# Juniper

## Exam Questions JN0-664

Service Provider - Professional (JNCIP-SP)

**NEW QUESTION 1**
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
                       Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                       Indirect path forwarding next hops: 2
                               Next hop type: Router
                               Next hop: 10.0.2.12 via et-0/0/0.0
                               Session Id: 0x0
                               Next hop: 10.0.2.22 via et-0/0/1.0
                               Session Id: 0x0
192.168.101.5/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 2
Nexthop: 10.0.2.12 via et-0/0/0.0
Session Id: 0
Nexthop: 10.0.2.22 via et-0/0/1.0
Session Id: 0
...
```

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.

**Answer:** BC

**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.

**NEW QUESTION 2**
Exhibit

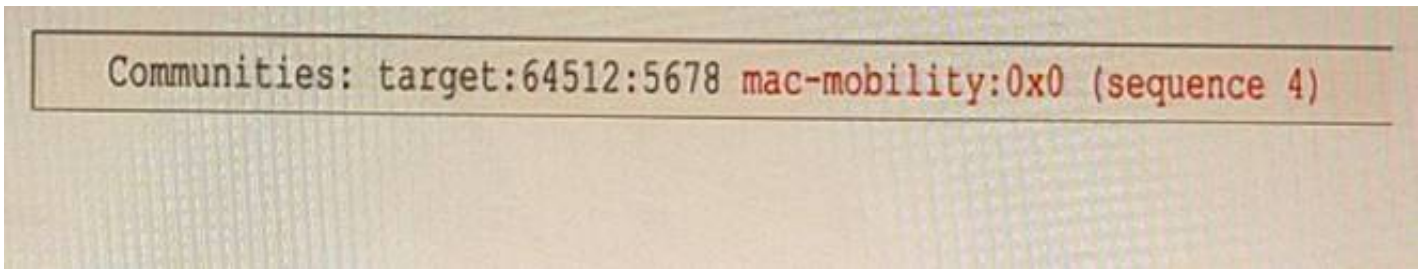


You have MAC addresses moving in your EVPN environment
Referring to the exhibit, which two statements are correct about the sequence number? (Choose two)

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

**Answer:** BC

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

**NEW QUESTION 3**
Exhibit

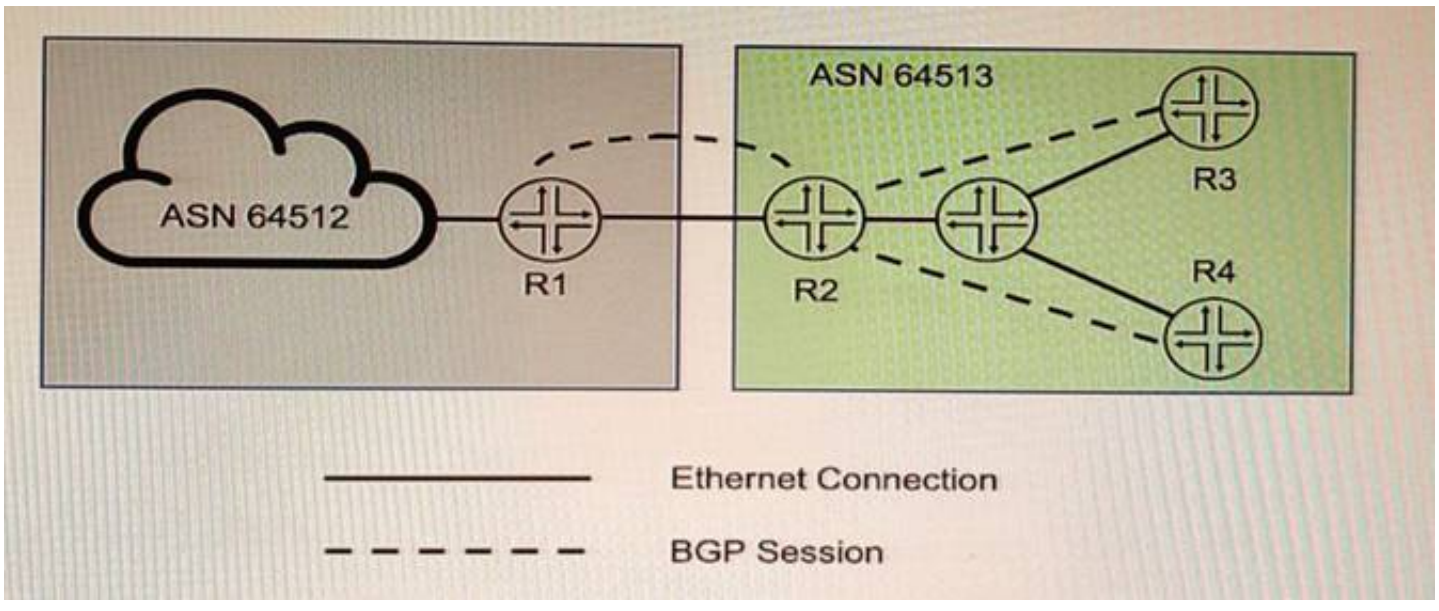


You want to implement the BGP Generalized TTL Security Mechanism (GTSM) on the network
Which three statements are correct in this scenario? (Choose three)

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 4**
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 5**
Exhibit

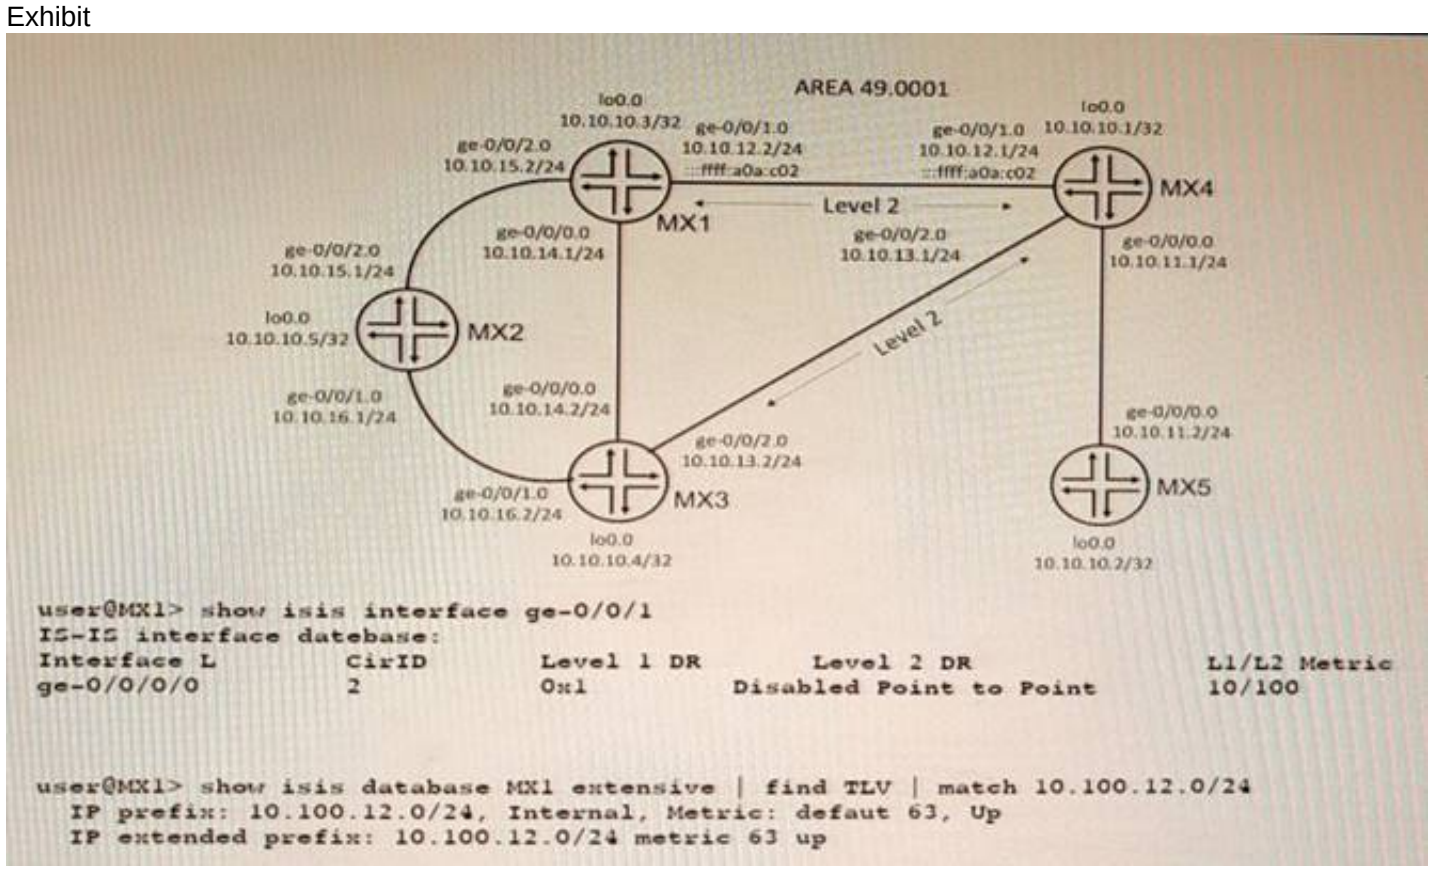


A network is using IS-IS for routing.
In this scenario, why are there two TLVs shown in the exhibit?

A. There are both narrow and wide metric devices in the topology
B. The interface specified a metric of 100 for L2.
C. Wide metrics have specifically been requested
D. Both IPv4 and IPv6 are being used in the topology

**Answer:** A

**Explanation:**
TLVs are tuples of (Type, Length, Value) that can be advertised in IS-IS packets. TLVs can carry different kinds of information in the Link State Packets (LSPs). IS- IS supports both narrow and wide metrics for link costs. Narrow metrics use a single octet to encode the link cost, while wide metrics use three octets. Narrow metrics have a maximum value of 63, while wide metrics have a maximum value of 16777215. If there are both narrow and wide metric devices in the topology, IS-IS will advertise two TLVs for each link: one with the narrow metric and one with the wide metric. This allows backward compatibility with older devices that only support narrow metrics12.

**NEW QUESTION 6**
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 7**
Exhibit

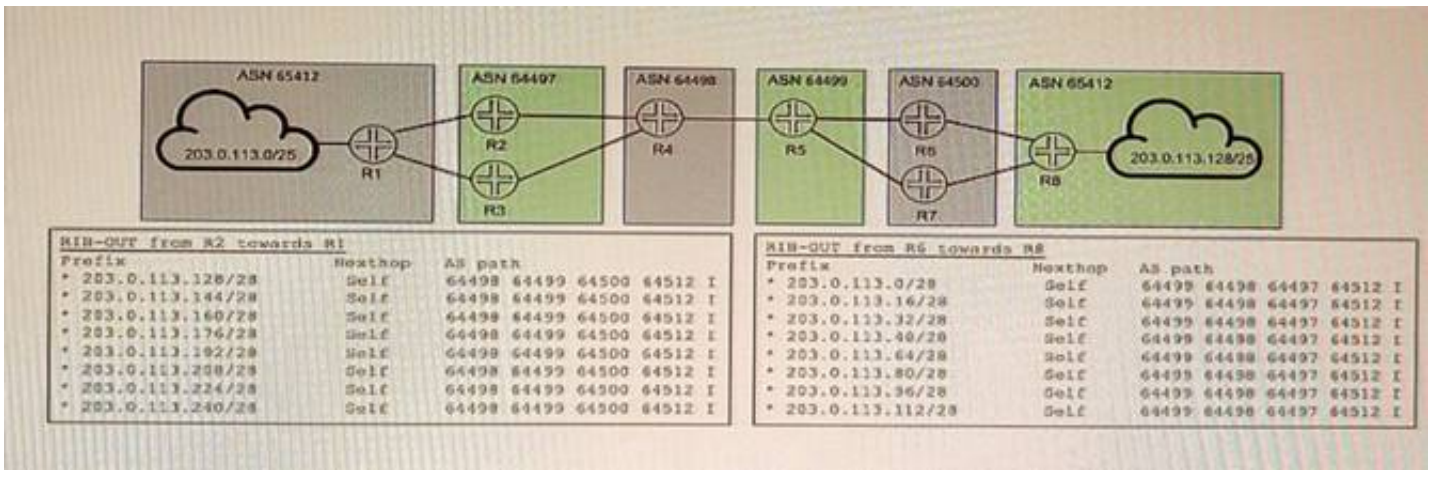


R1 and R8 are not receiving each other's routes
Referring to the exhibit, what are three configuration commands that would solve this problem? (Choose three.)

A. Configure loops and advertise-peer-as on routers in AS 64497 and AS 64450.
B. Configure loops on routers in AS 65412 and advertise-peer-as on routers in AS 64498.

C. Configure as-override on advertisement from AS 64500 toward AS 64512.
D. Configure remove-private on advertisements from AS 64497 toward AS 64498
E. Configure remove-private on advertisements from AS 64500 toward AS 64499

**Answer:** BDE

**Explanation:**
The problem in this scenario is that R1 and R8 are not receiving each other's routes because of private AS numbers in the AS path. Private AS numbers are not globally unique and are not advertised to external BGP peers. To solve this problem, you need to do the following:
? Configure loops on routers in AS 65412 and advertise-peer-as on routers in AS 64498. This allows R5 and R6 to advertise their own AS number (65412) instead of their peer's AS number (64498) when sending updates to R7 and R8. This prevents a loop detection issue that would cause R7 and R8 to reject the routes from R5 and 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 8**
What is the correct order of packet flow through configurable components in the Junos OS CoS features?

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

**Answer:** C

**Explanation:**
The correct order of packet flow through configurable components in the Junos OS CoS features is as follows:
? Behavior Aggregate Classifier: This component uses a single field in a packet header to classify traffic into different forwarding classes and loss priorities based on predefined or user-defined values.
? Input Policer: This component applies rate-limiting and marking actions to incoming traffic based on the forwarding class and loss priority assigned by the classifier.
? Multifield Classifier: This component uses multiple fields in a packet header to classify traffic into different forwarding classes and loss priorities based on user-defined values and filters.
? Forwarding Policy Options: This component applies actions such as load balancing, filtering, or routing to traffic based on the forwarding class and loss priority assigned by the classifier.
? Fabric Scheduler: This component schedules traffic across the switch fabric based on the forwarding class and loss priority assigned by the classifier.
? Output Policer: This component applies rate-limiting and marking actions to outgoing traffic based on the forwarding class and loss priority assigned by the classifier.
? Scheduler/Shaper/RED: This component schedules, shapes, and drops traffic at the egress interface based on the forwarding class and loss priority assigned by the classifier.
? Rewrite Marker: This component rewrites the code-point bits of packets leaving an interface based on the forwarding class and loss priority assigned by the classifier.

**NEW QUESTION 9**
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 10**
Exhibit

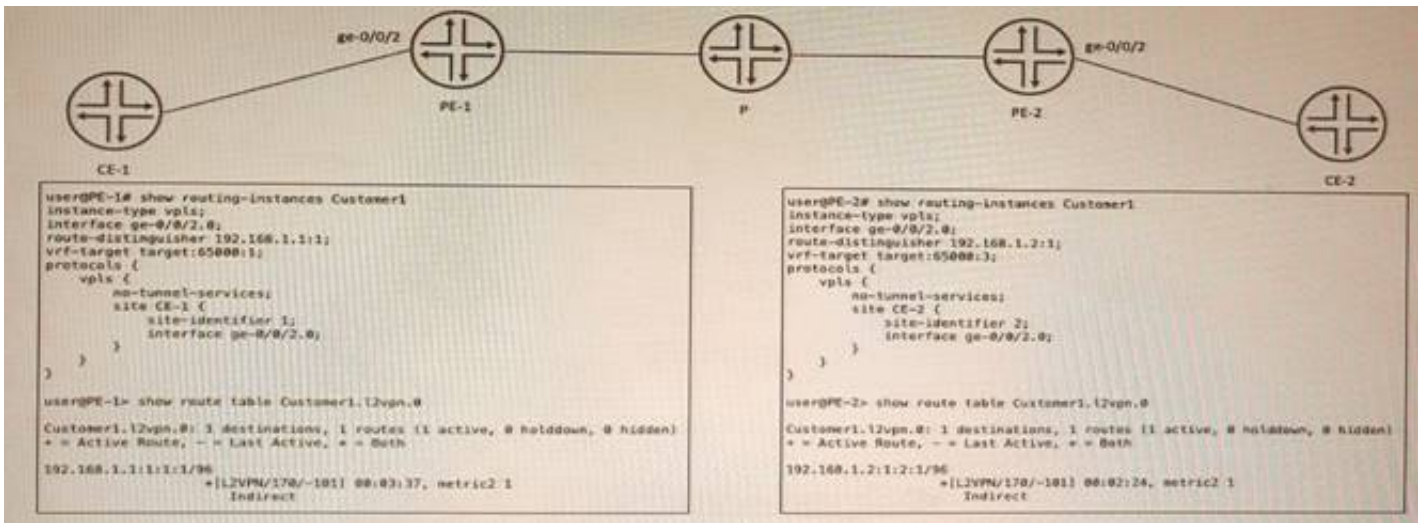


CE-1 and CE-2 are part of a VPLS called Customer1 No connectivity exists between CE-1 and CE-2. In the process of troubleshooting, you notice PE-1 is not learning any routes for this VPLS from PE-2, and PE-2 is not learning any routes for this VPLS from PE-1.

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.

**Answer:** A

**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.

**NEW QUESTION 10**
Exhibit

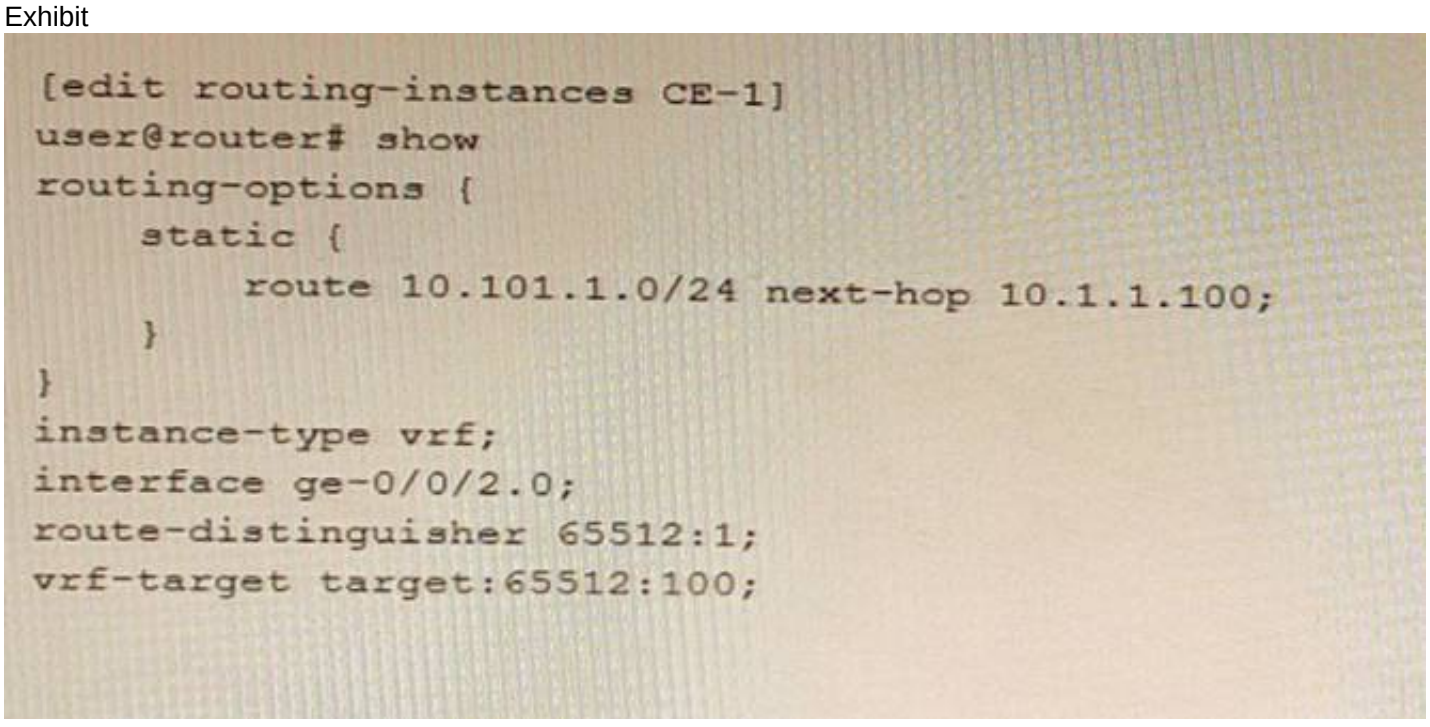


Referring to the exhibit, which statement is true?

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

**Answer:** D

**Explanation:**
The auto-export parameter is a routing option that allows a routing instance to share routes with other routing instances or the master routing table. The auto-export parameter automatically exports routes from one routing instance to another based on the route target communities attached to the routes. In this scenario, the 10.101.1.0/24 route will be shared if the auto-export parameter is configured under [edit routing-options] hierarchy level.

**NEW QUESTION 13**
By default, which statement is correct about OSPF summary LSAs?

A. All Type 2 and Type 7 LSAs will be summanzed into a single Type 5 LSA
B. The area-range command must be installed on all routers.
C. Type 3 LSAs are advertised for routes in Type 1 LSAs.
D. The metric associated with a summary route will be equal to the lowest metric associated with an individual contributing route

**Answer:** C

**Explanation:**
OSPF uses different types of LSAs to describe different aspects of the network topology. Type 1 LSAs are also known as router LSAs, and they describe the links and interfaces of a router within an area. Type 3 LSAs are also known as summary LSAs, and they describe routes to networks outside an area but within the same autonomous system (AS). By default, OSPF will summarize routes from Type 1 LSAs into Type 3 LSAs when advertising them across area boundaries .

**NEW QUESTION 18**
Exhibit

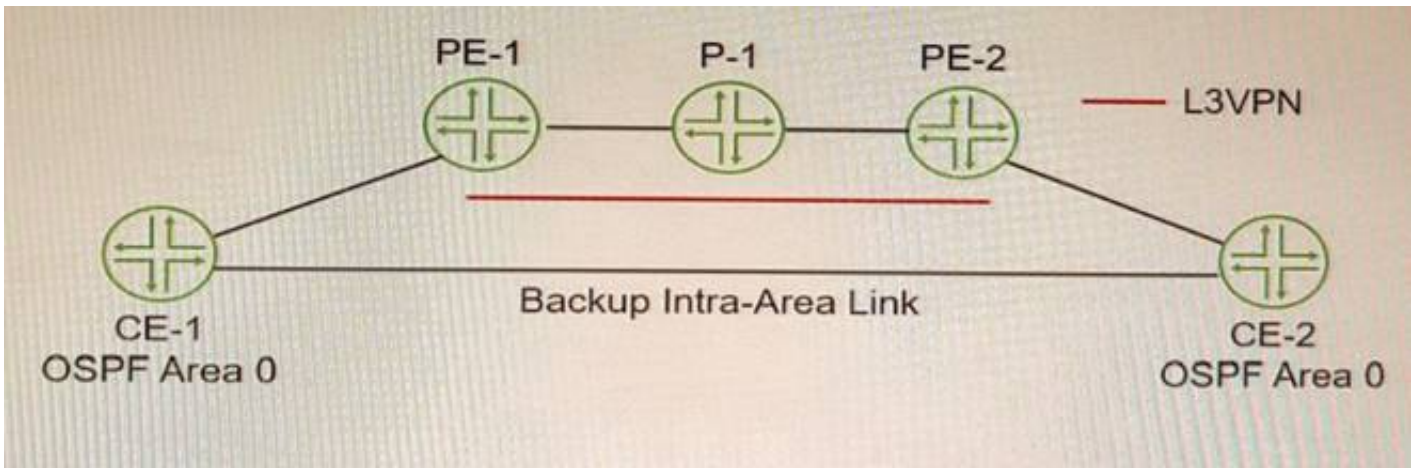


You must ensure that the VPN backbone is preferred over the back door intra-area link as long as the VPN is available. Referring to the exhibit, which action will accomplish this task?

A. Configure an import routing policy on the CE routers that rejects OSPF routes learned on the backup intra-area link.
B. Enable OSPF traffic-engineering.
C. Configure the OSPF metric on the backup intra-area link that is higher than the L3VPN link.
D. Create an OSPF sham link between the PE routers.

**Answer:** D

**Explanation:**
A sham link is a logical link between two PE routers that belong to the same OSPF area but are connected through an L3VPN. A sham link makes the PE routers appear as if they are directly connected, and prevents OSPF from preferring an intra-area back door link over the VPN backbone. To create a sham link, you need to configure the local and remote addresses of the PE routers under the [edit protocols ospf area area-id] hierarchy level1.

**NEW QUESTION 23**
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 25**
Exhibit

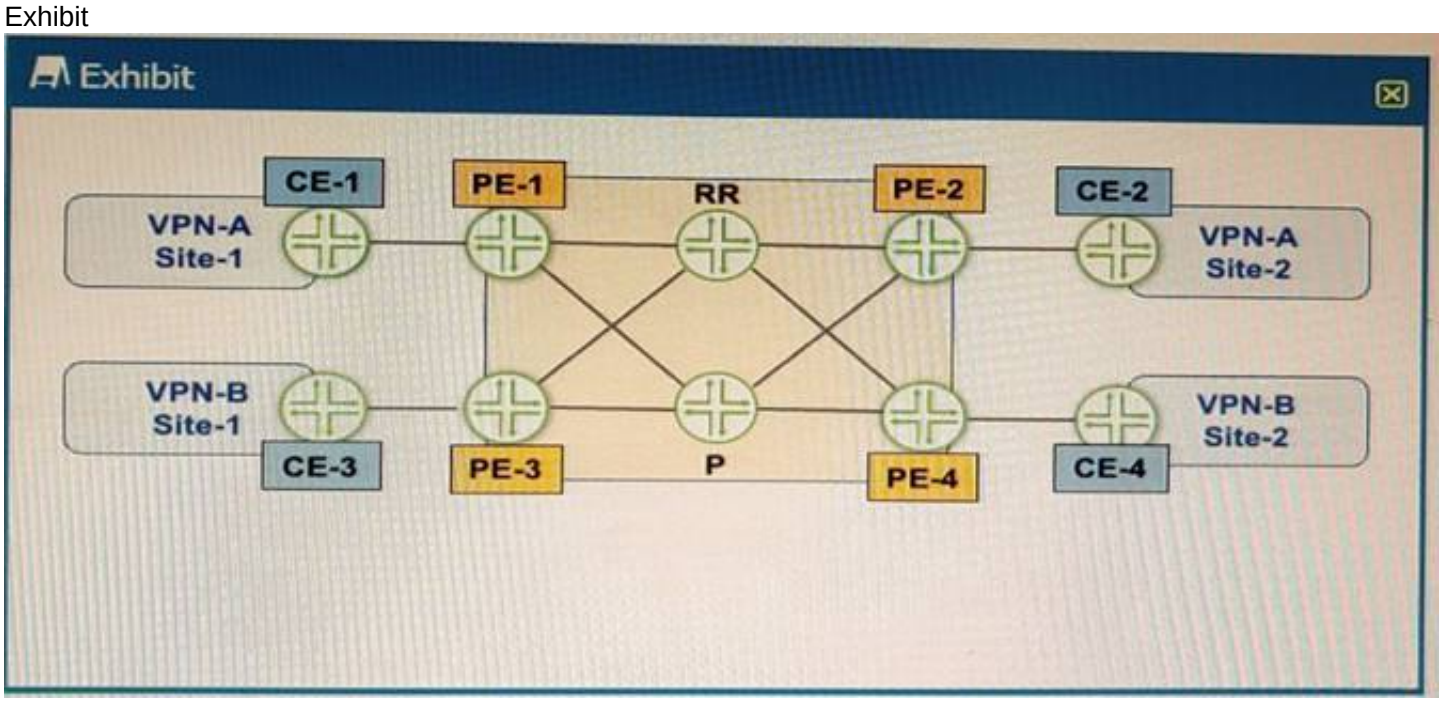


Referring to the exhibit, PE-1 and PE-2 are getting route updates for VPN-B when neither of them service that VPN Which two actions would optimize this process? (Choose two.)

A. Configure the family route-target statement on the PEs.
B. Configure the family route-target statement on the RR
C. Configure the resolution rib bgp . 13vpn . 0 resolution-ribs ine
D. 0 Statement on the PEs.
E. Configure the resolution rib bgp.l3vpn.O resolution-ribs ine
F. 0 Statement on the RR

**Answer:** BD

**Explanation:**
BGP route target filtering is a technique that reduces the number of routers that receive VPN routes and route updates, helping to limit the amount of overhead associated with running a VPN. BGP route target filtering is based on the exchange of the route-target address family, which contains information about the VPN membership of each PE device. Based on this information, a PE device can decide whether to accept or reject VPN routes from another PE device.
BGP route target filtering can be configured on PE devices or on route reflectors (RRs). Configuring BGP route target filtering on RRs is more efficient and scalable, as it reduces the number of BGP sessions and updates between PE devices. To configure BGP route target filtering on RRs, the following steps are required:
? Configure the family route-target statement under the BGP group or neighbor configuration on the RRs. This enables the exchange of the route-target address family between the RRs and their clients (PE devices).
? Configure the resolution rib bgp.l3vpn.0 resolution-ribs inet.0 statement under the routing-options configuration on the RRs. This enables the RRs to resolve next hops for VPN routes using the inet.0 routing table.
? Configure an export policy for BGP route target filtering under the routing-options configuration on the RRs. This policy controls which route targets are advertised to each PE device based on their VPN membership.

**NEW QUESTION 26**
Exhibit

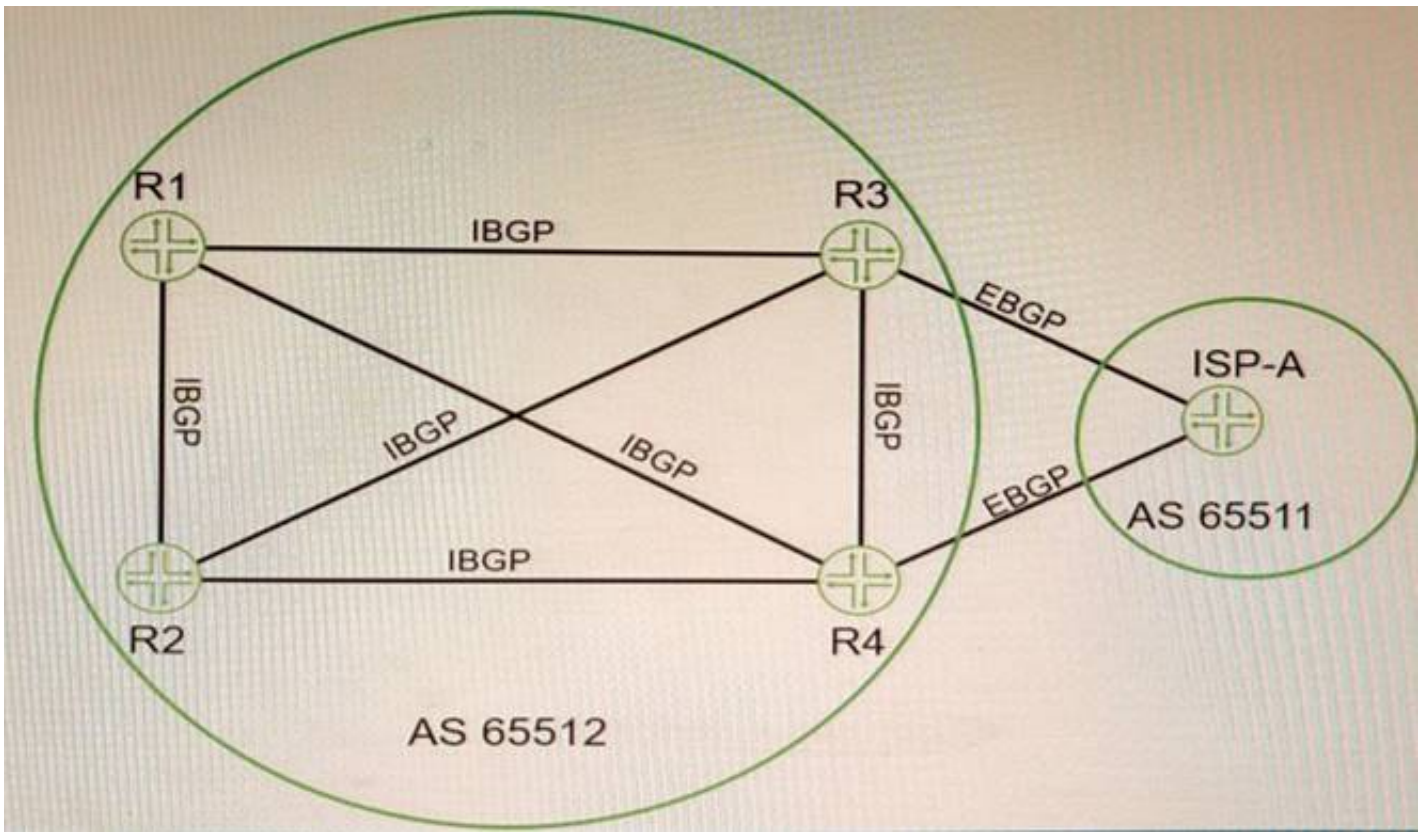


Click the Exhibit button-Referring to the exhibit, which two statements are correct about BGP routes on R3 that are learned from the ISP-A neighbor? (Choose two.)

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 30**
Exhibit

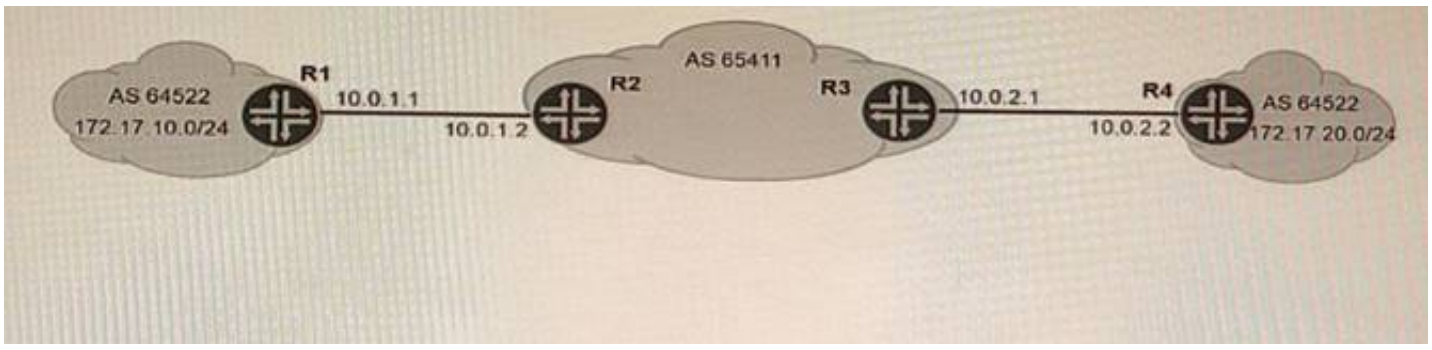


You are asked to exchange routes between R1 and R4 as shown in the exhibit. These two routers use the same AS number Which two steps will accomplish this task? (Choose two.)

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 33**
You are a network architect for a service provider and want to offer Layer 2 services to your customers You want to use EVPN for Layer 2 services in your existing MPLS network.
Which two statements are correct in this scenario? (Choose two.)

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

**Answer:** CD

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

**NEW QUESTION 36**
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 39**
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 40**
Exhibit

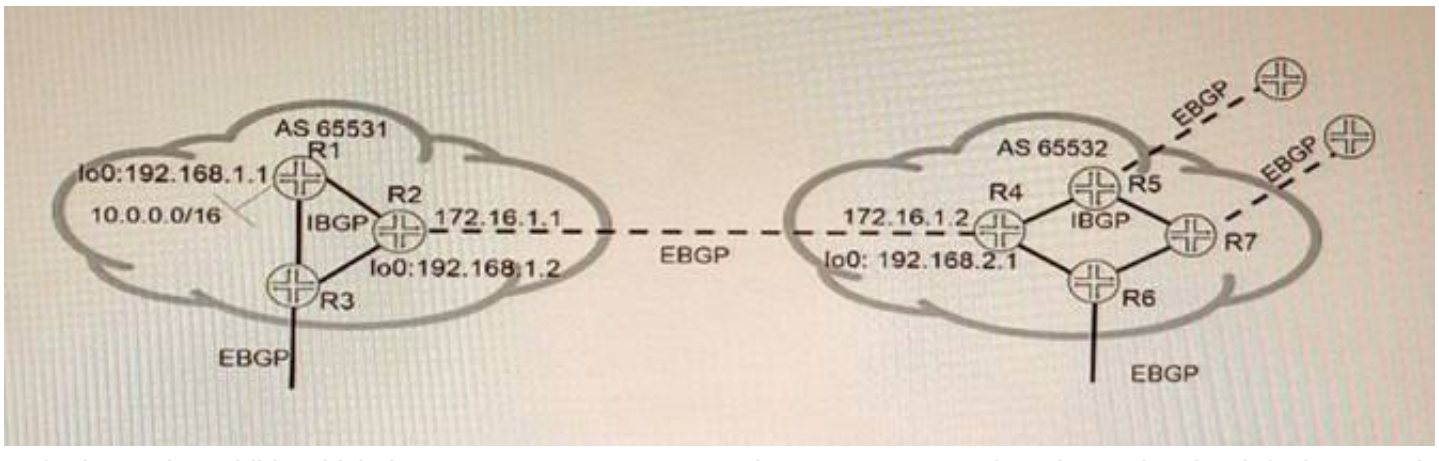


Referring to the exhibit, which three statements are correct about route 10 0 0.0/16 when using the default BGP advertisement rules'? (Choose three.)

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 42**
Exhibit

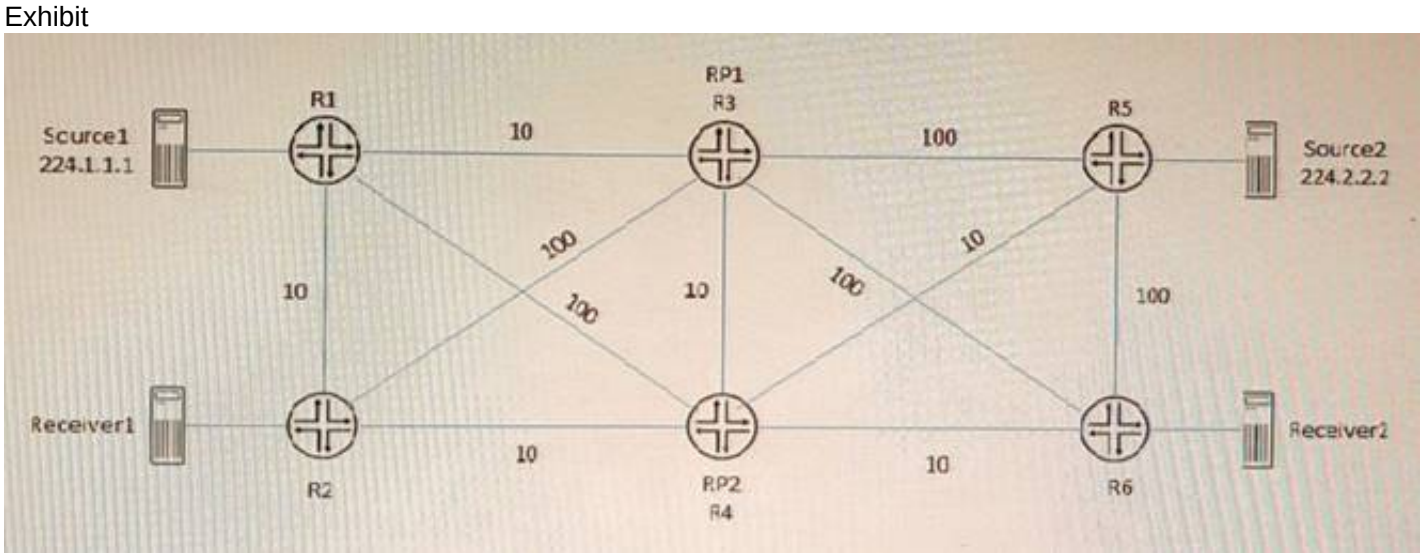


Referring to the exhibit, PIM-SM is configured on all routers, and Anycast-RP with Anycast- PIM is used for the discovery mechanism on RP1 and RP2. The interface metric values are shown for the OSPF area.
In this scenario, which two statements are correct about which RP is used? (Choose two.)

A. Source2 will use RP2 and Received will use RP2 for group 224.2.2.2.
B. Source2 will use RP1 and Receiver2 will use RP1 for group 224.2.2.2.
C. Source1 will use RP1 and Receiver1 will use RP1 for group 224.1.1.1.
D. Source1 will use RP1 and Receiver1 will use RP2 for group 224.1 1 1

**Answer:** AC

**Explanation:**
A sham link is a logical link between two PE routers that belong to the same OSPF area but are connected through an L3VPN. A sham link makes the PE routers appear as if they are directly connected, and prevents OSPF from preferring an intra-area back door link over the VPN backbone. A sham link creates an OSPF multihop neighborship between the PE routers using TCP port 646. The PEs exchange Type 1 OSPF LSAs instead of Type 3 OSPF LSAs for the L3VPN routes, which allows OSPF to use the correct metric for route selection1.

**NEW QUESTION 43**
.....

# 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

100% Actual & Verified — Instant Download, Please Click
Order The JN0-664 Practice Test Here