Free Questions for JN0-664 by actualtestdumps

Shared by Solomon on 29-01-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Exhibit

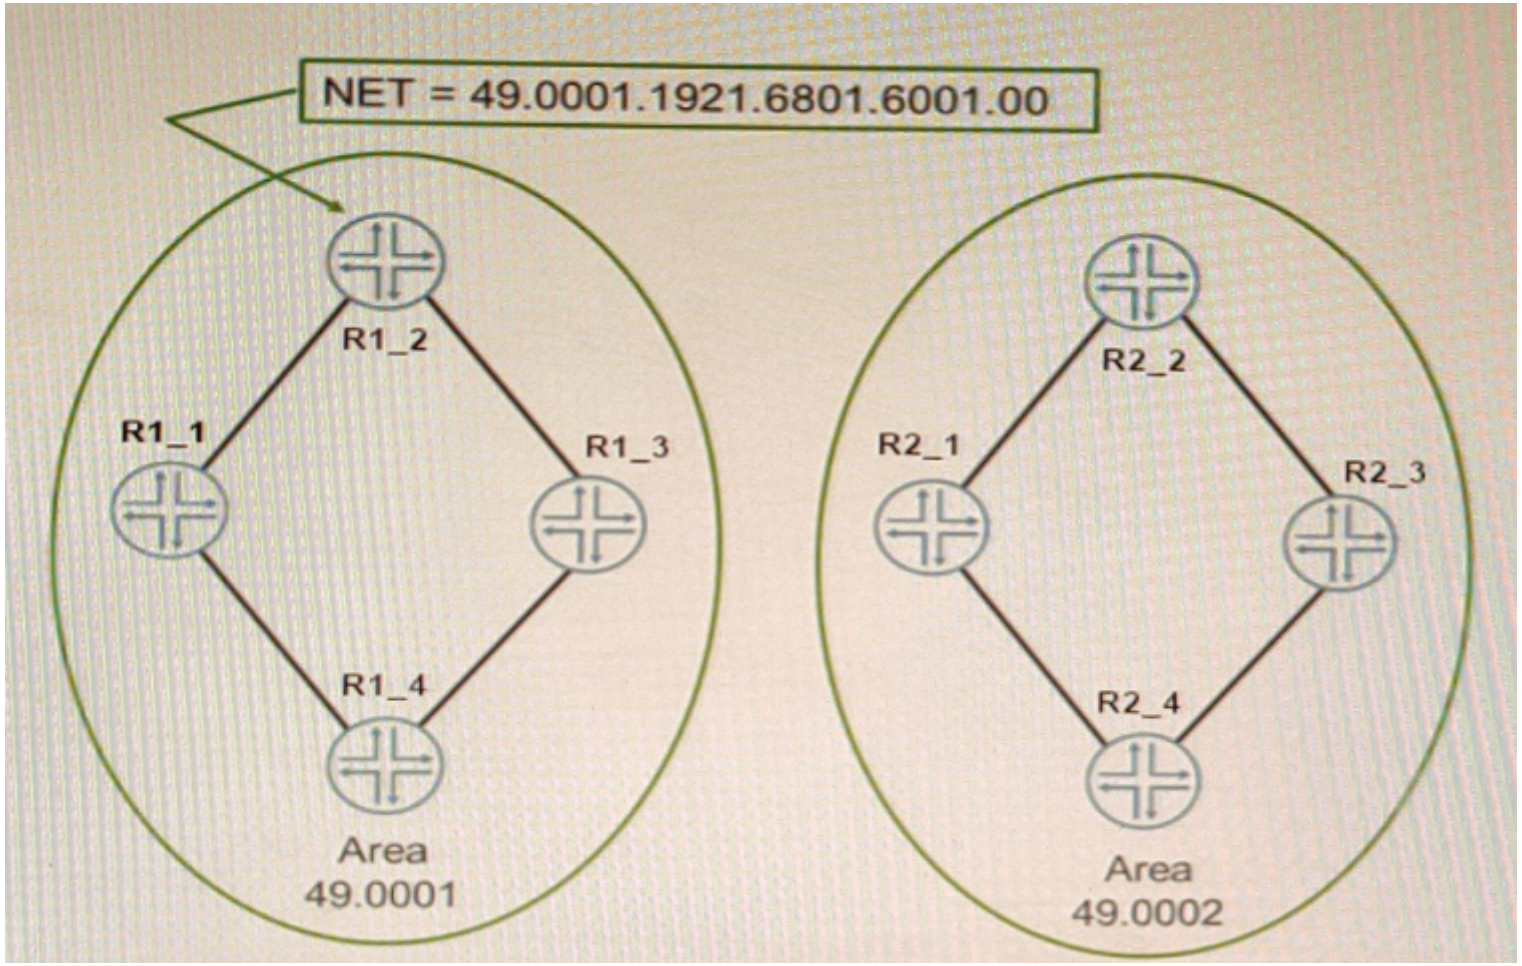

The network shown in the exhibit is based on IS-IS

Which statement is correct in this scenario?

Options:

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

Answer:

A

Explanation:

IS-IS is an interior gateway protocol that uses link-state routing to exchange routing information among routers within a single autonomous system. IS-IS uses two types of addresses to identify routers and areas: system ID and area address. The system ID is a unique identifier for each router in an IS-IS domain. The system ID is 6 octets long and can be derived from the MAC address or manually configured. The area address is a variable-length identifier for each area in an IS-IS domain. The area address can be 1 to 13 octets long and is composed of high-order octets of the address. An IS-IS instance may be assigned multiple area addresses, which are considered synonymous.Multiple synonymous area addresses are useful when merging or splitting areas in the 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 2

Question Type: MultipleChoice

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?

Options:

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.

Question 3

Question Type: MultipleChoice

When using OSPFv3 for an IPv4 environment, which statement is correct?

Options:

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.

Question 4

Question Type: MultipleChoice

Exhibit

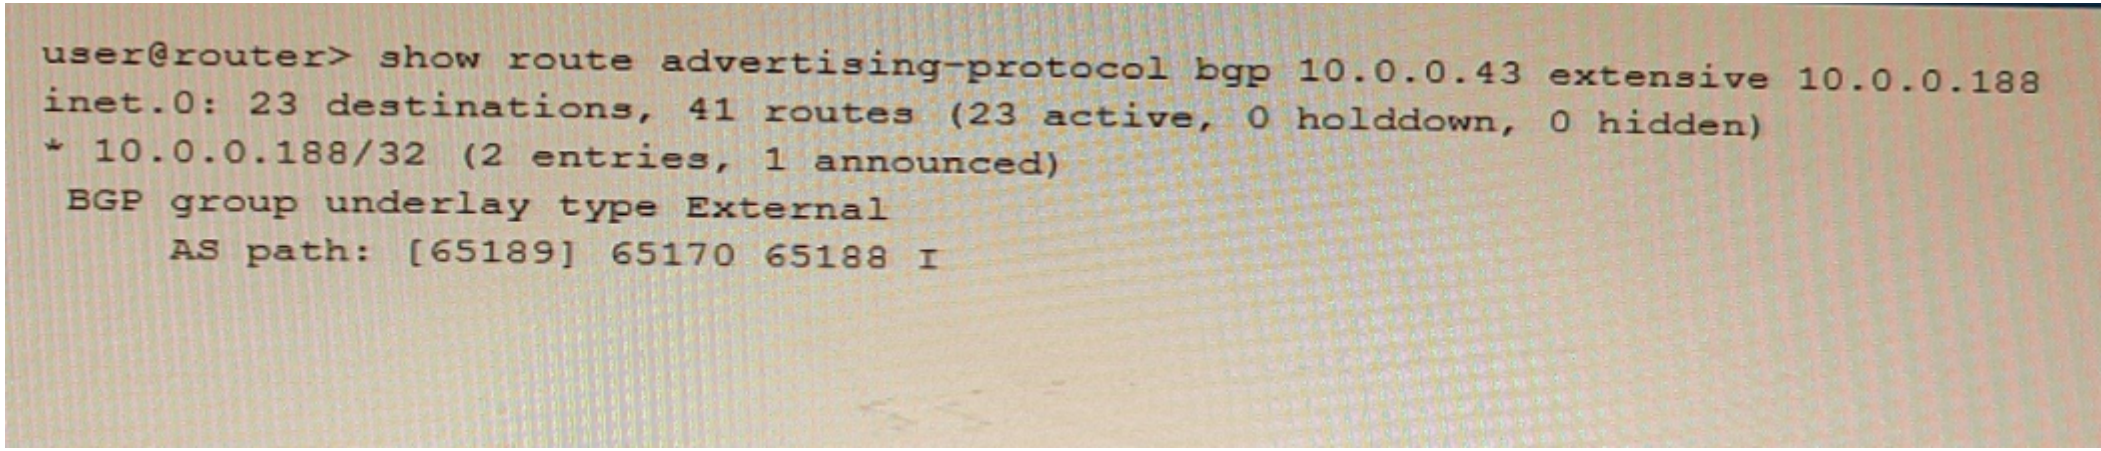

Referring to the exhibit, what do the brackets [ ] in the AS path identify?

Options:

A- They identify the local AS number associated with the AS path if configured on the router, or if AS path prepending is configured

B- They identify an AS set, which are groups of AS numbers in which the order does not matter

C- They identify that the autonomous system number is incomplete and awaiting more information from the BGP protocol.

D- They identify that a BGP confederation is being used to ensure that there are no routing loops.

Answer:

B

Explanation:

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

Question 5

Question Type: **MultipleChoice**

Exhibit

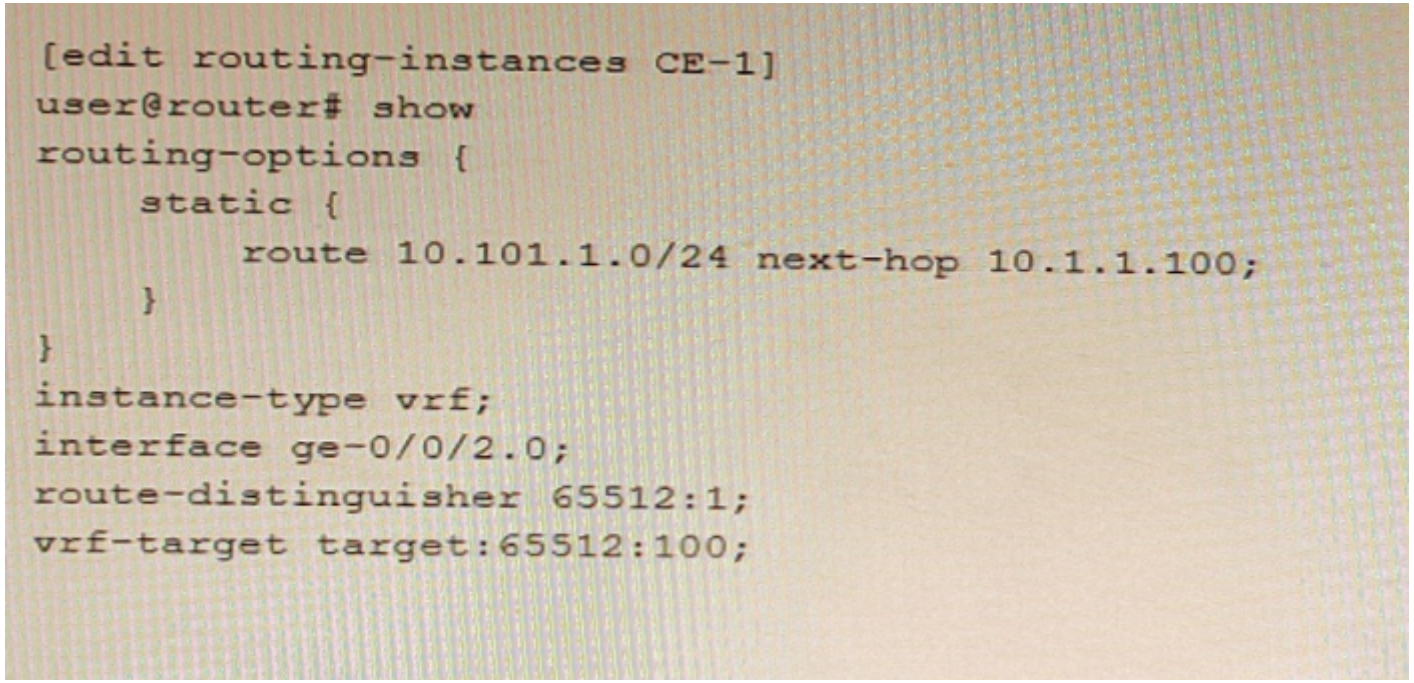

```
[edit routing-instances CE-1]
user@router# show
routing-options {
    static {
        route 10.101.1.0/24 next-hop 10.1.1.100;
    }
}
instance-type vrf;
interface ge-0/0/2.0;
route-distinguisher 65512:1;
vrf-target target:65512:100;
```

Referring to the exhibit, which statement is true?

Options:

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.

Question 6

Question Type: MultipleChoice

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

Options:

A- When the encapsulation is vlan-ccc or extended-vlan-ccc, the configured VLAN tag is not included in the control plane LDP advertisement

B- When the encapsulation is ethernet-ccc, only frames without a VLAN tag are accepted in the data plane

C- When the encapsulation is vLan-ccc or extended-vlan-ccc, the configured VLAN tag is included in the control plane LDP advertisement

D- When the encapsulation is ethemet-ccc, tagged and untagged frames are both accepted in the data plane.

Answer:

C, D

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 7

Question Type: **MultipleChoice**

Exhibit

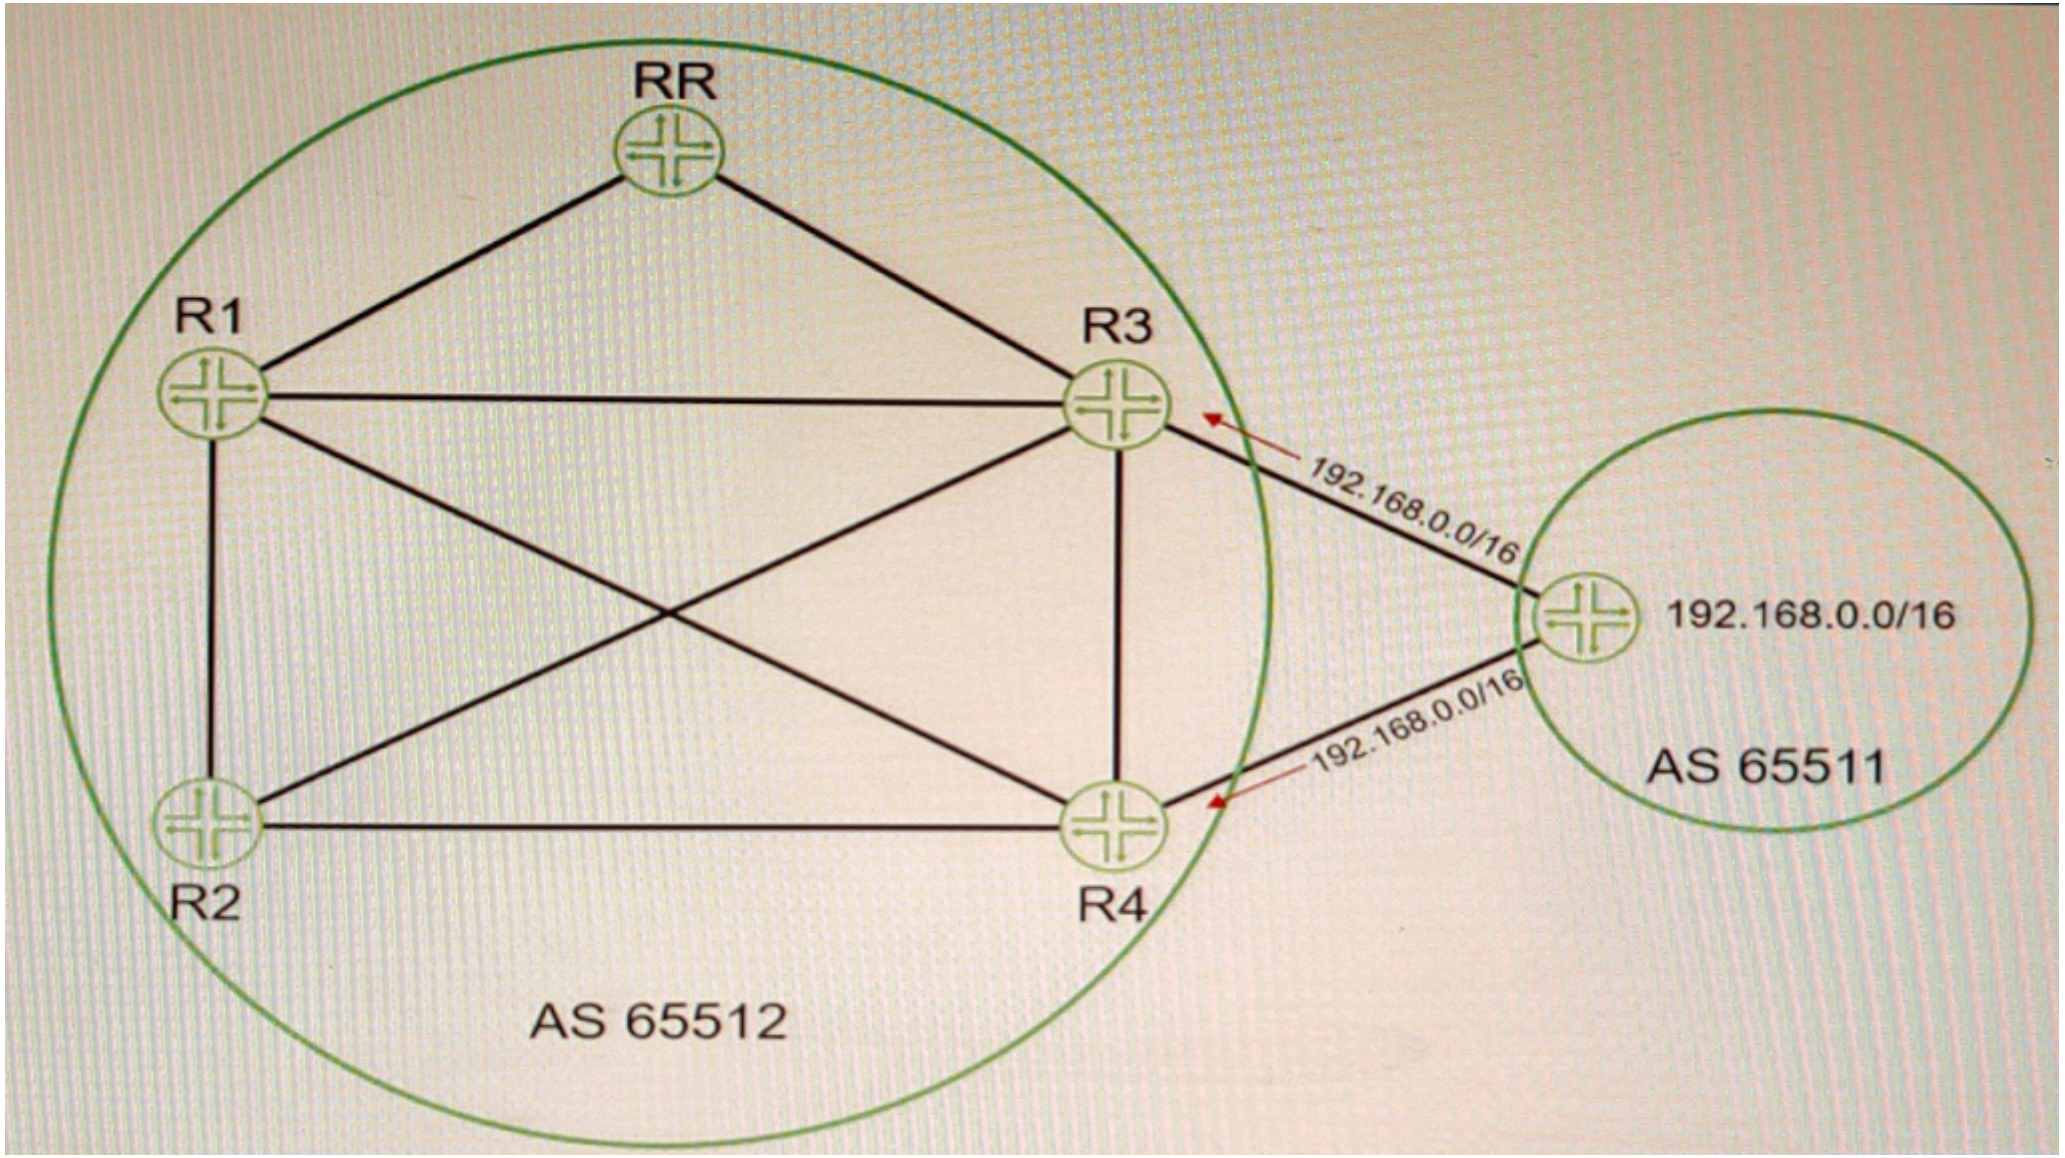
RR
R1
R3
192.168.0.0/16
192.168.0.0/16
192.168.0.0/16
AS 65511
R2
R4
AS 65512

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?

Options:

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.

Question 8

Question Type: MultipleChoice

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

Options:

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.

Answer:

C

Explanation:

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

Question 9

Question Type: MultipleChoice

Exhibit

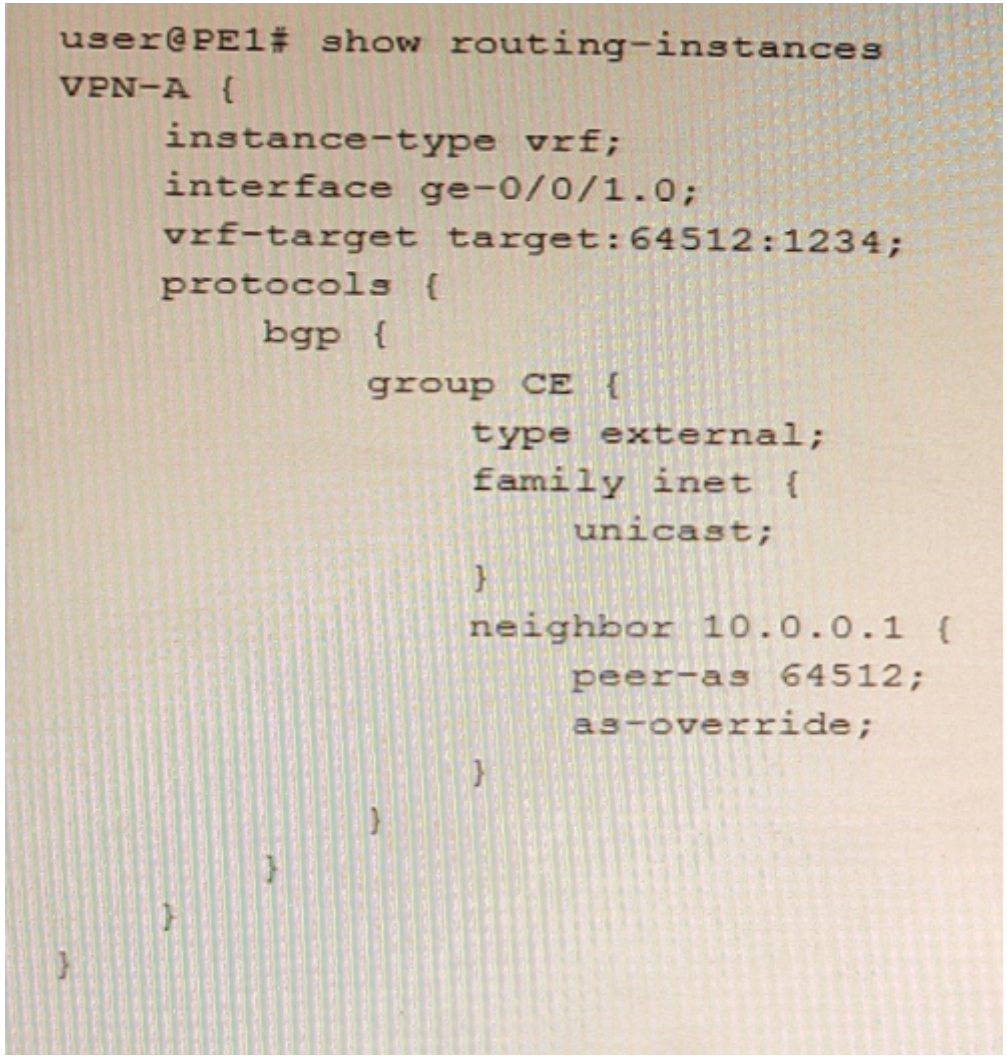

```
user@PE1# show routing-instances
VPN-A {
    instance-type vrf;
    interface ge-0/0/1.0;
    vrf-target target:64512:1234;
    protocols {
        bgp {
            group CE {
                type external;
                family inet {
                    unicast;
                }
           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

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

Options:

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.

Answer:

A, D

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.