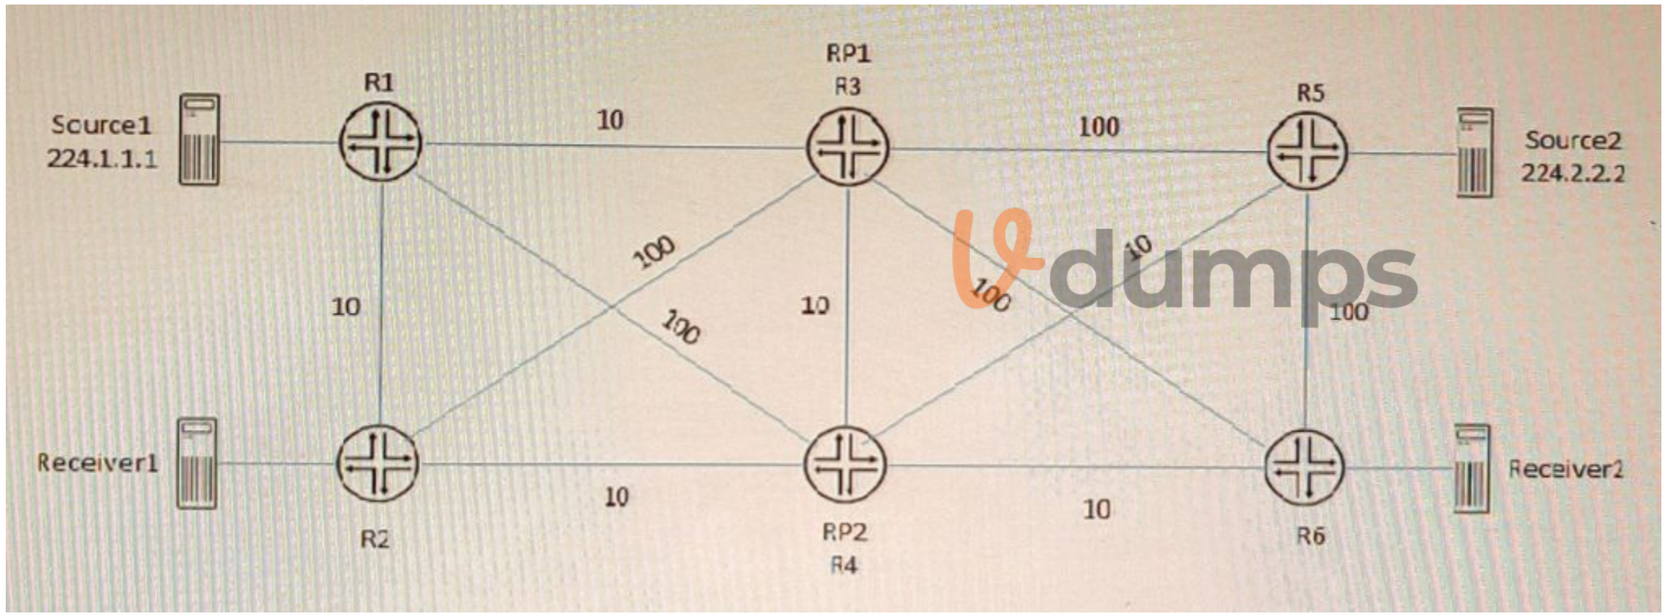

Referring to the exhibit, PIM-SM is configured on all routers, and Anycast-RP with Anycast-PIM is used for the discovery mechanism on RP1 and RP2. The interface metric values are shown for the OSPF area.
In this scenario, which two statements are correct about which RP is used? (Choose two.)

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

Correct Answer: A, C
Section:
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.

QUESTION 21
Exhibit

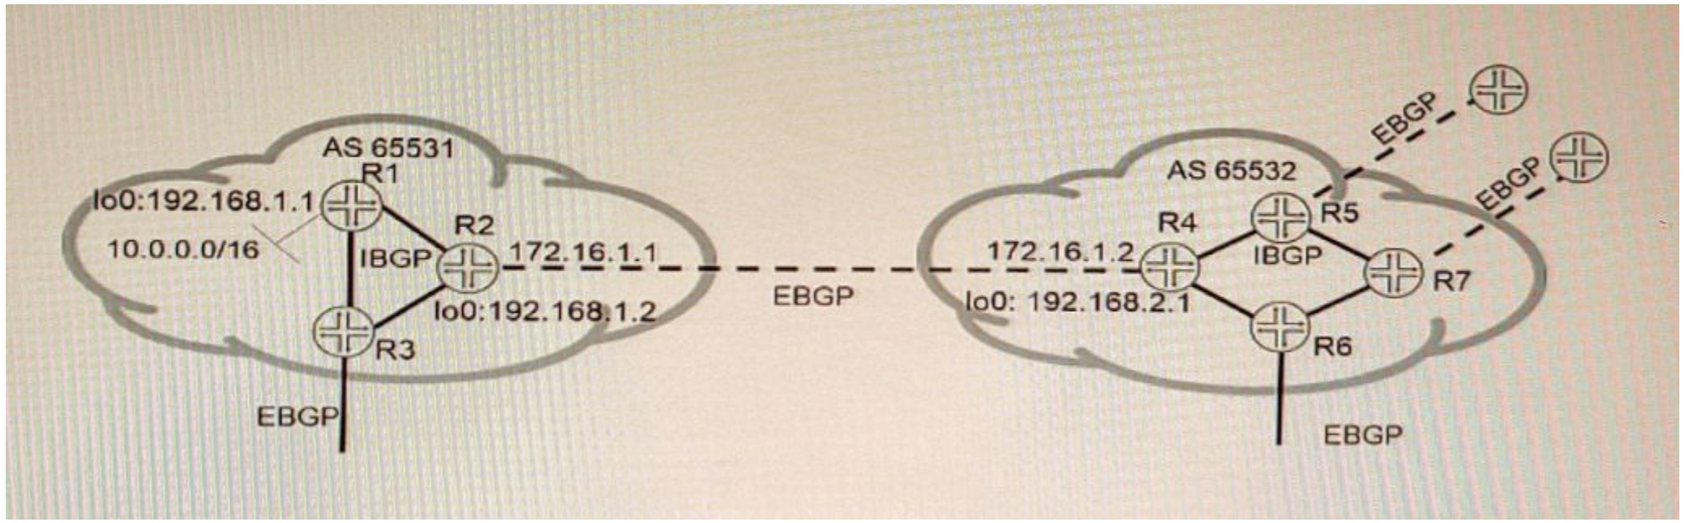

Referring to the exhibit, which three statements are correct about route 10 0 0.0/16 when using the default BGP advertisement rules'? (Choose three.)

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

Correct Answer: A, C, E
Section:

QUESTION 22
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 queueses
D. The exceeding queue will be considered to have positive bandwidth credit

Correct Answer: A, B
Section:
Explanation:
https://www.juniper.net/documentation/us/en/software/junos/cos-security-devices/topics/concept/cos-transmission-scheduling-security-overview.html

QUESTION 23
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.

Correct Answer: A, D
Section:
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

QUESTION 24
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;
```

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.

Correct Answer: D
Section:

QUESTION 25
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.

Correct Answer: C, D
Section:
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.

QUESTION 26
Exhibit

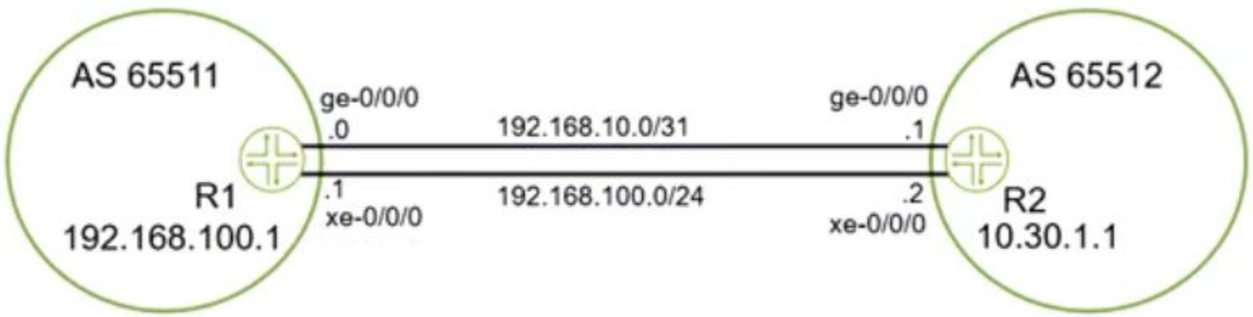

You want to use both links between R1 and R2 Because of the bandwidth difference between the two links, you must ensure that the links are used as much as possible.
Which action will accomplish this goal?

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.

Correct Answer: A
Section:
Explanation:
https://www.juniper.net/documentation/us/en/software/junos/sampling-forwarding-monitoring/bgp/topics/concept/bgp-multipath-unequal-understanding.html

QUESTION 27
Exhibit

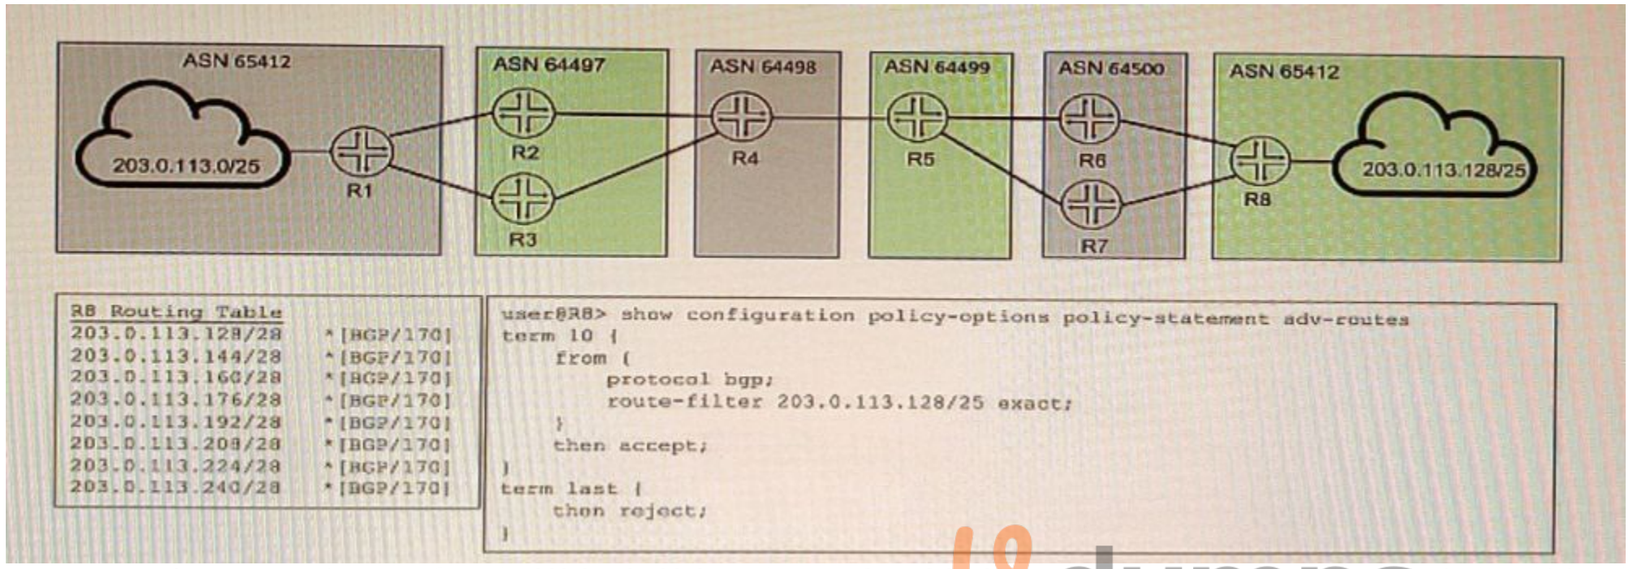

You are attempting to summarize routes from the 203.0.113.128/25 IP block on R8 to AS 64500. You implement the export policy shown in the exhibit and all routes from the routing table stop being advertised.
In this scenario, which two steps would you take to summarize the route in BGP? (Choose two.)

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 tables. -
C. Add the set routing-options static route 203.0.113.123/25 discard command.
D. Replace exact in the export policy with orlonger.

Correct Answer: C, D
Section:
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.

QUESTION 28
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.

Correct Answer: A, C
Section:
Explanation:
https://www.juniper.net/documentation/us/en/software/nce/feature-guide-virtual-private-lan-service/topics/task/vpls-ldp-signaling-solutions.html
https://www.juniper.net/documentation/us/en/software/junos/vpn-l2/topics/concept/vpns-configuring-vpls-routing-instances.html#id-11510150__id-11568648

QUESTION 29
Exhibit

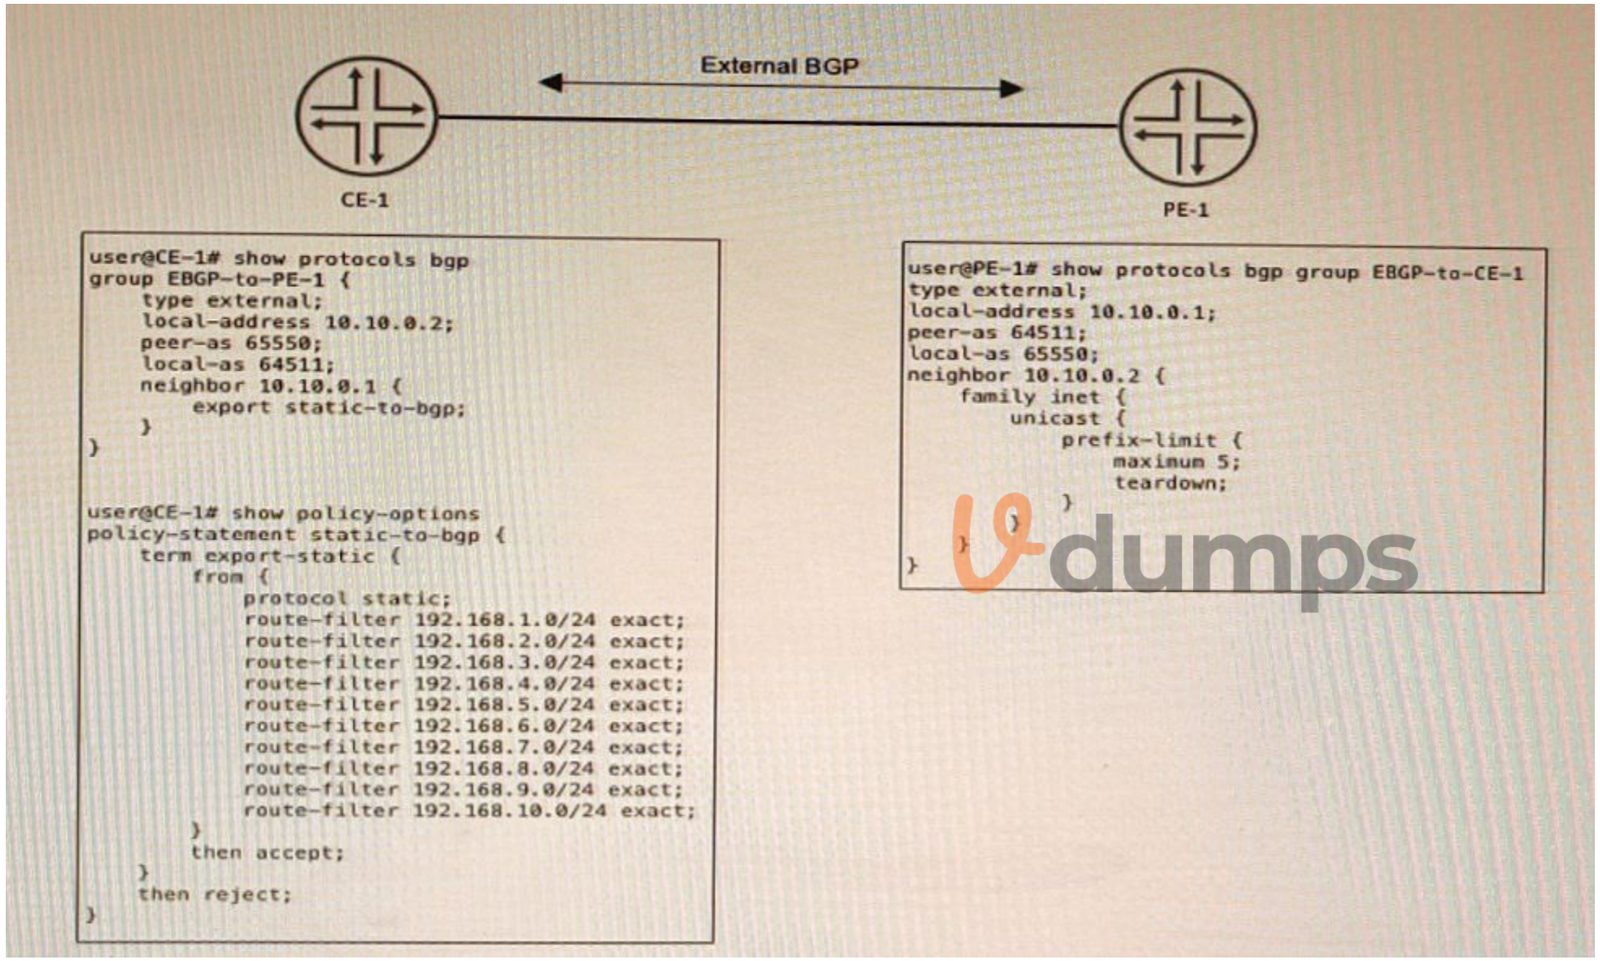

CE-1 must advertise ten subnets to PE-1 using BGP Once CE-1 starts advertising the subnets to PE-1, the BGP peering state changes to Active.
Referring to the CLI output shown in the exhibit, which statement is correct?

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

Correct Answer: B
Section:
Explanation:

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

QUESTION 30
Exhibit

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

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

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

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

QUESTION 31
Exhibit

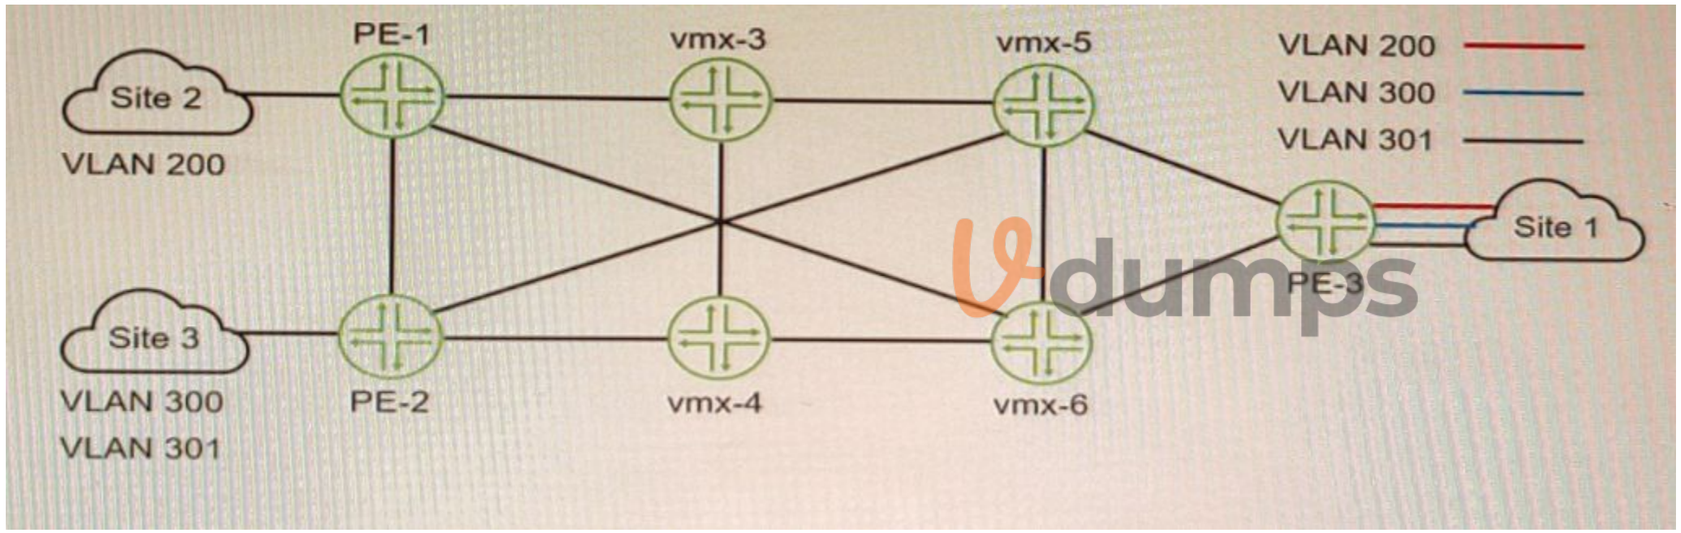

You want Site 1 to access three VLANs that are located in Site 2 and Site 3 The customer-facing interface on the PE-1 router is configured for Ethernet-VLAN encapsulation.
What is the minimum number of L2VPN routing instances to be configured to accomplish this task?

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

Correct Answer: B
Section:
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.

QUESTION 32

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

Correct Answer: C
Section:
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.

QUESTION 33
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

Correct Answer: C
Section:
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.

QUESTION 34
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-distinguiaher value on both PE devices.

Correct Answer: C
Section:
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.

QUESTION 35
Exhibit

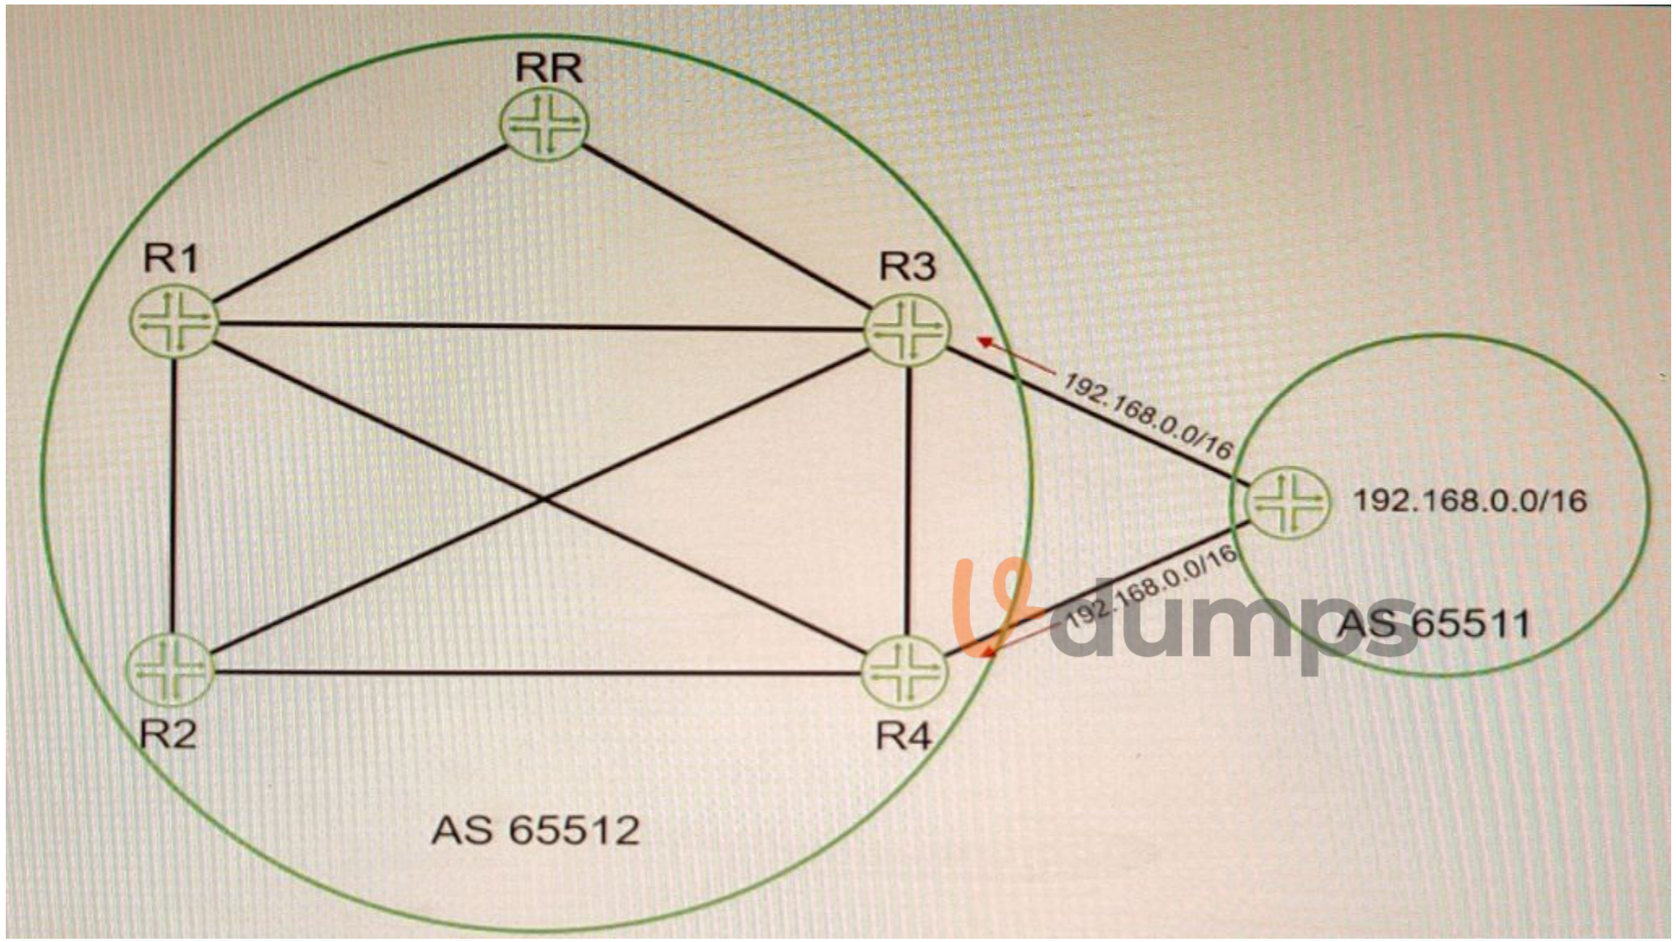

Referring to the exhibit, you are receiving the 192.168 0 0/16 route on both R3 and R4 from your EBGP neighbor You must ensure that R1 and R2 receive both BGP routes from the route reflector
In this scenario, which BGP feature should you configure to accomplish this behavior?

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

Correct Answer: A
Section:
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.

QUESTION 36
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.

Correct Answer: C, D
Section:
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:

QUESTION 37
Exhibit

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


Referring to the exhibit, which statement is true?

A. The 10.101.1.0/24 route will be shared if the vrf-table-label parameter is configured.
B. The 10.101.1.0/24 route will only be shared if BGP is configured in the routing instance
C. The 10.101.1 0/24 route will be shared if there are other VRFs that use the same route target community
D. The 10.101.1.0/24 route will be shared if the auto-export parameter is configured

Correct Answer: D
Section:
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.

QUESTION 38
Exhibit

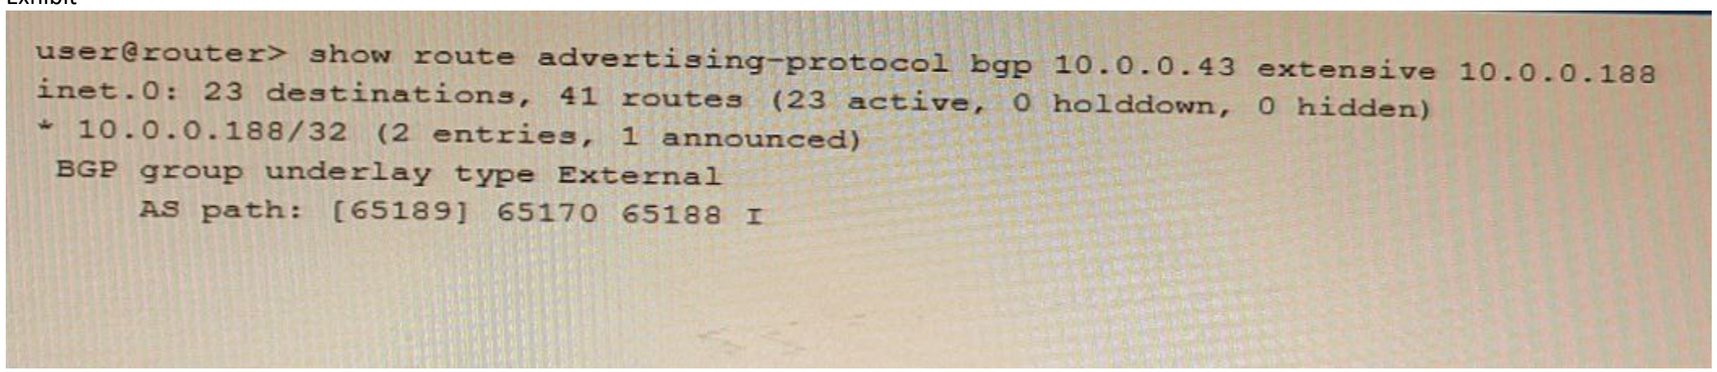

Referring to the exhibit, what do the brackets [ ] in the AS path identify?

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.

Correct Answer: B
Section:
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.

QUESTION 39
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

Correct Answer: D
Section:

QUESTION 40
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.

Correct Answer: A, B
Section:
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.

QUESTION 41
Exhibit

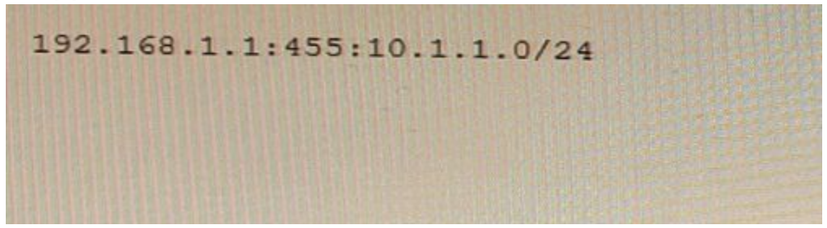

You are examining an L3VPN route that includes the information shown in the exhibit
Which statement is correct in this scenario?

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

Correct Answer: A
Section:
Explanation:
Type 1: When Type value is 1, the Administrator field is 4-bytes and Assigned Number field is 2-bytes. The Administrator field should be set to the IP address (public IP addresses should be used). The Assigned Number field contains a number from a numbering space that is administered by the enterprise to which the IP address has been assigned by the appropriate authority.

QUESTION 42
Exhibit

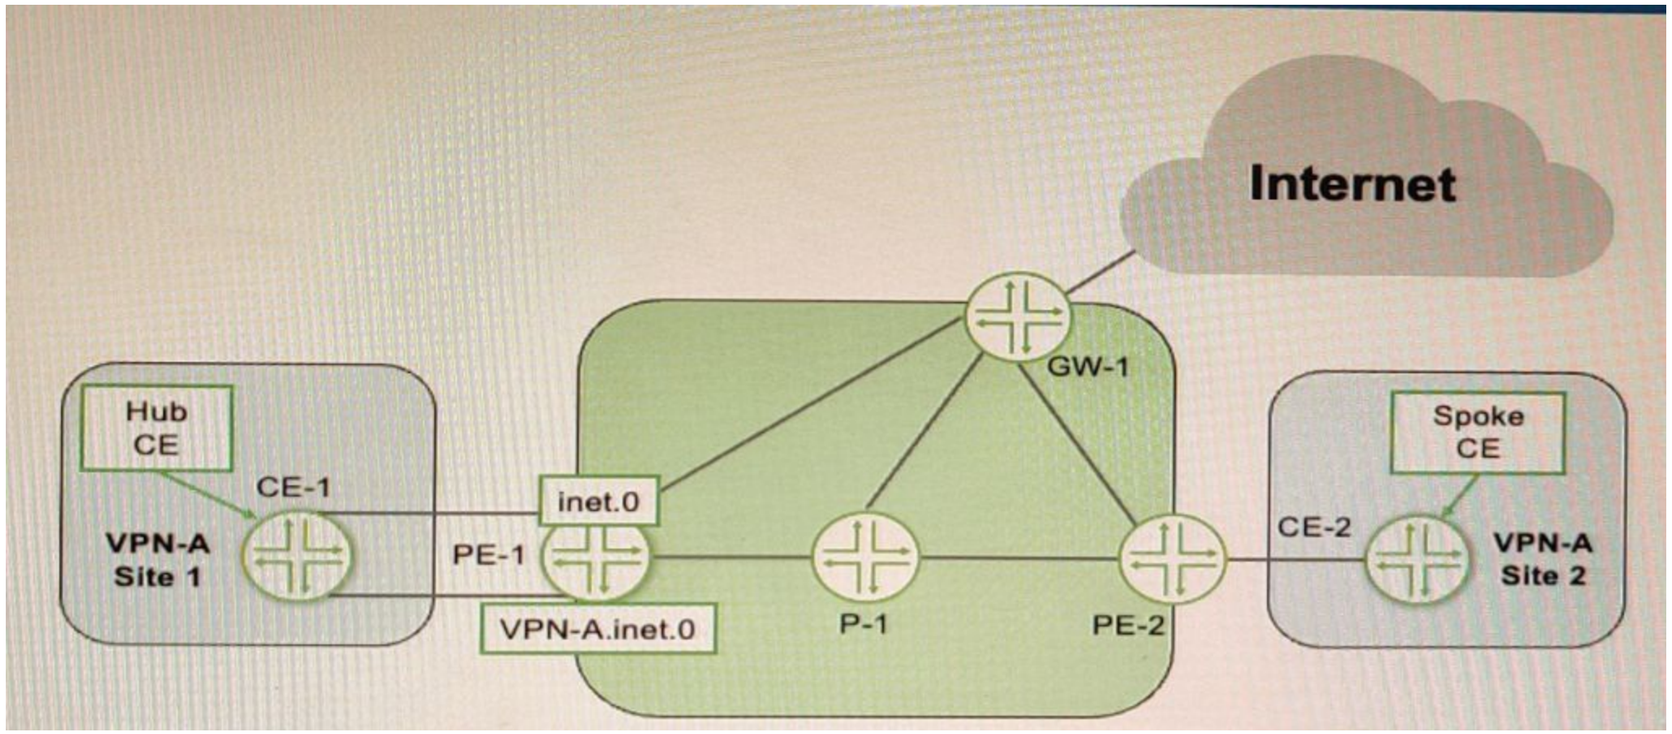

Referring to the exhibit, you must provide Internet access for VPN-A using CE-1 as the hub CE.
Which two statements are correct in this situation? (Choose two.)

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

Correct Answer: A, D
Section:
Explanation:
To provide Internet access for VPN-A using CE-1 as the hub CE, you need to do the following:
You must use RIB groups to leak routes between the inet.0 and vpn-a.inet.0 tables on PE-1 and CE-1. RIB groups are routing options that allow you to import routes from one routing table into another routing table based on certain criteria. In this scenario, you need to configure RIB groups on PE-1 and CE-1 to import Internet routes from inet.0 into vpn-a.inet.0 and vice versa.
Internet traffic from Site 2 takes the path of PE-2 -> PE-1 -> CE-1 -> PE-1 -> GW-1. This is because Site 2 does not have direct Internet access and needs to use CE-1 as its default gateway for Internet traffic. Site 2 sends its Internet traffic to PE-2, which forwards it to PE-1 based on VPN-A routes. PE-1 then sends it to CE-1 based on RIB group import policy. CE-1 then sends it back to PE-1 based on its default route pointing to GW-1. PE-1 then forwards it to GW-1 based on RIB group import policy again.

QUESTION 43
Exhibit

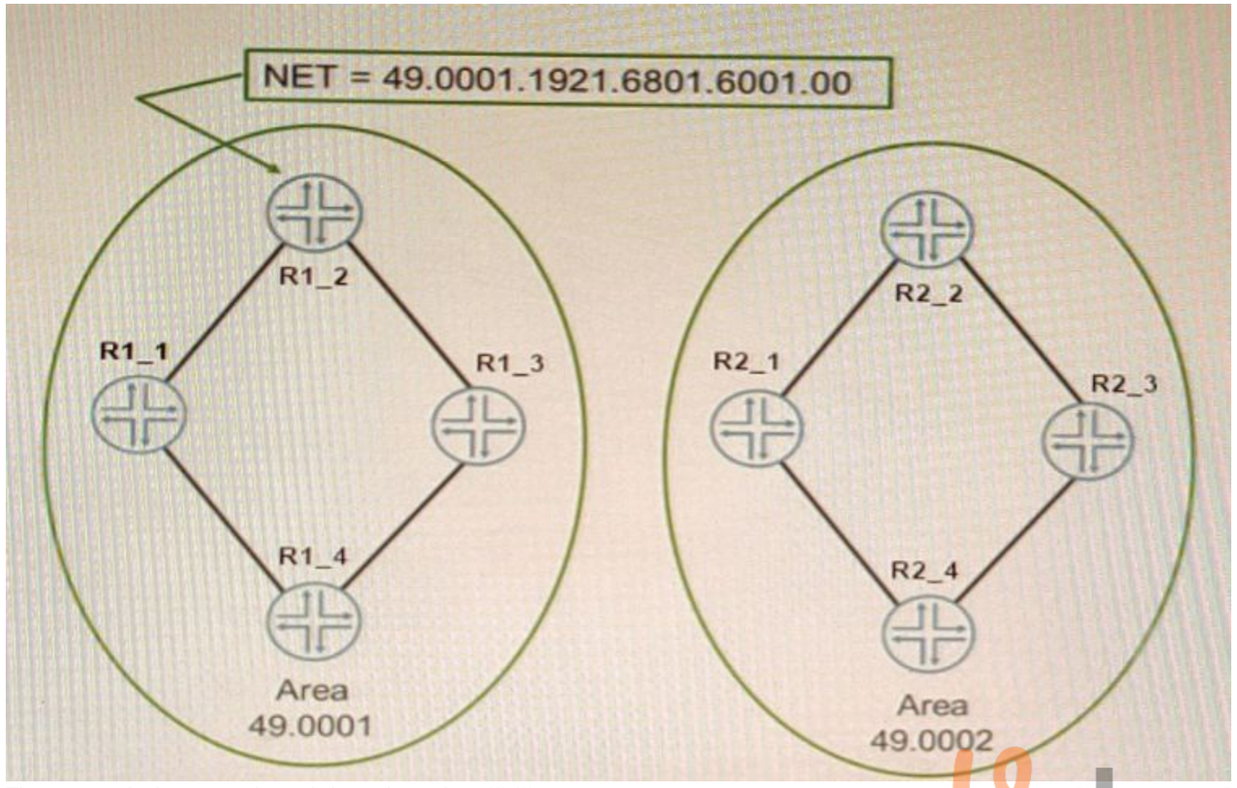

The network shown in the exhibit is based on IS-IS
Which statement is correct in this scenario?

A. The NSEL byte for Area 0001 is 00.
B. The area address is two bytes.
C. The routers are using unnumbered interfaces
D. The system IDofR1_2 is 192.168.16.1

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

QUESTION 44
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

Correct Answer: D
Section:
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.

QUESTION 45
Which two statements are correct regarding bootstrap messages that are forwarded within a PIM sparse mode domain? (Choose two.)

A. Bootstrap messages are forwarded only to routers that explicitly requested the messages within the PIM sparse-mode domain
B. Bootstrap messages distribute RP information dynamically during an RP election.
C. Bootstrap messages are used to notify which router is the PIM RP
D. Bootstrap messages are forwarded to all routers within a PIM sparse-mode domain.

Correct Answer: B, D
Section:

Explanation:
Bootstrap messages are PIM messages that are used to distribute rendezvous point (RP) information dynamically during an RP election. Bootstrap messages are sent by bootstrap routers (BSRs), which are routers that are elected to perform the RP discovery function for a PIM sparse-mode domain. Bootstrap messages contain information about candidate RPs and their multicast groups, as well as BSR priority and hash mask length. Bootstrap messages are forwarded to all routers within a PIM sparse-mode domain using hop-by-hop flooding.

QUESTION 46
Exhibit

```
user@PE1# show routing-instances
VPN-A {
    instance-type vrf;
    interface ge-0/0/1.0;
    vrf-target target:64512:1234;
    protocols {
        bgp {
            group CE {
                type external;
                family inet {
                    unicast;
                }
                neighbor 10.0.0.1 {
                    peer-as 64512;
                    as-override;
                }
            }
        }
    }
}
```


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

A. This VPN connects customer sites that use different AS numbers.
B. This VPN connects customer sites that use the same AS number
C. A Layer 2 VPN is configured.
D. A Layer 3 VPN is configured.

Correct Answer: A, D
Section:
Explanation:
The configuration shown in the exhibit is for a Layer 3 VPN that connects customer sites that use different AS numbers. A Layer 3 VPN is a type of VPN that uses MPLS labels to forward packets across a provider network and BGP to exchange routing information between PE routers and CE routers. A Layer 3 VPN allows customers to use different routing protocols and AS numbers at their sites, as long as they can peer with BGP at the PE-CE interface. In this example, CE-1 is using AS 65530 and CE-2 is using AS 65531, but they can still communicate through the VPN because they have BGP sessions with PE-1 and PE-2, respectively.

QUESTION 47

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

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

Correct Answer: A, C
Section:
Explanation:
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 checksums. CSNPs contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases.

QUESTION 48
Exhibit

```
(65001)R1------R2------R3(65003)
[edit]
user@R2# run show route 11.11.11.0/24
inet.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
11.11.11.0/24      *[BGP/170] 00:00:50, localpref 100
                      AS path: 65001 I, validation-state: unverified
                    >  to 172.16.1.1 via ge-0/0/0.0
                    [BGP/170] 00:00:50, localpref 100
                      AS path: 65003 I, validation-state: unverified
[edit]
user@R2# show protocols bgp
group R1 {
    neighbor 172.16.1.1 {
        peer-as 65001;
    }
}
group R3 {
    neighbor 172.16.2.1 {
        peer-as 65003;
    }
}
local-as 65002;
[edit]
user@R2# show policy-options
policy-statement lb {
    then {
        load-balance per-packet;
    }
}
policy-statement prepend {
    term 1 {
        then as-path-prepend 65001;
    }
}
[edit]
user@R2# show routing-options
forwarding-table {
    export lb;
}
```

R2 is receiving the same route from R1 and R3. You must ensure that you can load balance traffic for that route.
Referring to the exhibit, which configuration change will allow load balancing?

A. Configure the multipath parameter under the global BGP configuration.
B. Apply the prepend policy as an import policy under group R1.
C. Configure the multipath multiple-as parameter under the global BGP configuration.
D. Apply the prepend policy as an import policy under group R3.

Correct Answer: C
Section:

QUESTION 49
Exhibit

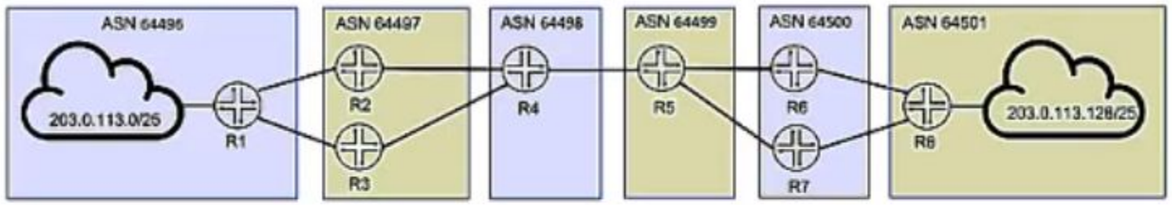


```
user@R1> show route forwarding-table matching 203.0.113.128/25
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
203.0.113.128/28   user     0 10.1.1.1           ucst      576    11 ge-0/0/4.0
203.0.113.144/28   user     0 10.1.1.1           ucst      576    11 ge-0/0/4.0
203.0.113.160/28   user     0 10.1.1.1           ucst      576    11 ge-0/0/4.0
203.0.113.176/28   user     0 10.1.1.1           ucst      576    11 ge-0/0/4.0
203.0.113.192/28   user     0 10.1.1.1           ucst      576    11 ge-0/0/4.0
203.0.113.208/28   user     0 10.1.1.1           ucst      576    11 ge-0/0/4.0
203.0.113.224/28   user     0 10.1.1.1           ucst      576    11 ge-0/0/4.0
203.0.113.240/28   user     0 10.1.1.1           ucst      576    11 ge-0/0/4.0
```

You are troubleshooting the connection between AS 64496 and AS 64497 and notice that only one of the paths is being used for traffic forwarding.
Referring to the exhibit, which three actions will ensure that R1 is configured properly for load balancing BGP routes? (Choose three.)

A. Verify that the routing table on R1 has BGP routes for 203.0.113.128/25 with multiple next hops.
B. Verify that the multipath option is configured under protocols bgp on both R2 and R3.
C. Verify that there is a load balancing export policy under routing-options for the received BGP routes on R1.
D. Verify that the multipath option is configured under protocols bgp on R1.
E. Verify that an import load balancing policy exists under protocols bgp for the received BGP routes on R1.

Correct Answer: A, B, D
Section: