Juniper

Exam Questions JN0-664

Service Provider - Professional (JNCIP-SP)

NEW QUESTION 1
You are configuring a BGP signaled Layer 2 VPN across your MPLS enabled core network. In this scenario, which statement is correct?

A. You must assign a unique site number to each attached site's configuration.
B. This type of VPN only supports Ethernet interfaces when connecting to CE devices.
C. This type of VPN requires the support of the inet-vpn NLRI on all core BGP devices
D. You must use the same route-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. References: 1: https://www.juniper.net/documentation/us/en/software/junos/vpn-l2/topics/concept/vpn-layer-2-overview.html 2: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l2_vpns/configuration/xe-16/mp-l2-vpns-xe-16-book/vpls-bgp-signaling-l2vpn- inter-as-option-a.html

NEW QUESTION 2
Exhibit

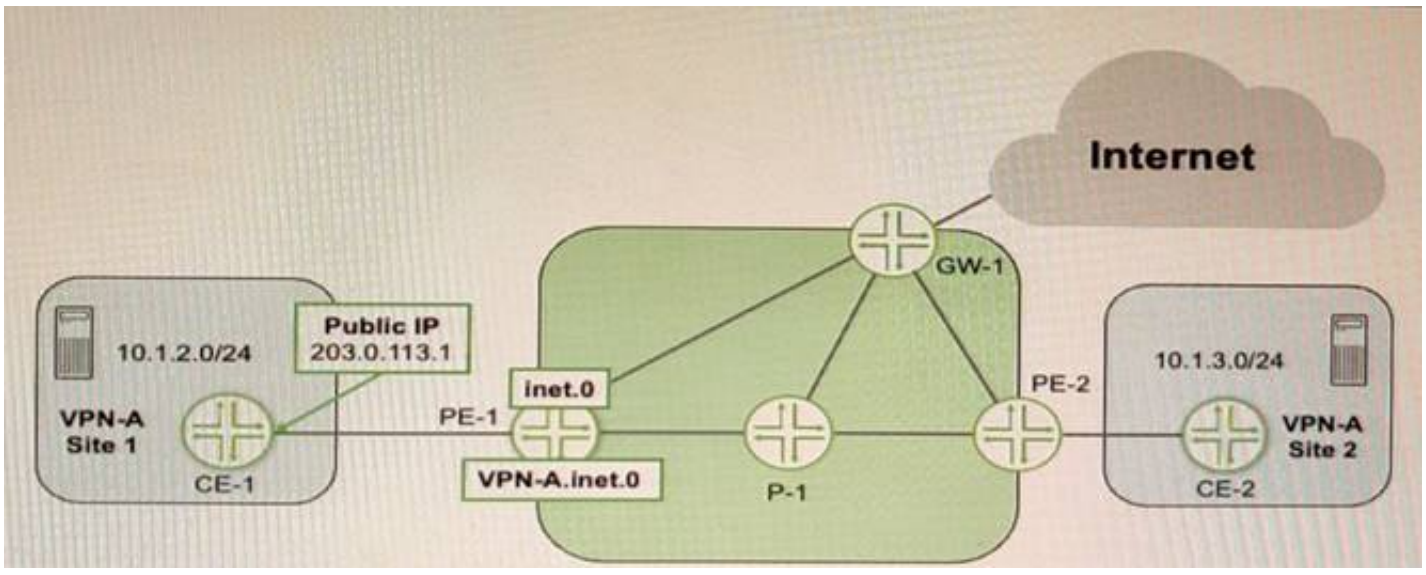

Referring to the exhibit, CE-1 is providing NAT services for the hosts at Site 1 and you must provide Internet access for those hosts
Which two statements are correct in this scenario? (Choose two.)

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 two statements describe PIM-SM? (Choose two)

A. Routers with receivers send join messages to their upstream neighbors.
B. Routers without receivers must periodically prune themselves from the SPT.
C. Traffic is initially flooded to all routers and an S,G is maintained for each group
D. Traffic is only forwarded to routers that request to join the distribution tree.

Answer: AD

Explanation:
PIM sparse mode (PIM-SM) is a multicast routing protocol that uses a pull model to deliver multicast traffic. In PIM-SM, routers with receivers send join messages to their upstream neighbors toward a rendezvous point (RP) or a source-specific tree (SPT). The RP or SPT acts as the root of a shared distribution tree for a multicast group. Traffic is only forwarded to routers that request to join the distribution tree by sending join messages. PIM-SM does not flood traffic to all routers or prune routers without receivers, as PIM dense mode does.

NEW QUESTION 4
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 5
Exhibit

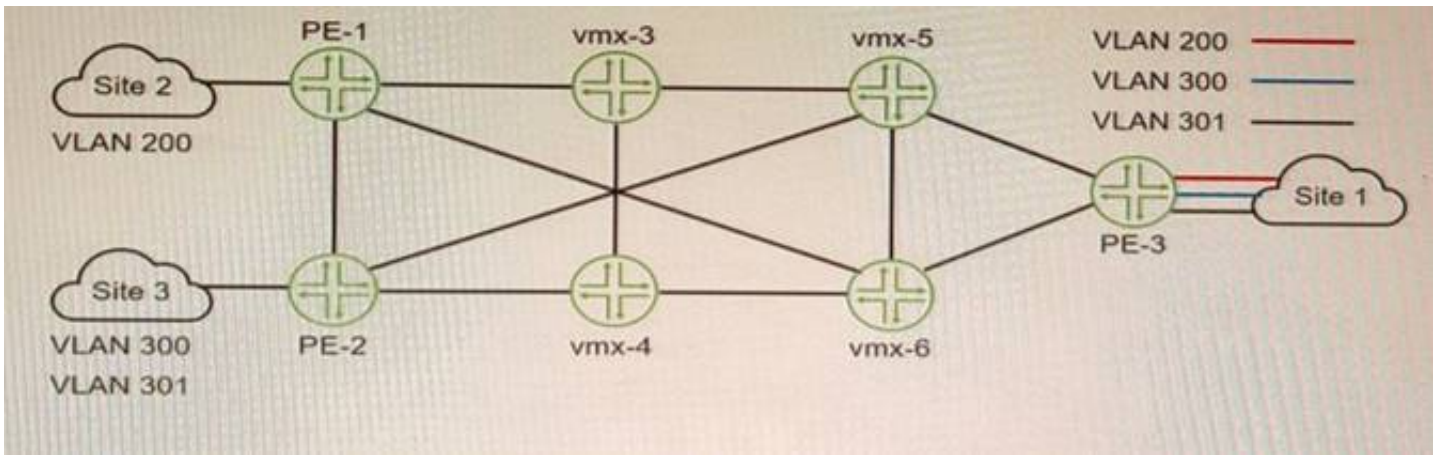

You want Site 1 to access three VLANs that are located in Site 2 and Site 3 The customer- facing interface on the PE-1 router is configured for Ethernet-VLAN encapsulation.
What is the minimum number of L2VPN routing instances to be configured to accomplish this task?

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 6
Exhibit.

Exhibit

```
user@R1# show interfaces
ge-1/2/3 {
    unit 0 {
        description to-R2;
        family inet {
            address 10.1.1.1/30;
        }
        family iso;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.16.1/32;
        }
        family iso {
            address 49.0001.1921.6801.6001.00;
        }
    }
}
user@R1# show protocols
isis {
    interface ge-1/2/3.0 {
        level 2 disable;
    }
}
...
user@R2# show interfaces
ge-1/2/3 {
    unit 0 {
        description to-R1;
        family inet {
            address 10.1.1.2/30;
        }
        family iso;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.16.2/32;
        }
        family iso {
            address 49.0001.1921.6801.6002.00;
        }
    }
}
user@R2# show protocols
isis {
    interface ge-1/2/3.0 {
        level 1 disable;
    }
    interface lo0.0 {
        level 1 disable;
    }
}
```

Referring to the exhib.t, what must be changed to establish a Level 1 adjacency between routers R1 and R2?

A. Change the level l disable parameter under the R1 protocols isis interface lo0.0 hierarchy to the level 2 disable parameter.
B. Remove the level i disable parameter under the R2 protocols isis interface loo . 0 configuration hierarchy.
C. Change the level 1 disable parameter under the R2 protocols isis interface ge-1/2/3 .0 hierarchy to the level 2 disable parameter
D. Add IP addresses to the interface ge-l/2/3 unit 0 family iso hierarchy on both R1 and R2.

Answer: B

Explanation:
IS-IS routers can form Level 1 or Level 2 adjacencies depending on their configuration and network topology. Level 1 routers are intra-area routers that share the same area address with their neighbors. Level 2 routers are inter-area routers that can connect different areas. Level 1-2 routers are both intra-area and inter-area routers that can form adjacencies with any other router.

In the exhibit, R1 and R2 are in different areas (49.0001 and 49.0002), so they cannot form a Level 1 adjacency. However, they can form a Level 2 adjacency if they are both configured as Level 1-2 routers. R1 is already configured as a Level 1-2 router, but R2 is configured as a Level 1 router only, because of the level 1 disable command under the lo0.0 interface. This command disables Level 2 routing on the loopback interface, which is used as the router ID for IS-IS. Therefore, to establish a Level 1 adjacency between R1 and R2, the level 1 disable command under the R2 protocols isis interface lo0.0 hierarchy must be removed. This will enable Level 2 routing on R2 and allow it to form a Level 2 adjacency with R1.

NEW QUESTION 7
Exhibit
user@Rl show configuration interpolated-profile { interpolate {
fill-level [50 75 drop—probability [> }
class-of-service drop-profiles
];
20 60];
Which two statements are correct about the class-of-service configuration shown in the exhibit? (Choose two.)

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

Answer: BC

Explanation:
class-of-service (CoS) is a feature that allows you to prioritize and manage network traffic based on various criteria, such as application type, user group, or packet loss priority. CoS uses different components to classify, mark, queue, schedule, shape, and drop traffic according to the configured policies.
One of the components of CoS is drop profiles, which define how packets are dropped when a queue is congested. Drop profiles use random early detection (RED) algorithm to drop packets randomly before the queue is full, which helps to avoid global synchronization and improve network performance. Drop profiles can be discrete or interpolated. A discrete drop profile maps a specific fill level of a queue to a specific drop probability. An interpolated drop profile maps a range of fill levels of a queue to a range of drop probabilities and interpolates the values in between.
In the exhibit, we can see that the class-of-service configuration shows an interpolated drop profile with two fill levels (50 and 75) and two drop probabilities (20 and 60). Based on this configuration, we can infer the following statements:
? The drop probability jumps immediately from 20% to 60% when the queue level reaches 75% full. This is not correct because the drop profile is interpolated, not discrete. This means that the drop probability gradually increases from 20% to 60% as the queue level increases from 50% full to 75% full. The drop probability for any fill level between 50% and 75% can be calculated by using linear interpolation formula.
? The drop probability gradually increases from 20% to 60% as the queue level increases from 50% full to 75% full. This is correct because the drop profile is interpolated and uses linear interpolation formula to calculate the drop probability for any fill level between 50% and 75%. For example, if the fill level is 60%, the drop probability is 28%, which is calculated by using the formula: (60 - 50) / (75 - 50) * (60 - 20) + 20 = 28.
? To use this drop profile, you reference it in a scheduler. This is correct because a scheduler is a component of CoS that determines how packets are dequeued from different queues and transmitted on an interface. A scheduler can reference a drop profile by using the random-detect statement under the [edit class-of-service schedulers] hierarchy level. For example: scheduler test { transmit-rate percent 10; buffer-size percent 10; random-detect test-profile; }
? To use this drop profile, you apply it directly to an interface. This is not correct because a drop profile cannot be applied directly to an interface. A drop profile can only be referenced by a scheduler, which can be applied to an interface by using the scheduler-map statement under the [edit class-of-service interfaces] hierarchy level. For example: interfaces ge-0/0/0 { unit 0 { scheduler-map test-map; } }

NEW QUESTION 8
Exhibit

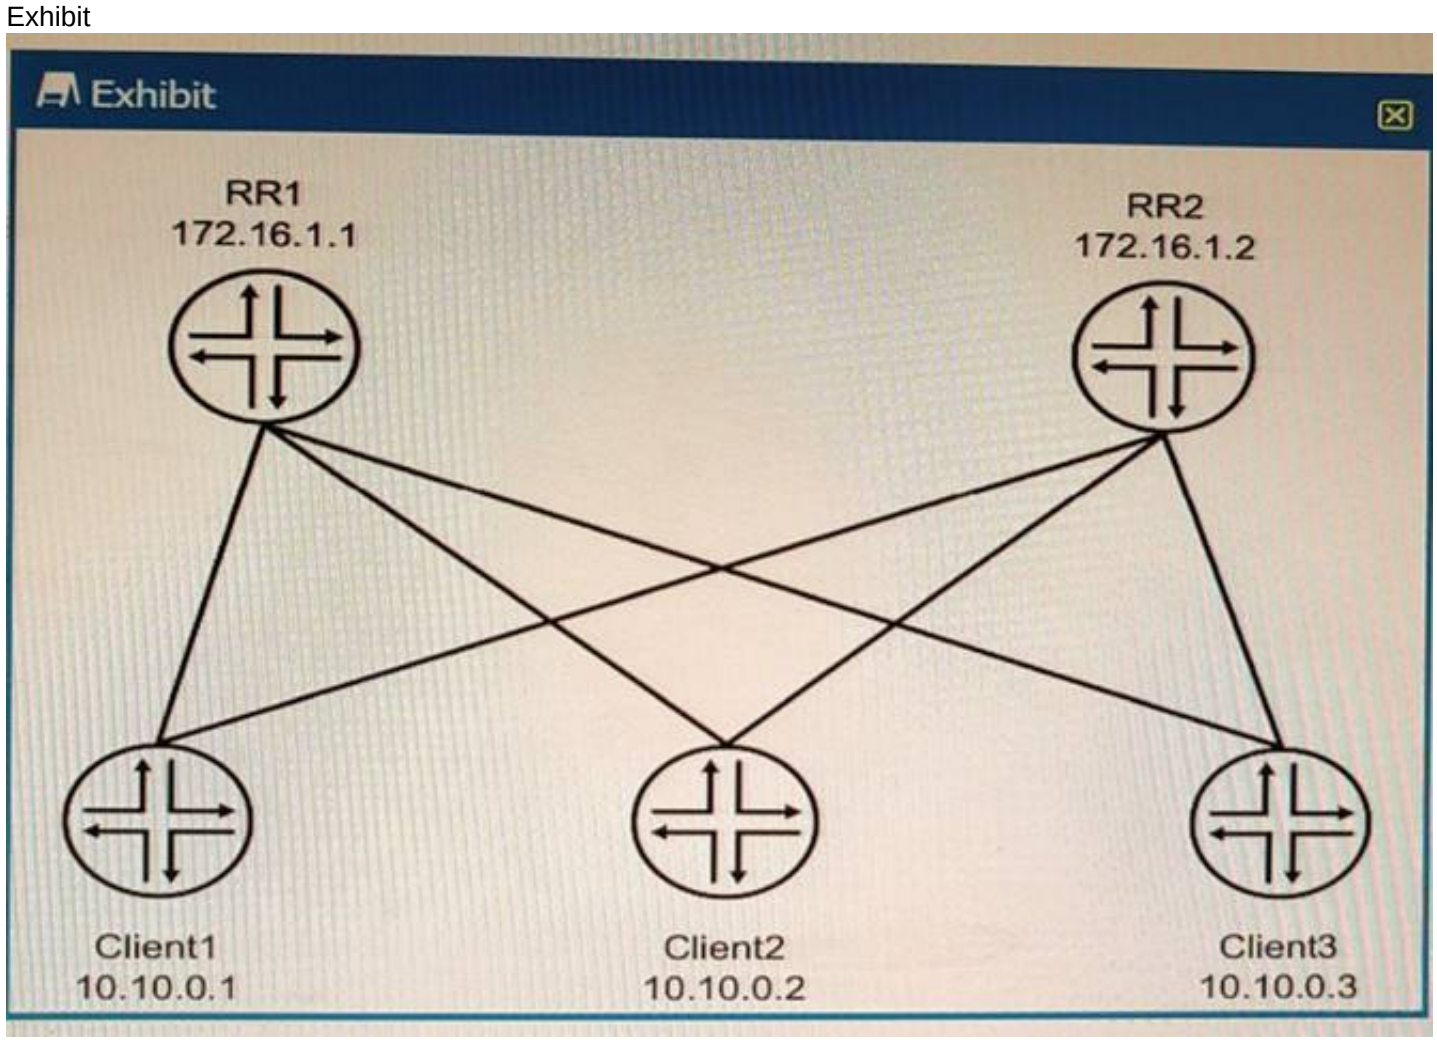

The environment is using BGP All devices are in the same AS with reachability redundancy Referring to the exhibit, which statement is correct?

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 9
Exhibit

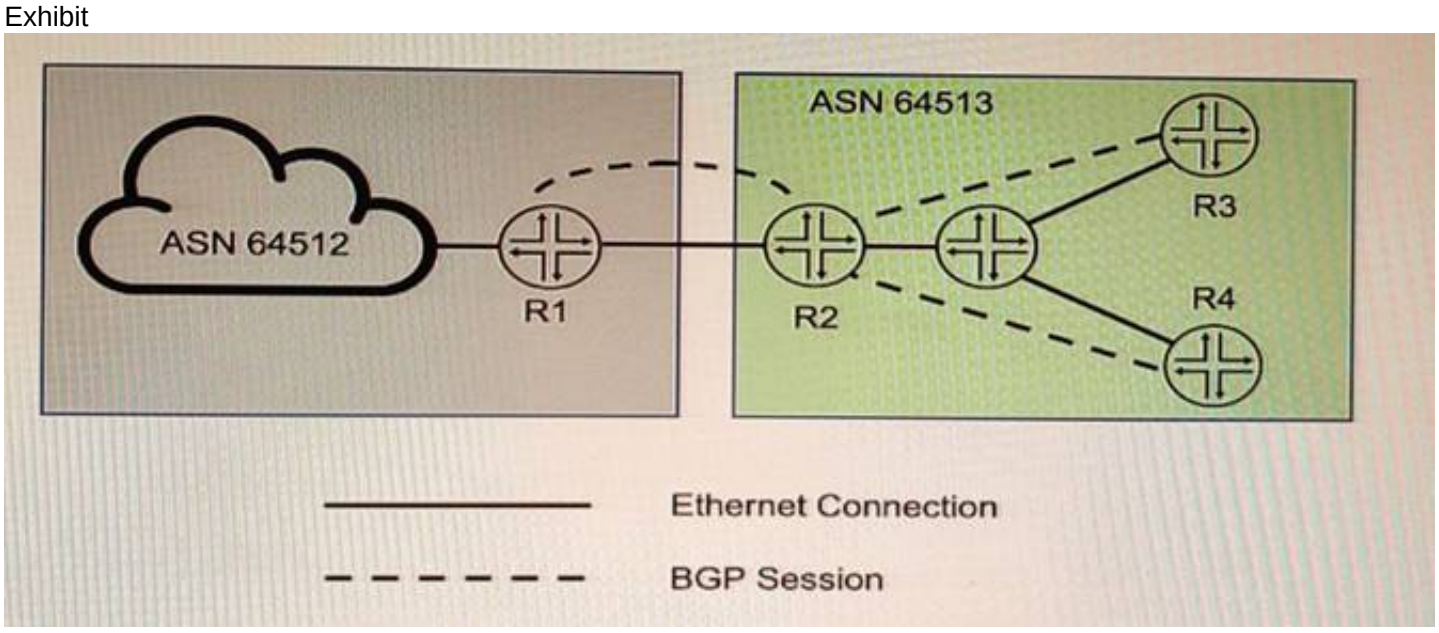

You want to implement the BGP Generalized TTL Security Mechanism (GTSM) on the network
Which three statements are correct in this scenario? (Choose three)

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

Answer: ADE

Explanation:
BGP GTSM is a technique that protects a BGP session by comparing the TTL value in the IP header of incoming BGP packets against a valid TTL range. If the TTL value is within the valid TTL range, the packet is accepted. If not, the packet is discarded. The valid TTL range is from 255 – the configured hop count + 1 to 255. When GTSM is configured, the BGP packets sent by the device have a TTL of 255. GTSM provides best protection for directly connected EBGP sessions, but not for multihop EBGP or IBGP sessions because the TTL of packets might be modified by intermediate devices.
In the exhibit, we can see that R2, R3, and R4 are in the same AS (AS 20) and R1 is in a different AS (AS 10). Based on this information, we can infer the following statements:
? You can implement BGP GTSM between R2, R3, and R4. This is not correct because R2, R3, and R4 are IBGP peers and GTSM does not provide effective protection for IBGP sessions. The TTL of packets between IBGP peers might be changed by intermediate devices or routing protocols.
? BGP GTSM requires a firewall filter to discard packets with incorrect TTL. This is not correct because BGP GTSM does not require a firewall filter to discard packets with incorrect TTL. BGP GTSM uses TCP option 19 to negotiate GTSM capability between peers and uses TCP option 20 to carry the expected TTL value in each packet. The receiver checks the expected TTL value against the actual TTL value and discards packets with incorrect TTL values.
? You can implement BGP GTSM between R2 and R1. This is correct because R2 and R1 are EBGP peers and GTSM provides effective protection for directly connected EBGP sessions. The TTL of packets between directly connected EBGP peers is not changed by intermediate devices or routing protocols.
? BGP GTSM requires a TTL of 1 to be configured between neighbors. This is not correct because BGP GTSM requires a TTL of 255 to be configured between neighbors. The sender sets the TTL of packets to 255 and the receiver expects the TTL of packets to be 255 minus the configured hop count.
? BGP GTSM requires a TTL of 255 to be configured between neighbors. This is correct because BGP GTSM requires a TTL of 255 to be configured between neighbors. The sender sets the TTL of packets to 255 and the receiver expects the TTL of packets to be 255 minus the configured hop count.

NEW QUESTION 10
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

Answer: D

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.

NEW QUESTION 10
Exhibit

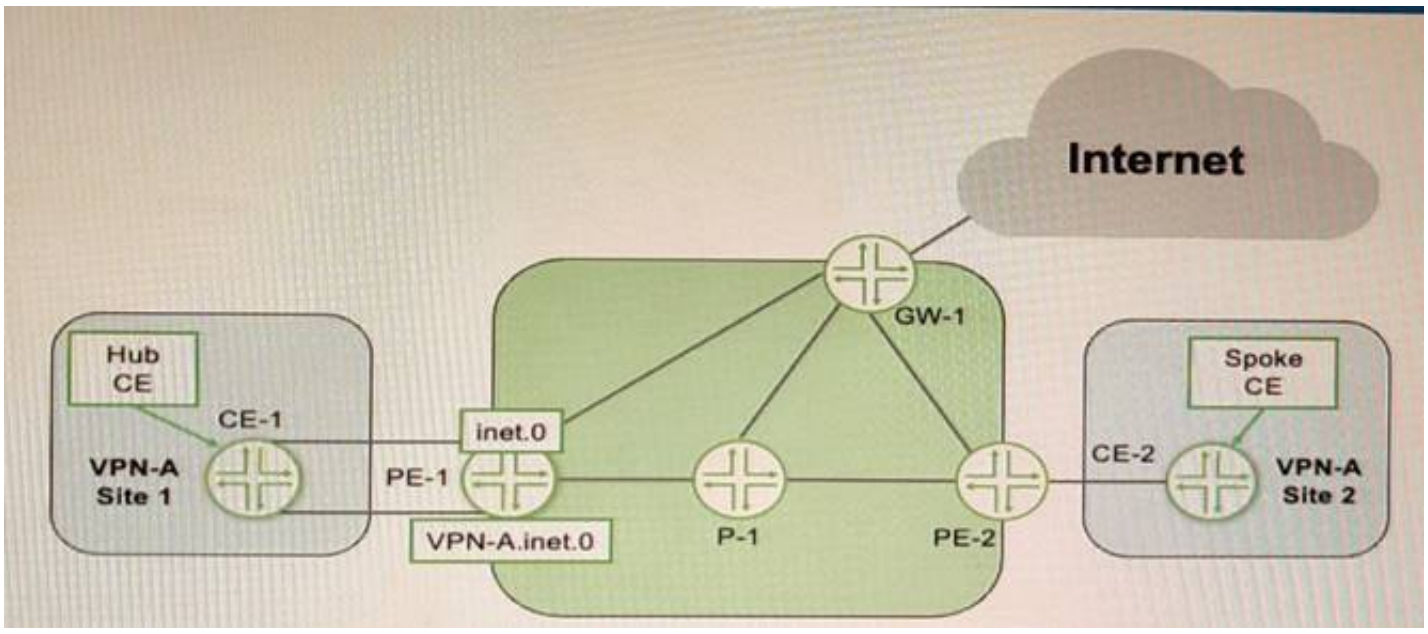

Referring to the exhibit, you must provide Internet access for VPN-A using CE-1 as the hub CE.
Which two statements are correct in this situation? (Choose two.)

A. You must use RIB groups to leak routes between the ine
B. o and vpn-
C. ine
D. o tables.
E. RIB groups are not needed to leak routes between the ine
F. 0 and VPN—
G. ine
H. 0 tables,
I. Internet traffic from Site 2 takes the path of PE-2 -> PE-1 -> GW-1.
J. Internet traffic from Site 2 takes the path of PE-2 -> PE-1 -> CE-1 -> PE-1 -> GW-1.

Answer: AD

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.

NEW QUESTION 11
Which two EVPN route types are used to advertise a multihomed Ethernet segment? (Choose two)

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

Answer: AC

Explanation:
EVPN is a solution that provides Ethernet multipoint services over MPLS networks. EVPN uses BGP to distribute endpoint provisioning information and set up pseudowires between PE devices. EVPN uses different route types to convey different information in the control plane. The following are the main EVPN route types:
? Type 1 - Ethernet Auto-Discovery Route: This route type is used for network-wide messaging and discovery of other PE devices that are part of the same EVPN instance. It also carries information about the redundancy mode and load balancing algorithm of the PE devices.
? Type 2 - MAC/IP Advertisement Route: This route type is used for MAC and IP address learning and advertisement between PE devices. It also carries information about the Ethernet segment identifier (ESI) and the label for forwarding traffic to the MAC or IP address.
? Type 3 - Inclusive Multicast Ethernet Tag Route: This route type is used for broadcast, unknown unicast, and multicast (BUM) traffic forwarding. It also carries information about the multicast group and the label for forwarding BUM traffic.
? Type 4 - Ethernet Segment Route: This route type is used for multihoming scenarios, where a CE device is connected to more than one PE device. It also carries information about the ESI and the designated forwarder (DF) election process.

NEW QUESTION 12
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 17
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.

Answer: C

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.

NEW QUESTION 19
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 22
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.

Answer: AD

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.

NEW QUESTION 27
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 29
In which two ways does OSPF prevent routing loops in multi-area networks? (Choose two.)

A. All areas are required to connect as a full mesh.
B. The LFA algorithm prunes all looped paths within an area.
C. All areas are required to connect to area 0.
D. The SPF algorithm prunes looped paths within an area.

Answer: CD

Explanation:
OSPF is an interior gateway protocol that uses link-state routing to exchange routing information among routers within a single autonomous system. OSPF prevents routing loops in multi-area networks by using two methods: area hierarchy and SPF algorithm. Area hierarchy is the concept of dividing a large OSPF network into smaller areas that are connected to a backbone area (area 0). This reduces the amount of routing information that each router has to store and process, and also limits the scope of link-state updates within each area. All areas are required to connect to area 0 either directly or through virtual links2. SPF algorithm is the method that OSPF uses to calculate the shortest path to each destination in the network based on link-state information. The SPF algorithm runs on each router and builds a shortest-path tree that represents the topology of the network from the router's perspective. The SPF algorithm prunes looped paths within an area by choosing only one best path for each destination3.
References: 2: https://www.juniper.net/documentation/us/en/software/junos/ospf/topics/concept/ospf-area- overview.html 3: https://www.juniper.net/documentation/us/en/software/junos/ospf/topics/concept/ospf-spf- algorithm-overview.html

NEW QUESTION 32
Exhibit.

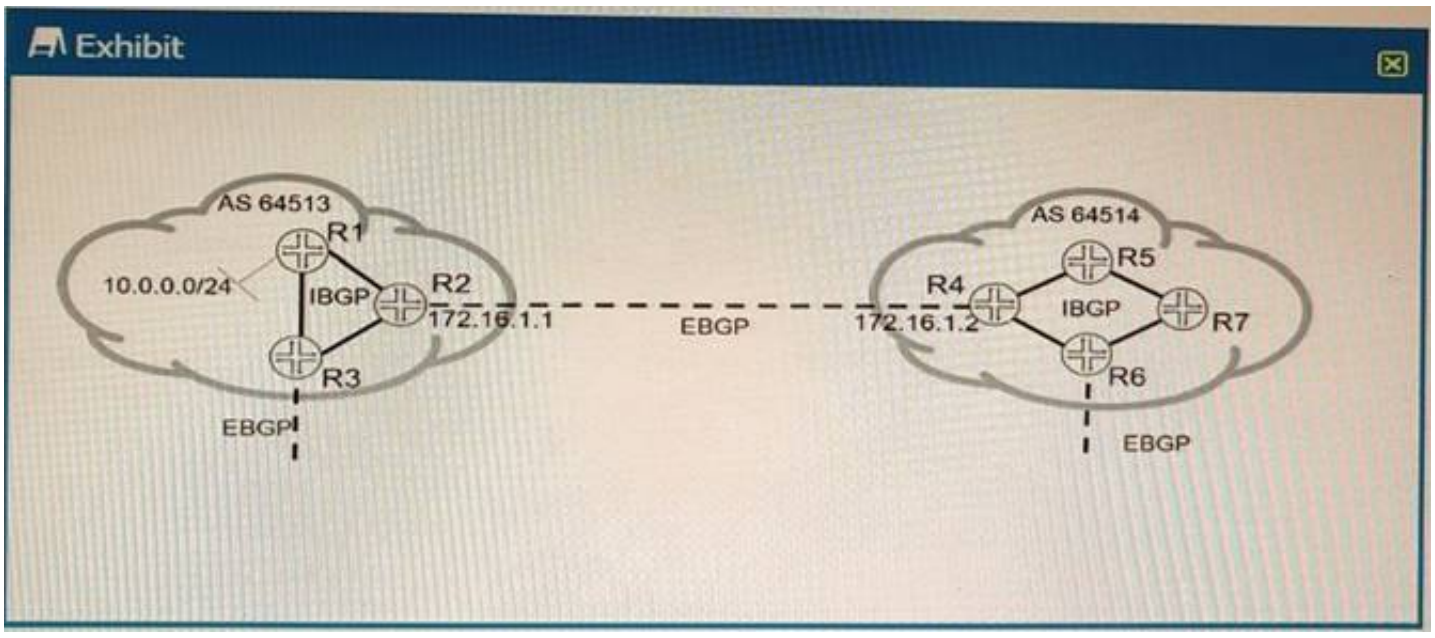

Referring to the exhibit; the 10.0.0.0/24 EBGP route is received on R5; however, the route is being hidden.
What are two solutions that will solve this problem? (Choose two.)

A. On R4, create a policy to change the BGP next hop to itself and apply it to IBGP as an export policy
B. Add the external interface prefix to the IGP routing tables
C. Add the internal interface prefix to the BGP routing tables.
D. On R4, create a policy to change the BGP next hop to 172.16.1.1 and apply it to IBGP as an export policy

Answer: AB

Explanation:
the default behavior for iBGP is to propagate EBGP-learned prefixes without changing the next-hop. This can cause issues if the next-hop is not reachable via the IGP. One solution is to use the next-hop self command on R4, which will change the next-hop attribute to its own loopback address. This way, R5 can reach the next-hop via the IGP and install the route in its routing table.
Another solution is to add the external interface prefix (120.0.4.16/30) to the IGP routing tables of R4 and R5. This will also make the next-hop reachable via the IGP and allow R5 to use the route. According to2, this is a possible workaround for a pure IP network, but it may not work well for an MPLS network.

NEW QUESTION 35
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

Answer: AB

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

NEW QUESTION 37
Exhibit

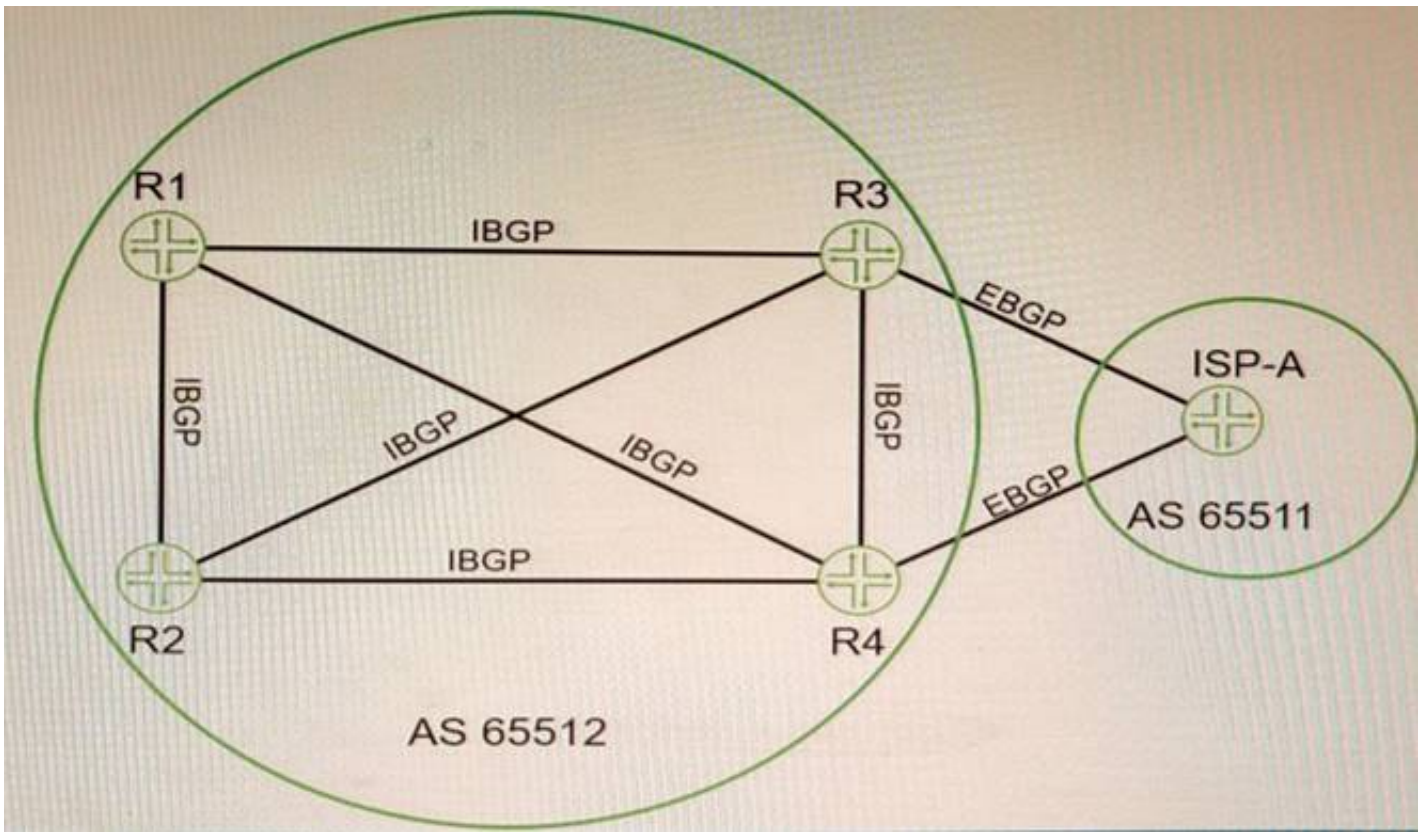

Click the Exhibit button-Referring to the exhibit, which two statements are correct about BGP routes on R3 that are learned from the ISP-A neighbor? (Choose two.)

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

Answer: AB

Explanation:
BGP is an exterior gateway protocol that uses path vector routing to exchange routing information among autonomous systems. BGP uses various attributes to select the best path to each destination and to propagate routing policies. Some of the common BGP attributes are AS path, next hop, local preference, MED, origin, weight, and community. BGP attributes can be classified into four categories: well-known mandatory, well-known discretionary, optional transitive, and optional nontransitive. Well-known mandatory attributes are attributes that must be present in every BGP update message and must be recognized by every BGP speaker. Well-known discretionary attributes are attributes that may or may not be present in a BGP update message but must be recognized by every BGP speaker. Optional transitive attributes are attributes that may or may not be present in a BGP update message and may or may not be recognized by a BGP speaker. If an optional transitive attribute is not recognized by a BGP speaker, it is passed along to the next BGP speaker. Optional nontransitive attributes are attributes that may or may not be present in a BGP update message and may or may not be recognized by a BGP speaker. If an optional nontransitive attribute is not recognized by a BGP speaker, it is not passed along to the next BGP speaker. In this question, we have four routers (R1, R2, R3, and R4) that are connected in a full mesh topology and running IBGP. R3 receives the 192.168.0.0/16 route from its EBGP neighbor and advertises it to R1 and R4 with different BGP attribute values. We are asked which statements are correct about the BGP routes on R3 that are learned from the ISP-A neighbor. Based on the information given, we can infer that the correct statements are:
? By default, the next-hop value for these routes is not changed by ISP-A before being sent to R3. This is because the default behavior of EBGP is to preserve the next-hop attribute of the routes received from another EBGP neighbor. The next- hop attribute indicates the IP address of the router that should be used as the next hop to reach the destination network.
? The BGP local-preference value that is used by ISP-A is not advertised to R3. This is because the local-preference attribute is a well-known discretionary attribute that is used to influence the outbound traffic from an autonomous system. The local- preference attribute is only propagated within an autonomous system and is not advertised to external neighbors.
References: : https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol- bgp/13753-25.html : https://www.cisco.com/c/en/us/support/docs/ip/border-gateway- protocol-bgp/13762-40.html : https://www.cisco.com/c/en/us/support/docs/ip/border- gateway-protocol-bgp/13759-37.html

NEW QUESTION 40
Exhibit

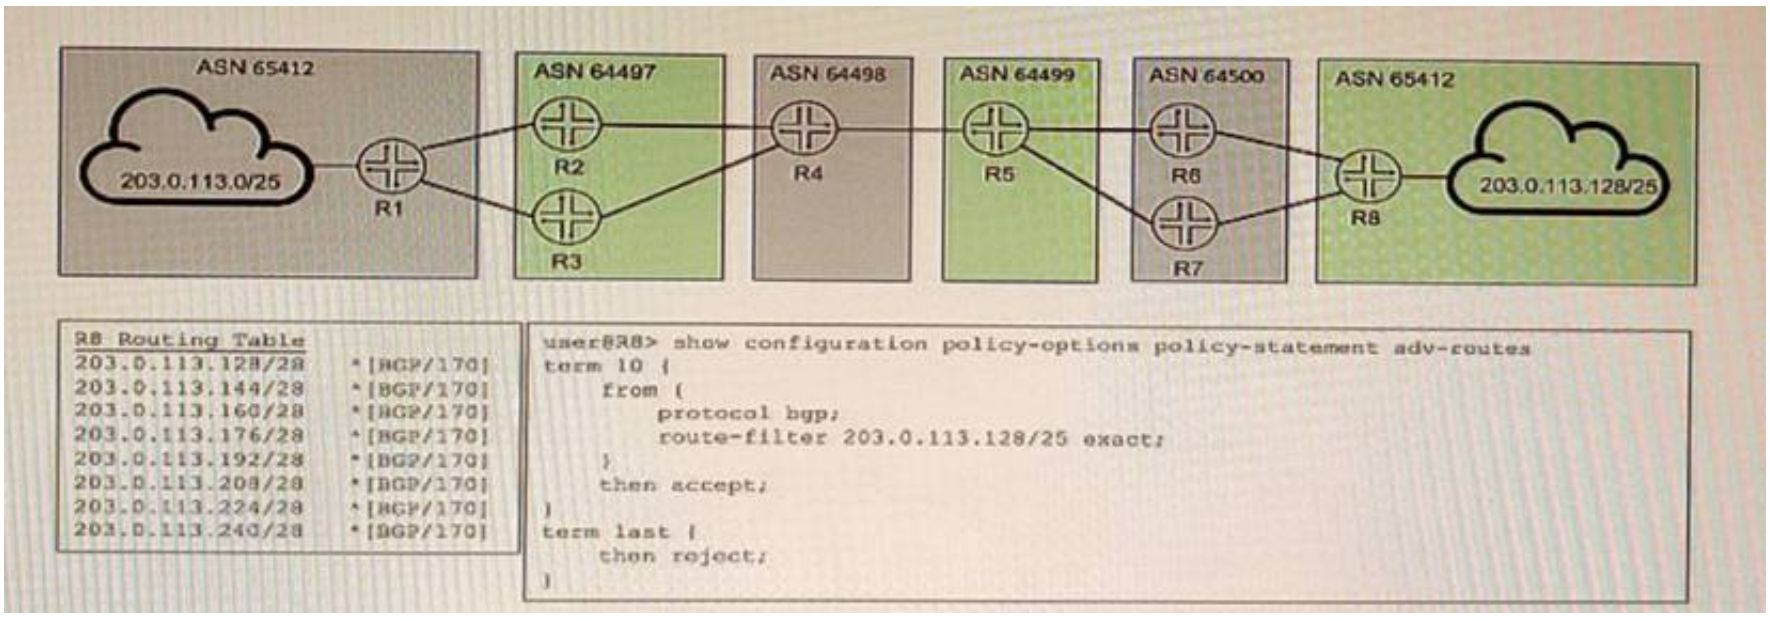

You are attempting to summarize routes from the 203.0.113.128/25 IP block on R8 to AS 64500. You implement the export policy shown in the exhibit and all routes from the routing table stop being advertised.
In this scenario, which two steps would you take to summarize the route in BGP? (Choose two.)

A. Remove the from protocol bgp command from the export policy.
B. Add the set protocols bgp family inet unicast add-path command to allow additional routes to the RIB table

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

Answer: CD

Explanation:
To summarize routes from the 203.0.113.128/25 IP block on R8 to AS 64500, you need to do the following:
? Add the set routing-options static route 203.0.113.128/25 discard command. This creates a static route for the summary prefix and discards any traffic destined to it. This is necessary because BGP can only advertise routes that are present in the routing table.
? Replace exact in the export policy with orlonger. This allows R8 to match and advertise any route that is equal or more specific than the summary prefix. The exact term only matches routes that are exactly equal to the summary prefix, which is not present in the routing table.

NEW QUESTION 42
Exhibit

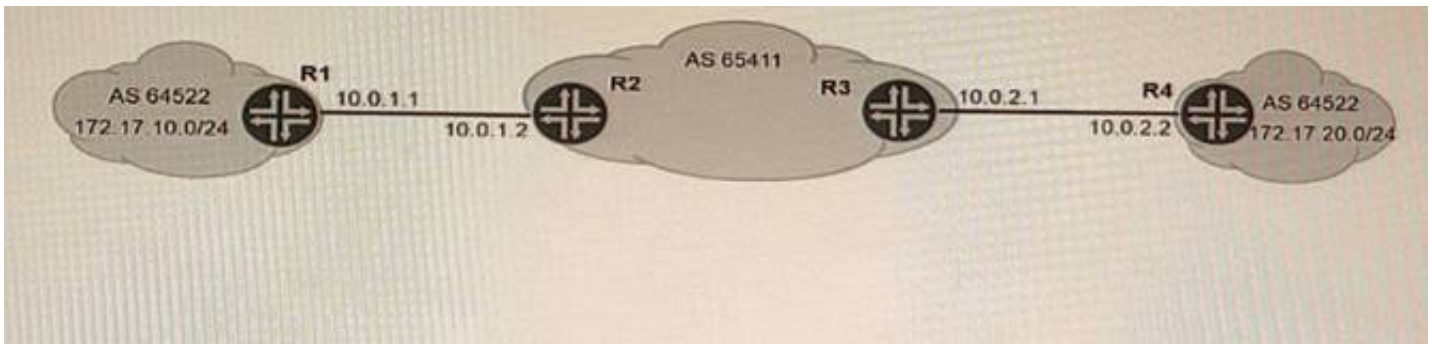

You are asked to exchange routes between R1 and R4 as shown in the exhibit. These two routers use the same AS number Which two steps will accomplish this task? (Choose two.)

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

Answer: AB

Explanation:
The advertise-peer-as parameter allows a router to advertise its peer's AS number as part of the AS path attribute when sending BGP updates to other peers. This parameter is useful when two routers in the same AS need to exchange routes through another AS, such as in the case of R1 and R4. By configuring this parameter on R1 and R4, they can advertise each other's AS number to R2 and R3, respectively.
The as-override parameter allows a router to replace the AS number of its peer with its own AS number when receiving BGP updates from that peer. This parameter is useful when two routers in different ASes need to exchange routes through another AS that has the same AS number as one of them, such as in the case of R2 and R3. By configuring this parameter on R2 and R3, they can override the AS number of R1 and R4 with their own AS number when sending BGP updates to each other.

NEW QUESTION 45
Which two statements are correct about a sham link? (Choose two.)

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

Answer: AC

Explanation:
A sham link is a logical link between two PE routers that belong to the same OSPF area but are connected through an L3VPN. A sham link makes the PE routers appear as if they are directly connected, and prevents OSPF from preferring an intra-area back door link over the VPN backbone. A sham link creates an OSPF multihop 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 48
Which statement is correct about IS-IS when it performs the Dijkstra algorithm?

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

Answer: A

Explanation:
IS-IS is a link-state routing protocol that uses the Dijkstra algorithm to compute the shortest paths between nodes in a network. The Dijkstra algorithm maintains three data structures: a tree database, a candidate database, and a link-state database (LSDB). The tree database contains the nodes that have been visited and their shortest distances from the source node. The candidate database contains the nodes that have not been visited yet and their tentative distances from the source node. The LSDB contains the topology information of the network, such as the links and their costs.
The Dijkstra algorithm works as follows:
? The local router moves its own local tuples into the tree database. A tuple consists of a node ID, a distance, and a parent node ID. The local router's tuple has a distance of zero and no parent node.
? The local router moves its neighbors' tuples into the candidate database. The neighbors' tuples have distances equal to the costs of the links to them and parent node IDs equal to the local router's node ID.

? The local router selects the tuple with the lowest distance from the candidate database and moves it to the tree database. This tuple becomes the current node.
? The local router updates the distances of the current node's neighbors in the candidate database by adding the current node's distance to the link costs. If a shorter distance is found, the parent node ID is also updated.
? The algorithm repeats steps 3 and 4 until either the destination node is reached or the candidate database is empty.

NEW QUESTION 49
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 50
Exhibit

```
user@router> show l2vpn connections
Layer-2 VPN connections:
Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not
CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby        SN -- Static Neighbor
LB -- Local site not best-site   RB -- Remote site not best-site
VM -- VLAN ID mismatch           HS -- Hot-standby Connection
Legend for interface status
Up -- operational
Dn -- down
Instance: vpn-A
Edge protection: Not-Primary
  Local site: CE1-2 (2)
    connection-site Type St     Time last up                    # Up trans
    1               rmt  Up     Apr 11 14:35:27 2020                   1
      Remote PE: 172.17.20.1, Negotiated control-word: Yes (Null)
      Incoming label: 21, Outgoing label: 22
      Local interface: ge-0/0/6.610, Status: Up, Encapsulation: VLAN
      Flow Label Transmit: No, Flow Label Receive: No
```

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

A. The PE is attached to a single local site.
B. The connection has not flapped since it was initiated.
C. There has been a VLAN ID mismatch.
D. The PE router has the capability to pop flow labels

Answer: AD

Explanation:
According to 1 and2, BGP Layer 2 VPNs use BGP to distribute endpoint provisioning information and set up pseudowires between PE devices. BGP uses the Layer 2 VPN (L2VPN) Routing Information Base (RIB) to store endpoint provisioning information, which is updated each time any Layer 2 virtual forwarding instance (VFI) is configured. The prefix and path information is stored in the L2VPN database, which allows BGP to make decisions about the best path.
In the output shown in the exhibit, we can see some information about the L2VPN RIB and the pseudowire state. Based on this information, we can infer the following statements:
? The PE is attached to a single local site. This is correct because the output shows only one local site ID (1) under the L2VPN RIB section. A local site ID is a unique identifier for a site within a VPLS domain. If there were multiple local sites attached to the PE, we would see multiple local site IDs with different prefixes.
? The connection has not flapped since it was initiated. This is correct because the output shows that the uptime of the pseudowire is equal to its total uptime (1w6d). This means that the pseudowire has been up for one week and six days without any interruption or flap.
? There has been a VLAN ID mismatch. This is not correct because the output shows that the remote and local VLAN IDs are both 0 under the pseudowire state section. A VLAN ID mismatch occurs when the remote and local VLAN IDs are different, which can cause traffic loss or misdelivery. If there was a VLAN ID mismatch, we would see different values for the remote and local VLAN IDs.
? The PE router has the capability to pop flow labels. This is correct because the output shows that the flow label pop bit is set under the pseudowire state section. The flow label pop bit indicates that the PE router can pop (remove) the MPLS flow label from the packet before forwarding it to the CE device. The flow label is an

optional MPLS label that can be used for load balancing or traffic engineering purposes.

NEW QUESTION 51
Exhibit

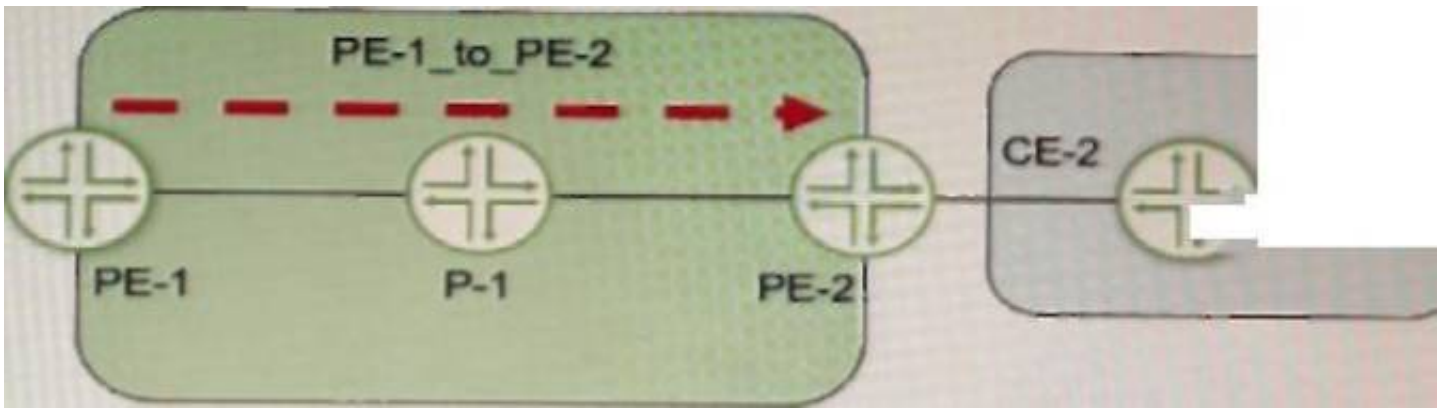

Referring to the exhibit, a working L3VPN exists that connects VPN-A sites CoS is configured correctly to match on the MPLS EXP bits of the LSP, but when traffic is sent from Site-1 to Site-2, PE-2 is not classifying the traffic correctly
What should you do to solve the problem?

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 55
Exhibit

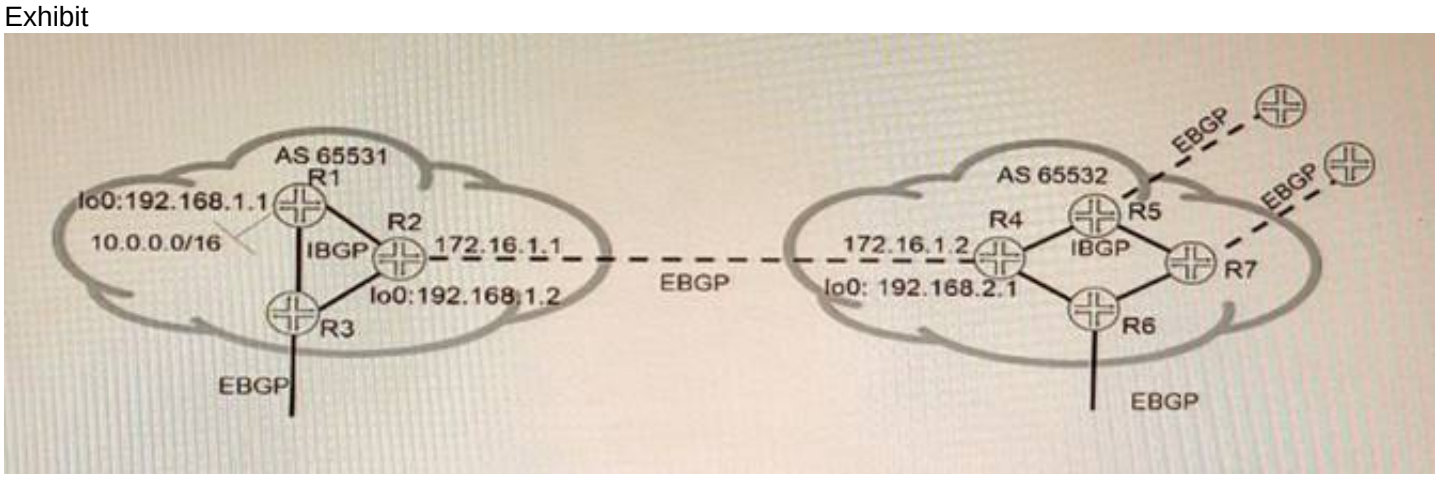

Referring to the exhibit, which three statements are correct about route 10 0 0.0/16 when using the default BGP advertisement rules'? (Choose three.)

A. R1 will prepend AS 65531 when advertising 10 0.0 0/16 to R2.
B. R1 will advertise 10.0.0.0/16 to R2 with 192 168 1 1 as the next hop.
C. R2 will advertise 10.0.0.0/16 to R3 with 192.168.1 1 as the next hop
D. R4 will advertise 10 0.0 0/16 to R6 with 172.16 1 1 as the next hop
E. R2 will advertise 10.0.0.0/16 to R4 with 172.16.1.1 as the next hop

Answer: BDE

Explanation:
The problem in this scenario is that R1 and R8 are not receiving each other's routes because of private AS numbers in the AS path. Private AS numbers are not globally unique and are not advertised to external BGP peers. To solve this problem, you need to do the following:
? Configure loops on routers in AS 65412 and advertise-peer-as on routers in AS 64498. This allows R5 and R6 to advertise their own AS number (65412) instead of their peer's AS number (64498) when sending updates to R7 and R8. This prevents a loop detection issue that would cause R7 and R8 to reject the routes from R5 and R62.
? Configure remove-private on advertisements from AS 64497 toward AS 64498 and from AS 64500 toward AS 64499. This removes any private AS numbers from the AS path before sending updates to external BGP peers. This allows R2 and R3 to receive the routes from R1 and R4, respectively3.

NEW QUESTION 59
Exhibit

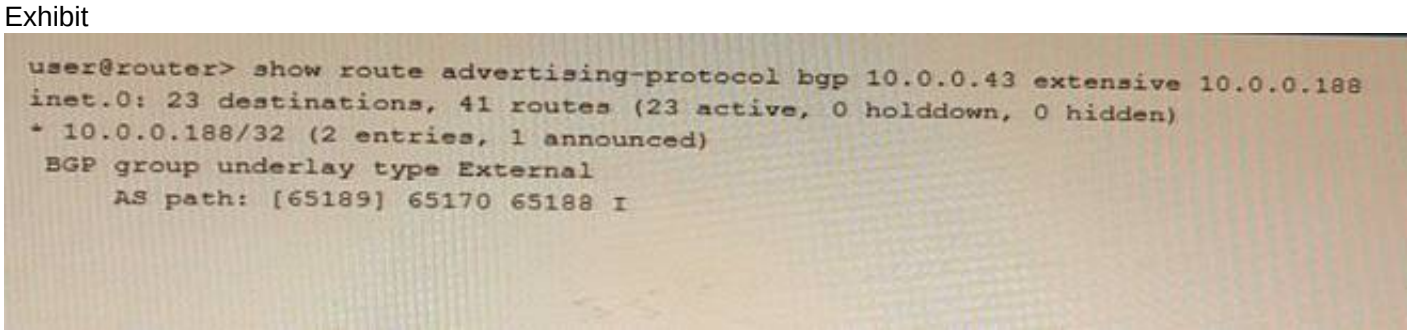

Referring to the exhibit, what do the brackets [ ] in the AS path identify?

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

Answer: B

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

NEW QUESTION 60
.....

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