# Juniper.JN0-664.v2024-07-12.q48

| | |
|---|---|
| **Exam Code:** | JN0-664 |
| **Exam Name:** | Service Provider, Professional (JNCIP-SP) |
| **Certification Provider:** | Juniper |
| **Free Question Number:** | 48 |
| **Version:** | v2024-07-12 |
| **# of views:** | 547 |
| **# of Questions views:** | 480 |
| https://www.freeqas.com/qa/Juniper/JN0-664/Juniper.JN0-664.v2024-07-12.q48.html | |

**NEW QUESTION: 1**
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: C,D (LEAVE A REPLY)**
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: 2**
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 (LEAVE A REPLY)**
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: 3**
Exhibit

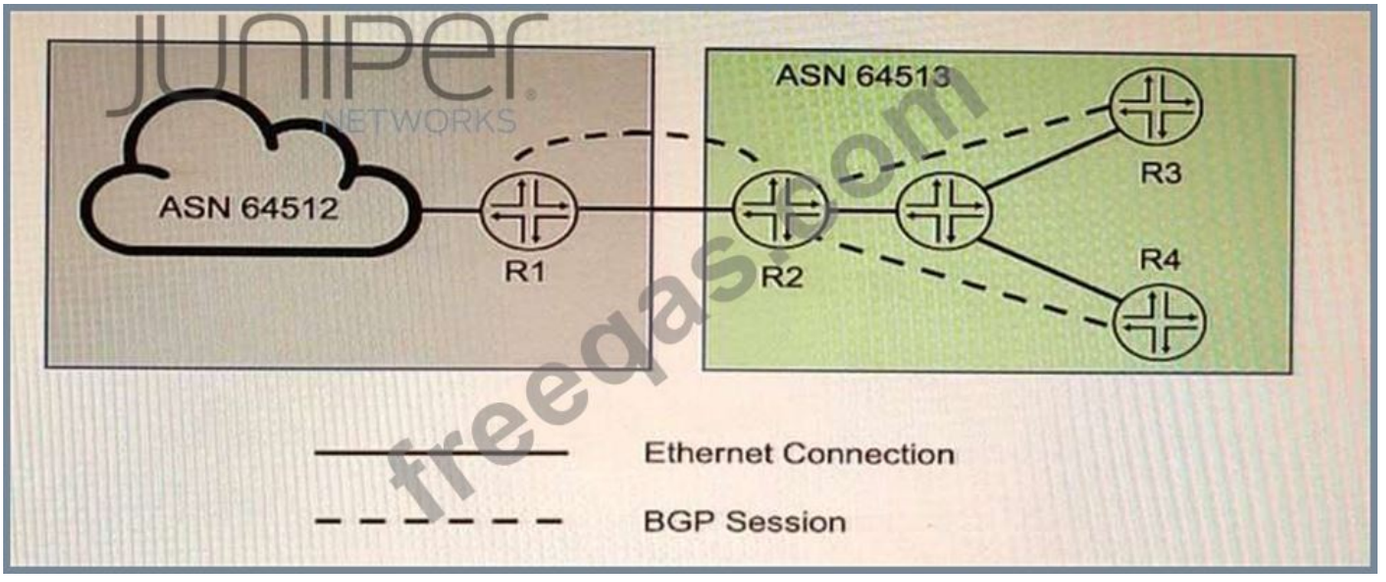


You want to implement the BGP Generalized TTL Security Mechanism (GTSM) on the network
Which three statements are correct in this scenario? (Choose three)
**A.** You can implement BGP GTSM between R2, R3, and R4
**B.** BGP GTSM requires a firewall filter to discard packets with incorrect TTL.
**C.** 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: (SHOW ANSWER)**
https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/ref/statement/multihop-edit-protocols-bg

**NEW QUESTION: 4**
You are responding to an RFP for a new MPLS VPN implementation. The solution must use LDP for signaling and support Layer 2 connectivity without using BGP The solution must be scalable and support multiple VPN connections over a single MPLS LSP The customer wants to maintain all routing for their Private network In this scenario, which solution do you propose?
**A.** circuit cross-connect
**B.** BGP Layer 2 VPN
**C.** LDP Layer 2 circuit
**D.** translational cross-connect
**Answer: (SHOW ANSWER)**
Explanation
AToM (Any Transport over MPLS) is a framework that supports various Layer 2 transport types over an MPLS network core. One of the transport types supported by AToM is LDP Layer 2 circuit, which is a point-to-point Layer 2 connection that uses LDP for signaling and MPLS for forwarding. LDP Layer 2 circuit can support Layer 2 connectivity without using BGP and can be scalable and efficient by using a single MPLS LSP for multiple VPN connections. The customer can maintain all routing for their private network by using their own CE switches.

**NEW QUESTION: 5**
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: B,D,E (LEAVE A REPLY)**
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: 6**
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: C,D (LEAVE A REPLY)**
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: 7**
Exhibit

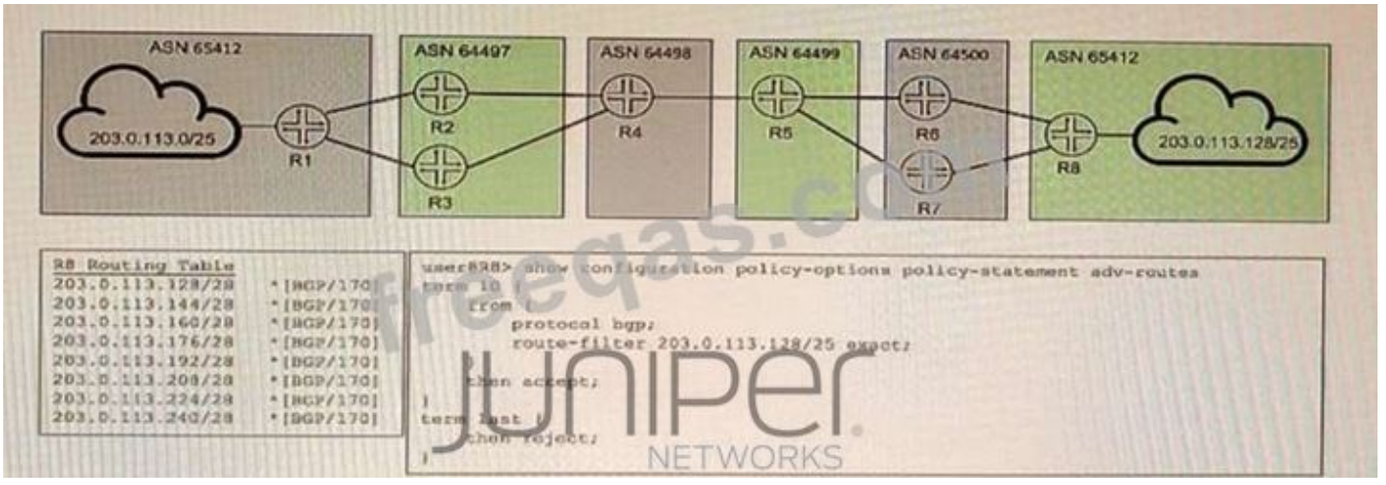

You are attempting to summarize routes from the 203.0.113.128/25 IP block on R8 to AS 64500. You implement the export policy shown in the exhibit and all routes from the routing table stop being advertised.
In this scenario, which two steps would you take to summarize the route in BGP? (Choose two.)
**A.** Remove the from protocol bgp command from the export policy.
**B.** Add the set protocols bgp family inet unicast add-path command to allow additional routes to the RIB tables. -
**C.** Add the set routing-options static route 203.0.113.123/25 discard command.

**D.** Replace exact in the export policy with orlonger.

**Answer: C,D (LEAVE A REPLY)**

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: 8**

Exhibit

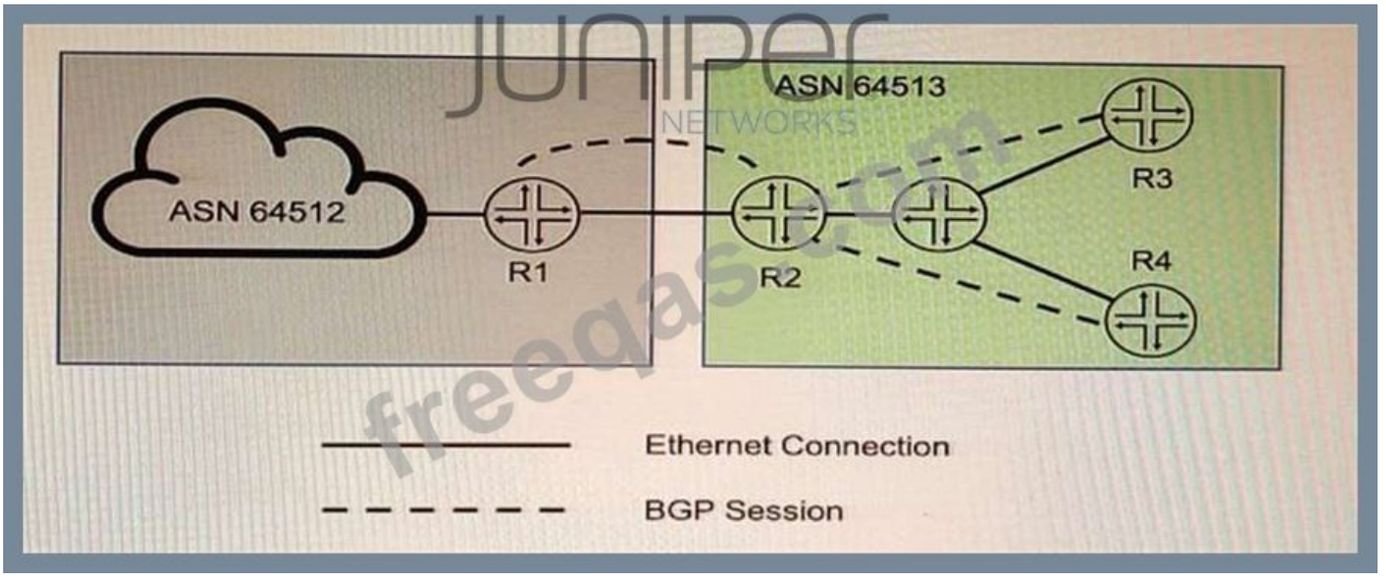


You want to implement the BGP Generalized TTL Security Mechanism (GTSM) on the network

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: A,D,E (LEAVE A REPLY)**

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: 9**
Exhibit

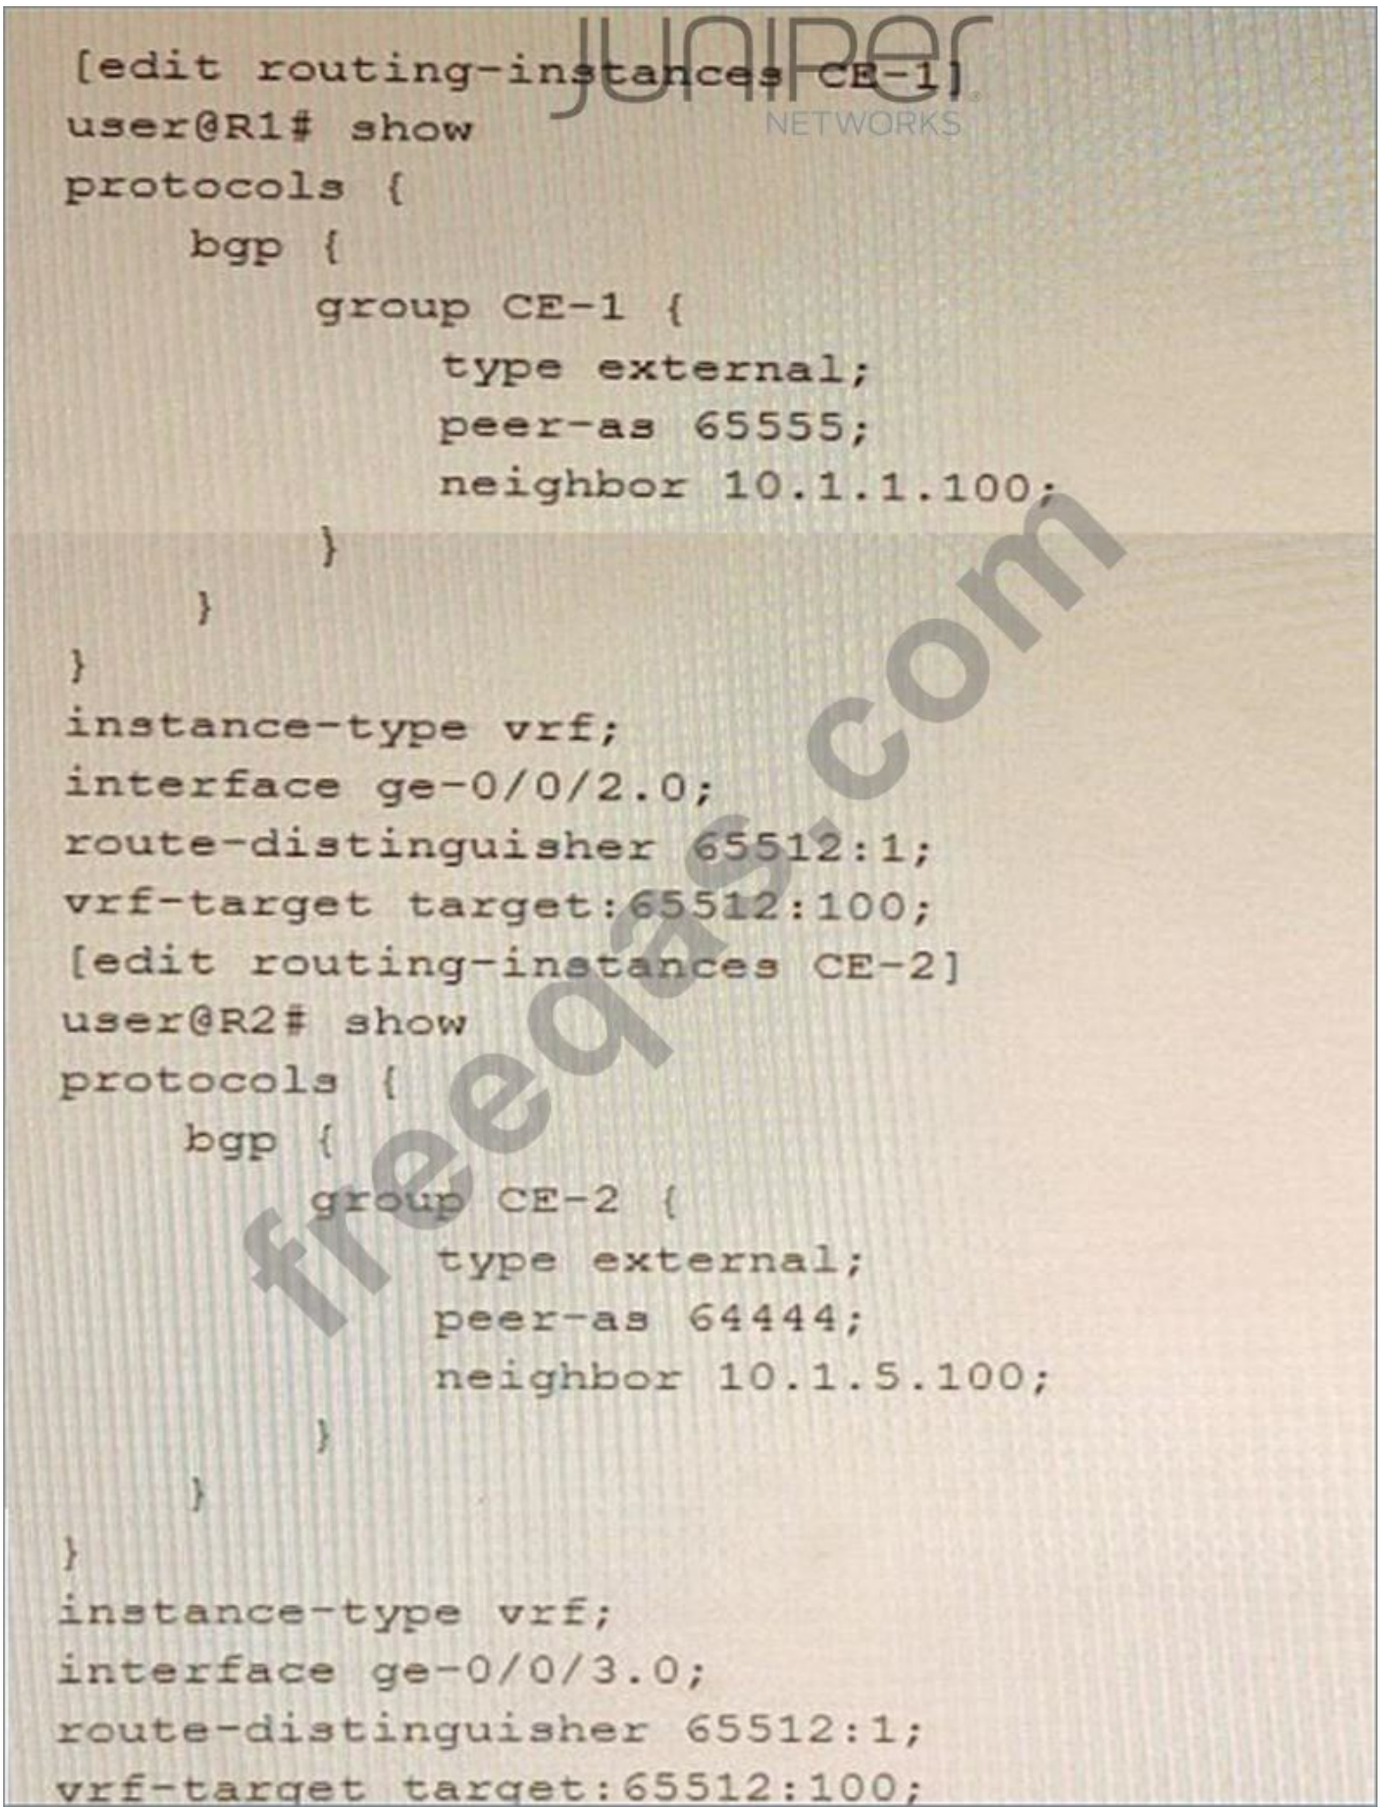

```
[edit routing-instances CE-1]
user@R1# show
protocols {
    bgp {
        group CE-1 {
            type external;
            peer-as 65555;
            neighbor 10.1.1.100;
        }
    }
}
instance-type vrf;
interface ge-0/0/2.0;
route-distinguisher 65512:1;
vrf-target target:65512:100;
[edit 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 (LEAVE A REPLY)**
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: 10**
Exhibit

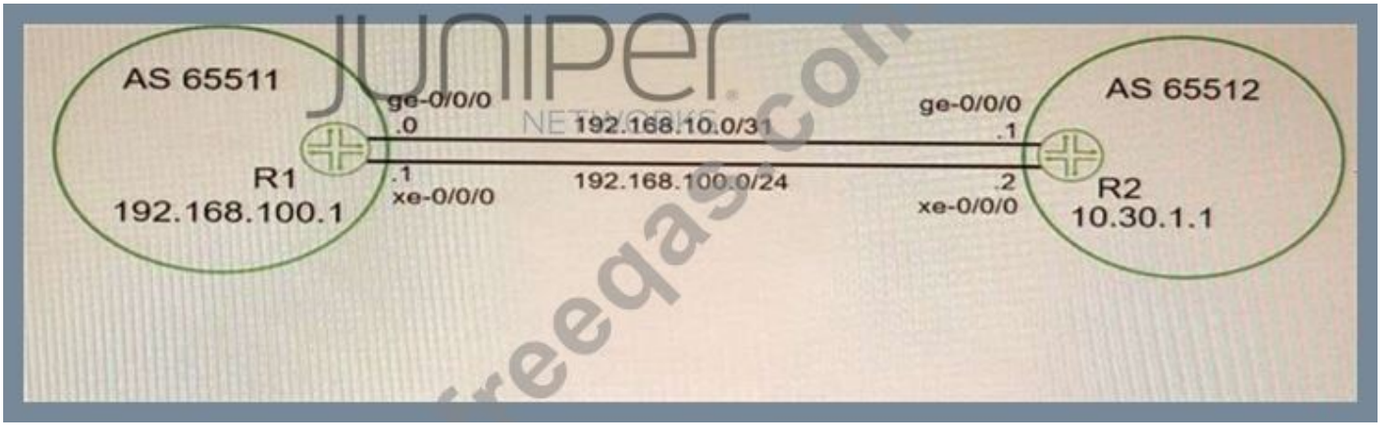


You want to use both links between R1 and R2 Because of the bandwidth difference between the two links, you must ensure that the links are used as much as possible.
Which action will accomplish this goal?
**A.** Define a policy to tag routes with the appropriate bandwidth community.
**B.** Disable multipath.
**C.** Ensure that the metric-out parameter on the Gigabit Ethernet interface is higher than the 10 Gigibit Ethernet interface.
**D.** Enable per-prefix load balancing.
**Answer: (SHOW ANSWER)**
Explanation
VPLS is a Layer 2 VPN technology that allows multiple sites to connect over a shared IP/MPLS network as if they were on the same LAN. VPLS tunnels can be signaled using either Label Distribution Protocol (LDP) or Border Gateway Protocol (BGP). In this question, we have two links between R1 and R2 with different bandwidths (10 Gbps and 1 Gbps). We want to use both links

as much as possible for VPLS traffic. To achieve this, we need to enable per-prefix load balancing on both routers. Per-prefix load balancing is a feature that allows a router to distribute traffic across multiple equal-cost or unequal-cost paths based on the destination prefix of each packet. This improves the utilization of multiple links and provides better load sharing than per-flow load balancing, which distributes traffic based on a hash of source and destination addresses4. Per-prefix load balancing can be enabled globally or per interface using the load-balance per-packet command.

**NEW QUESTION: 11**
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 (LEAVE A REPLY)**
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: 12**

Exhibit

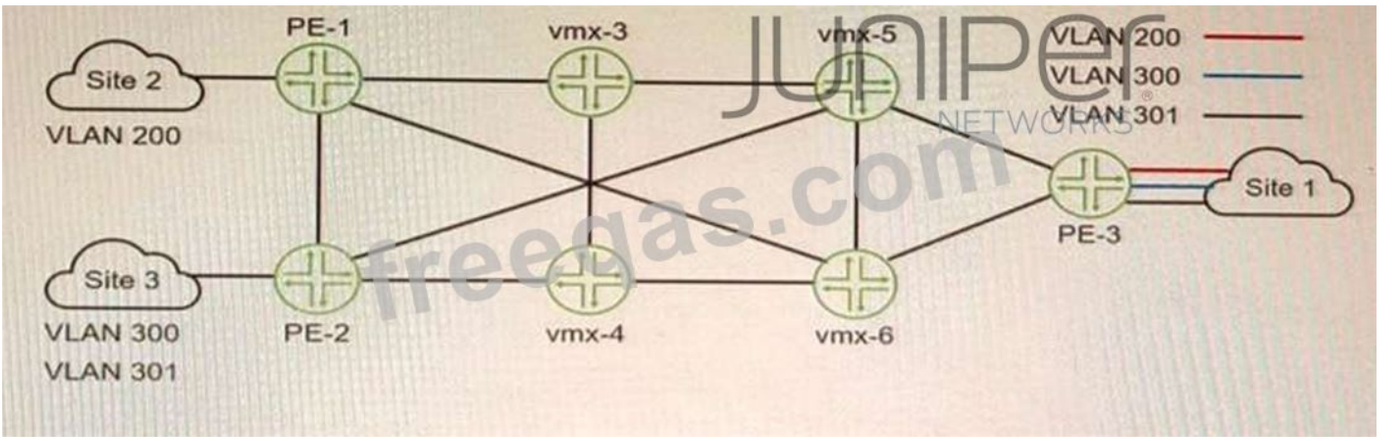


You want Site 1 to access three VLANs that are located in Site 2 and Site 3 The customer-facing interface on the PE-1 router is configured for Ethernet-VLAN encapsulation.

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 (LEAVE A REPLY)**

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: 13**

Exhibit

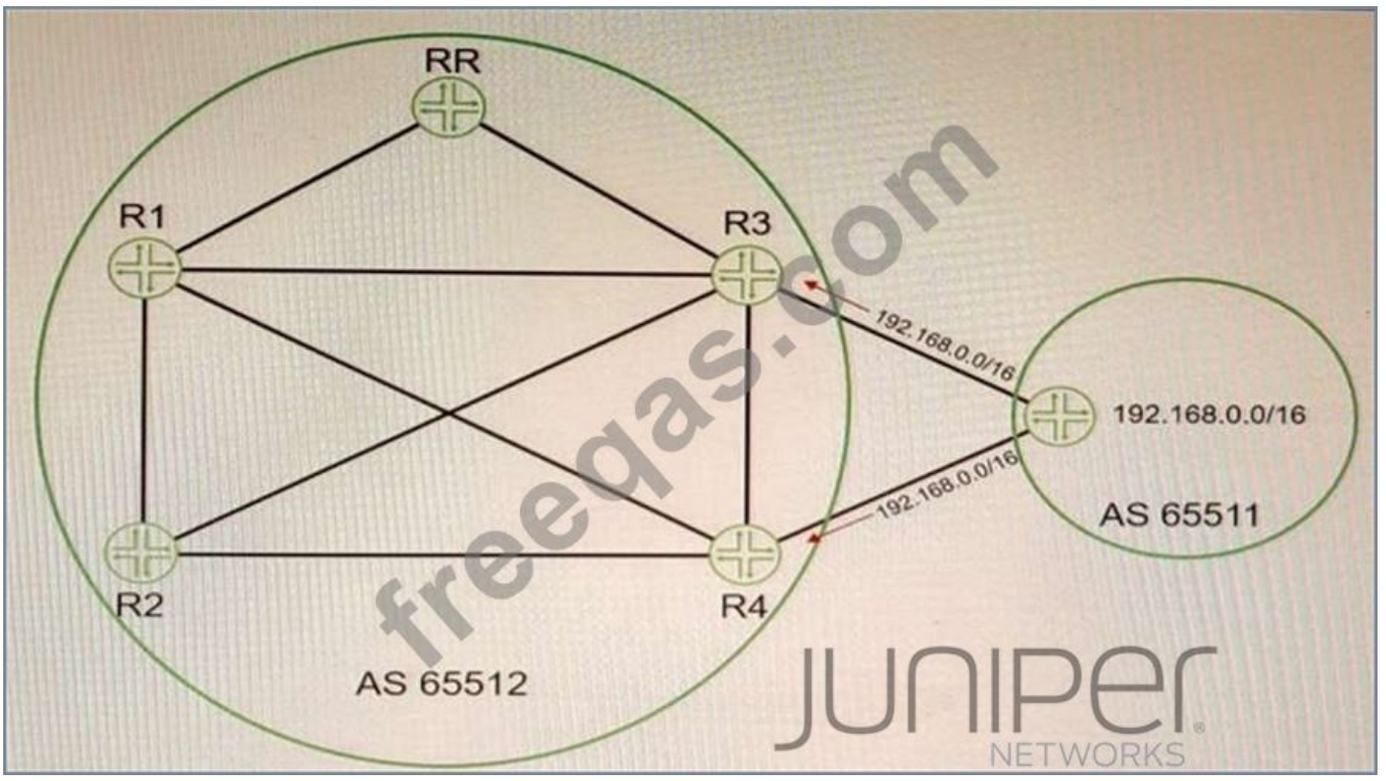


Referring to the exhibit, you are receiving the 192.168 0 0/16 route on both R3 and R4 from your EBGP neighbor You must ensure that R1 and R2 receive both BGP routes from the route reflector In this scenario, which BGP feature should you configure to accomplish this behavior?

**A.** add-path

**B.** multihop

**C.** multipath

**D.** route-target

**Answer: (SHOW ANSWER)**

Explanation

BGP add-path is a feature that allows the advertisement of multiple paths through the same peering session for the same prefix without the new paths implicitly replacing any previous paths. This behavior promotes path diversity and reduces multi-exit discriminator (MED) oscillations. BGP add-path is implemented by adding a path identifier to each path in the NLRI. The path identifier can be considered as something similar to a route distinguisher in VPNs, except that a path ID can apply to any address family. Path IDs are unique to a peering session and are generated for each network3. In this question, we have a route reflector (RR) that receives two routes for the same prefix (192.168.0.0/16) from an EBGP neighbor. By default, the RR will only advertise its best path to its clients (R1 and R2). However, we want R1 and R2 to receive both routes from the RR. To achieve this, we need to configure BGP add-path on the RR and enable it to send multiple paths for the same prefix to its clients.

**NEW QUESTION: 14**

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

**A.** Bootstrap messages are forwarded only to routers that explicitly requested the messages within the PIM sparse-mode domain
**B.** Bootstrap messages distribute RP information dynamically during an RP election.
**C.** Bootstrap messages are used to notify which router is the PIM RP
**D.** Bootstrap messages are forwarded to all routers within a PIM sparse-mode domain.
**Answer: B,D (LEAVE A REPLY)**
Bootstrap messages are PIM messages that are used to distribute rendezvous point (RP) information dynamically during an RP election. Bootstrap messages are sent by bootstrap routers (BSRs), which are routers that are elected to perform the RP discovery function for a PIM sparse-mode domain. Bootstrap messages contain information about candidate RPs and their multicast groups, as well as BSR priority and hash mask length. Bootstrap messages are forwarded to all routers within a PIM sparse-mode domain using hop-by-hop flooding.

**NEW QUESTION: 15**
Exhibit

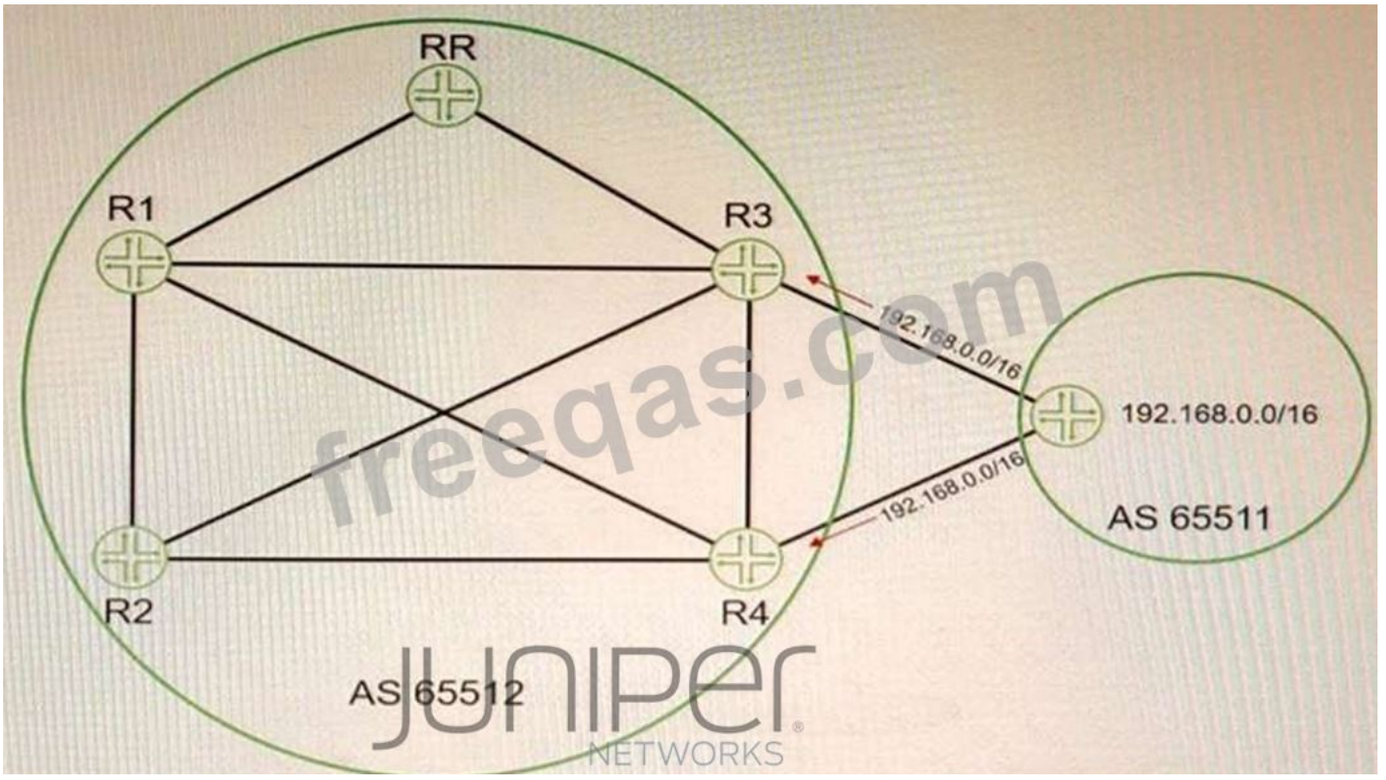


Referring to the exhibit, you are receiving the 192.168 0 0/16 route on both R3 and R4 from your EBGP neighbor You must ensure that R1 and R2 receive both BGP routes from the route reflector In this scenario, which BGP feature should you configure to accomplish this behavior?
**A.** add-path
**B.** multihop
**C.** multipath
**D.** route-target
**Answer: A (LEAVE A REPLY)**

BGP add-path is a feature that allows the advertisement of multiple paths through the same peering session for the same prefix without the new paths implicitly replacing any previous paths. This behavior promotes path diversity and reduces multi-exit discriminator (MED) oscillations. BGP add-path is implemented by adding a path identifier to each path in the NLRI. The path identifier can be considered as something similar to a route distinguisher in VPNs, except that a path ID can apply to any address family. Path IDs are unique to a peering session and are generated for each network3. In this question, we have a route reflector (RR) that receives two routes for the same prefix (192.168.0.0/16) from an EBGP neighbor. By default, the RR will only advertise its best path to its clients (R1 and R2). However, we want R1 and R2 to receive both routes from the RR. To achieve this, we need to configure BGP add-path on the RR and enable it to send multiple paths for the same prefix to its clients.

**NEW QUESTION: 16**

Exhibit

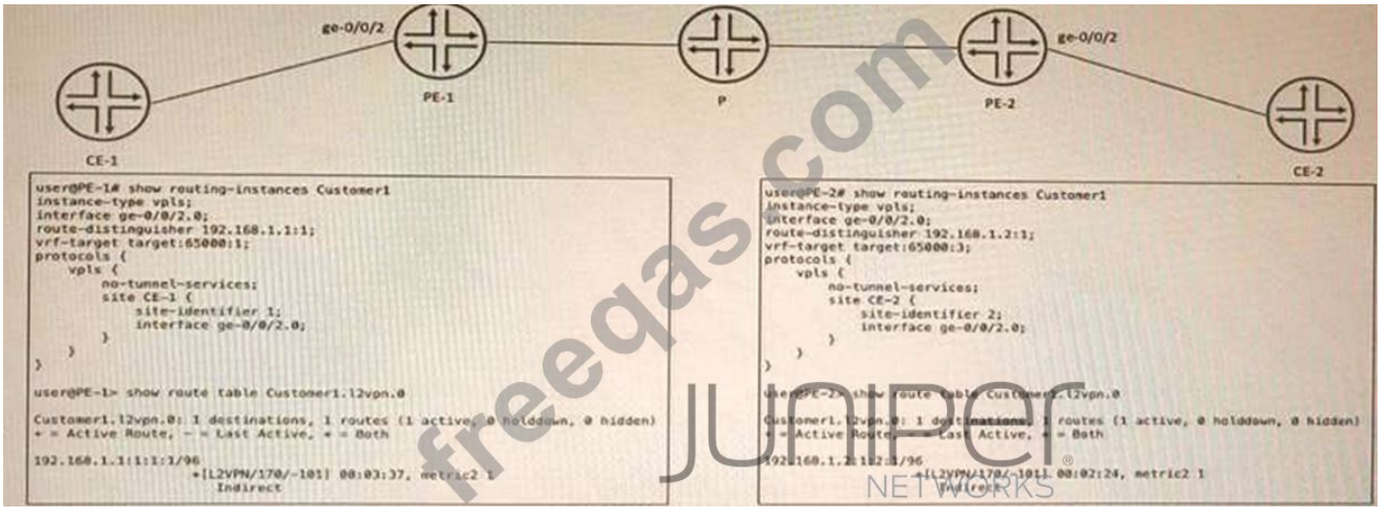


CE-1 and CE-2 are part of a VPLS called Customer1 No connectivity exists between CE-1 and CE-2. In the process of troubleshooting, you notice PE-1 is not learning any routes for this VPLS from PE-2, and PE-2 is not learning any routes for this VPLS from PE-1.

**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 (LEAVE A REPLY)**

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: 17**
When building an interprovider VPN, you notice on the PE router that you have hidden routes which are received from your BGP peer with family inet labeled-unica3t configured.
Which parameter must you configure to solve this problem?
**A.** Under the family inet labeled-unicast hierarchy, add the explicit null parameter.
**B.** Under the protocols ospf hierarchy, add the traffic-engineering parameter.
**C.** Under the family inet labeled-unicast hierarchy, add the resolve-vpn parameter.
**D.** Under the protocols mpls hierarchy, add the traffic-engineering parameter
**Answer: (SHOW ANSWER)**
Explanation
The resolve-vpn parameter is a BGP option that allows a router to resolve labeled VPN-IPv4 routes using unlabeled IPv4 routes received from another BGP peer with family inet labeled-unicast configured. This option enables interprovider VPNs without requiring MPLS labels between ASBRs or using VRF tables on ASBRs. In this scenario, you need to configure the resolve-vpn parameter under [edit protocols bgp group external family inet labeled-unicast] hierarchy level on both ASBRs.

**NEW QUESTION: 18**
Exhibit.

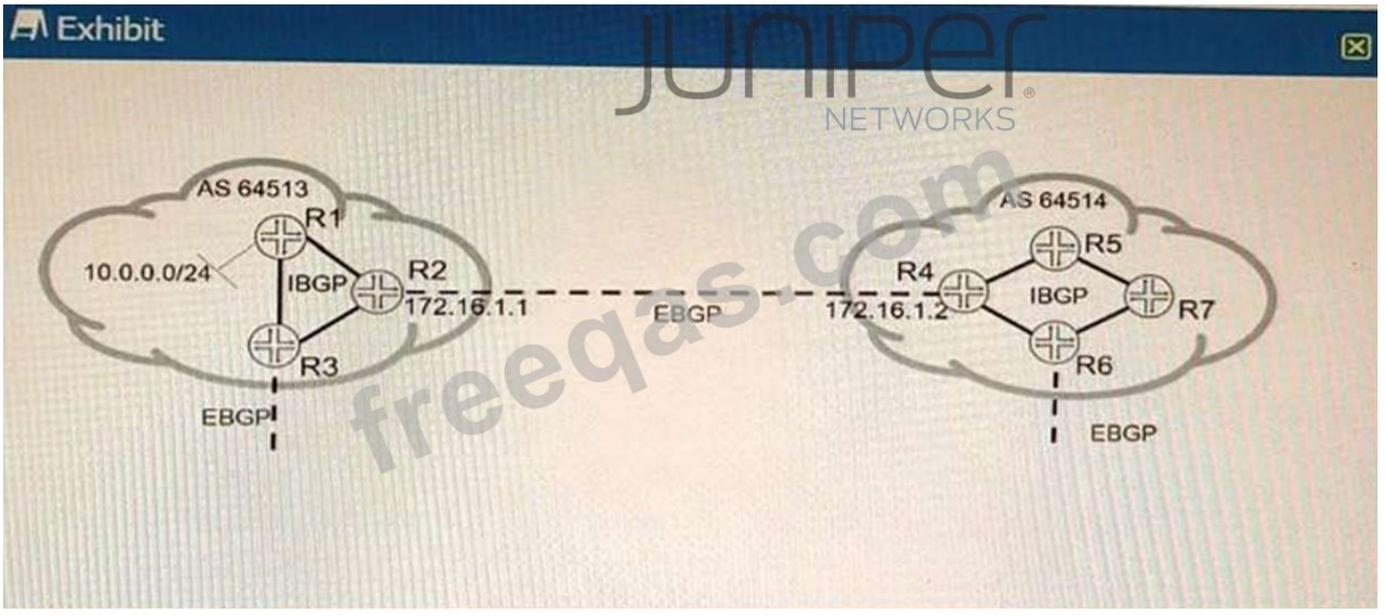

Referring to the exhibit; the 10.0.0.0/24 EBGP route is received on R5; however, the route is being hidden.
What are two solutions that will solve this problem? (Choose two.)
**A.** On R4, create a policy to change the BGP next hop to itself and apply it to IBGP as an export policy
**B.** Add the external 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: A,B (LEAVE A REPLY)**
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 to 2, this is a possible workaround for a pure IP network, but it may not work well for an MPLS network.
The reason why the route is being hidden is that R5 cannot reach the BGP next hop 10.0.0.1, which is the address of R1. R5 does not have a route to 10.0.0.0/24 in its routing table, and neither does R4. Therefore, R5 cannot resolve the BGP next hop and marks the route as hidden.
There are two solutions that will solve this problem:
Option A: On R4, create a policy to change the BGP next hop to itself and apply it to IBGP as an export policy. This way, R5 will receive the route with a next hop of 172.16.1.2, which is reachable via the IGP. This solution is also known as next-hop-self1.
Option B: Add the external interface prefix to the IGP routing tables. This way, R4 and R5 will learn a route to 10.0.0.0/24 via the IGP and be able to resolve the BGP next hop. This solution is also known as recursive lookup2.
Option C is not correct because adding the internal interface prefix to the BGP routing tables will not help R5 reach the BGP next hop 10.0.0.1.
Option D is not correct because changing the BGP next hop to 172.16.1.1 on R4 will not help R5 either, since R5 does not have a route to 172.16.1.1 in its routing table.
References: 1: Configuring Next-Hop-Self for IBGP Peers 2: Understanding Recursive Lookup

**NEW QUESTION: 19**
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 R2 and R3.
**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 R1 and R4.
**D.** Configure the BGP group with the as-override parameter on R1 and R4
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 20**
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: A,C (LEAVE A REPLY)**
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: 21**
You are asked to protect your company's customers from amplification attacks. In this scenario, what is Juniper's recommended protection method?
**A.** ASN prepending
**B.** BGP FlowSpec
**C.** destination-based Remote Triggered Black Hole
**D.** unicast Reverse Path Forwarding
**Answer: C (LEAVE A REPLY)**

amplification attacks are a type of distributed denial-of-service (DDoS) attack that exploit the characteristics of certain protocols to amplify the traffic sent to a victim. For example, an attacker can send a small DNS query with a spoofed source IP address to a DNS server, which will reply with a much larger response to the victim. This way, the attacker can generate a large amount of traffic with minimal resources.
One of the methods to protect against amplification attacks is destination-based Remote Triggered Black Hole (RTBH) filtering. This technique allows a network operator to drop traffic destined to a specific IP address or prefix at the edge of the network, thus preventing it from reaching the victim and consuming bandwidth and resources. RTBH filtering can be implemented using BGP to propagate a special route with a next hop of
192.0.2.1 (a reserved address) to the edge routers. Any traffic matching this route will be discarded by the edge routers.

**NEW QUESTION: 22**
Exhibit.

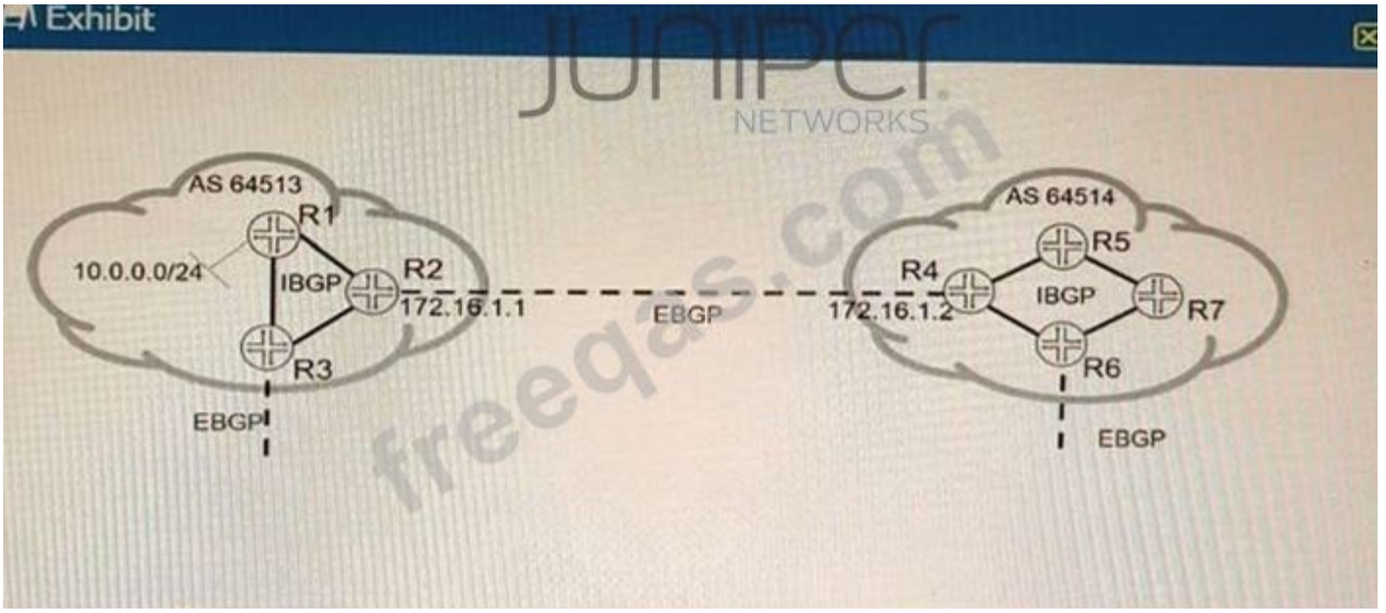


Referring to the exhibit; the 10.0.0.0/24 EBGP route is received on R5; however, the route is being hidden.
What are two solutions that will solve this problem? (Choose two.)
**A.** On R4, create a policy to change the BGP next hop to itself and apply it to IBGP as an export policy
**B.** Add the external 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: A,B (LEAVE A REPLY)**
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 to 2, this is a possible workaround for a pure IP network, but it may not work well for an MPLS network.

**NEW QUESTION: 23**
Your network is receiving the 203.0.113.0/24 network using EBGP from AS 64500 and AS 64501. Both of these advertisements have identical local-preference values, AS-path lengths, and BGP origin codes. You want to influence the way your AS sends traffic to the 203.0.113.0/24 network. In this scenario, which attribute would you consider next when selecting the best path?
**A.** router ID
**B.** MED value
**C.** peer IP address
**D.** IGP metric
**Answer: (SHOW ANSWER)**
as by default, the MED attribute is only compared for routes received from the same neighbouring AS. The next feasible tiebreaker in the BGP route selection algorithm would be Router ID.

**NEW QUESTION: 24**
Which statement is true regarding BGP FlowSpec?
**A.** It uses a remote triggered black hole to protect a network from a denial-of-service attack.
**B.** It uses dynamically created routing policies to protect a network from denial-of-service attacks
**C.** It is used to protect a network from denial-of-service attacks dynamically
**D.** It verifies that the source IP of the incoming packet has a resolvable route in the routing table
**Answer: B (LEAVE A REPLY)**
Explanation
BGP FlowSpec is a feature that extends the Border Gateway Protocol (BGP) to enable routers to exchange traffic flow specifications, allowing for more precise control of network traffic. The BGP FlowSpec feature enables routers to advertise and receive information about specific flows in the network, such as those originating from a particular source or destined for a particular destination. Routers can then use this information to construct traffic filters that allow or deny packets of a certain type, rate limit flows, or perform other actions1. BGP FlowSpec can also help in filtering traffic and taking action against distributed denial of service (DDoS) attacks by dropping the DDoS traffic or diverting it to an analyzer2. BGP FlowSpec rules are internally converted to equivalent Cisco Common Classification Policy Language (C3PL) representing corresponding match and action parameters2. Therefore, BGP FlowSpec uses dynamically created routing policies to protect a network from denial-of-service attacks.
References: 1: https://www.networkingsignal.com/what-is-bgp-flowspec/ 2:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-16/irg-xe-16-book/bgp-flowspe

**NEW QUESTION: 25**
Which origin code is preferred by BGP?
**A.** External
**B.** Null
**C.** Incomplete
**D.** Internal
**Answer: (SHOW ANSWER)**
Explanation
BGP uses several attributes to select the best path for a destination prefix. One of these attributes is origin, which indicates how BGP learned about a route. The origin attribute can have one of three values: IGP, EGP, or Incomplete. IGP means that the route was originated by a network or aggregate statement within BGP or by redistribution from an IGP into BGP. EGP means that the route was learned from an external BGP peer (this value is obsolete since BGP version 4). Incomplete means that the route was learned by some other means, such as redistribution from a static route into BGP. BGP prefers routes with lower origin values, so Incomplete is preferred over EGP, which is preferred over IGP.

**NEW QUESTION: 26**
Exhibit
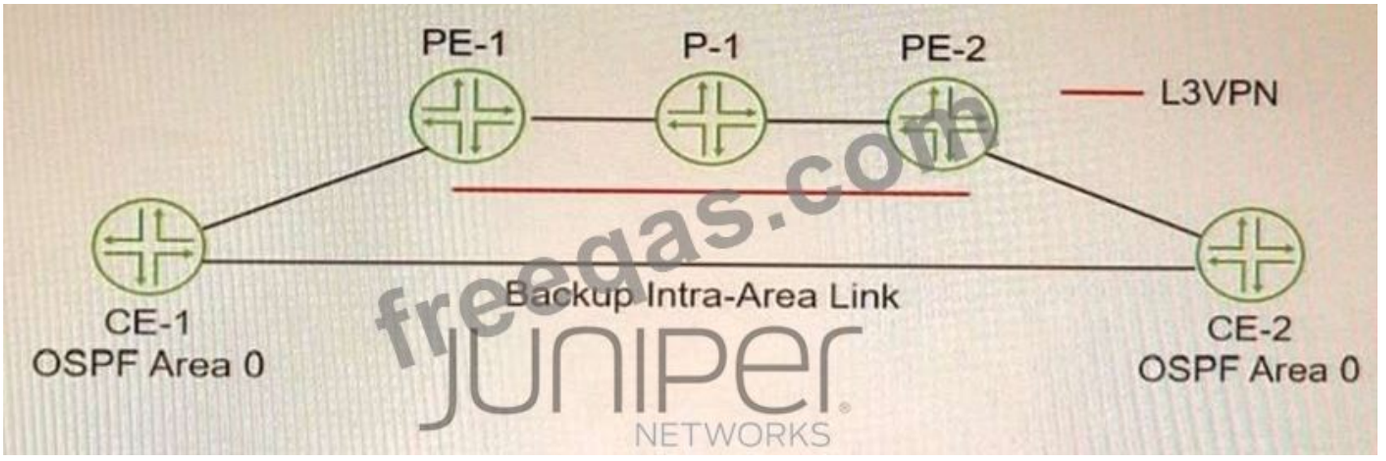

You must ensure that the VPN backbone is preferred over the back door intra-area link as long as the VPN is available. Referring to the exhibit, which action will accomplish this task?
**A.** Configure an import routing policy on the CE routers that rejects OSPF routes learned on the backup intra-area 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 (LEAVE A REPLY)**

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: 27**

Exhibit

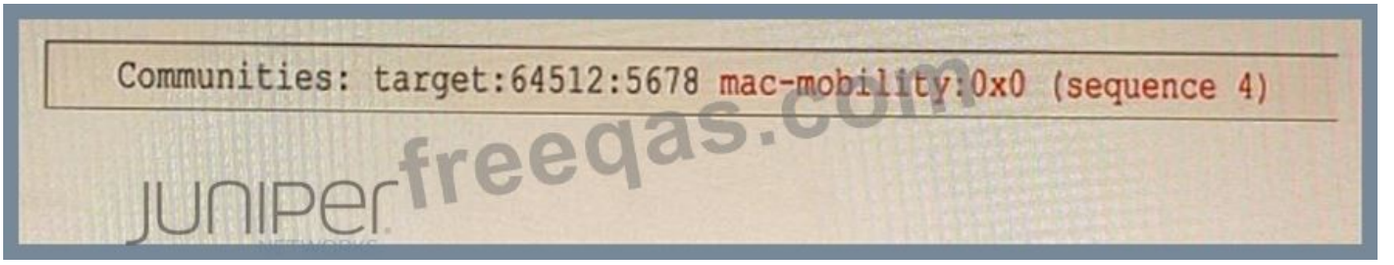



You have MAC addresses moving in your EVPN environment

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: B,C (LEAVE A REPLY)**

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: 28**

Exhibit

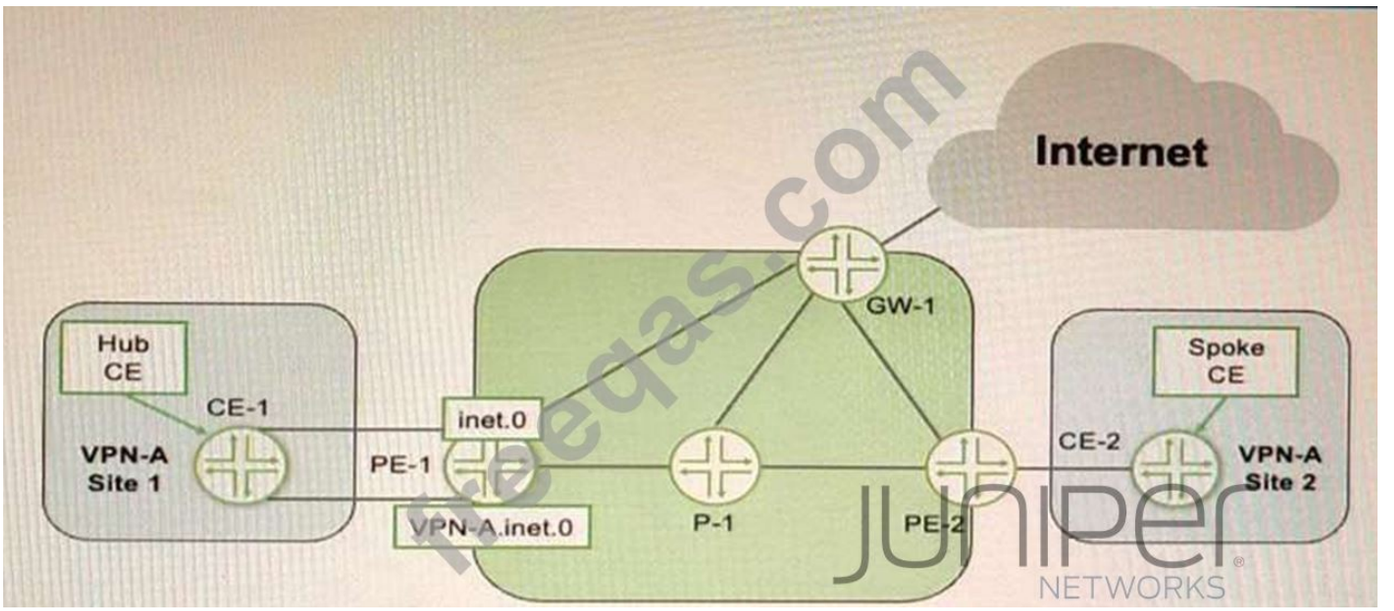


Referring to the exhibit, you must provide Internet access for VPN-A using CE-1 as the hub CE. Which two statements are correct in this situation? (Choose two.)

**A.** You must use RIB groups to leak routes between the inet. o and vpn-a. inet. o tables.

**B.** RIB groups are not needed to leak routes between the inet. 0 and VPN-A. inet. 0 tables,

**C.** Internet traffic from Site 2 takes the path of PE-2 -> PE-1 -> GW-1.

**D.** Internet traffic from Site 2 takes the path of PE-2 -> PE-1 -> CE-1 -> PE-1 -> GW-1.

**Answer: A,D (LEAVE A REPLY)**

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: 29**

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

**A.** CSNPs are flooded periodically.

**B.** PSNPs are flooded periodically.

**C.** PSNPs contain only descriptions of LSPs.

**D.** CSNPs contain only descriptions of LSPs.

**Answer: (SHOW ANSWER)**

LSPs contain information about the state and cost of links in the network, and are flooded periodically throughout the network. PSNPs are used to acknowledge receipt of LSPs and request retransmission of missing or corrupted LSPs. PSNPs contain only descriptions of LSPs, such as their sequence numbers and checksums. CSNPs contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases.

**NEW QUESTION: 30**
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: A,B (LEAVE A REPLY)**
https://www.juniper.net/documentation/us/en/software/junos/cos-security-devices/topics/concept/cos-transmissio

**NEW QUESTION: 31**
Exhibit

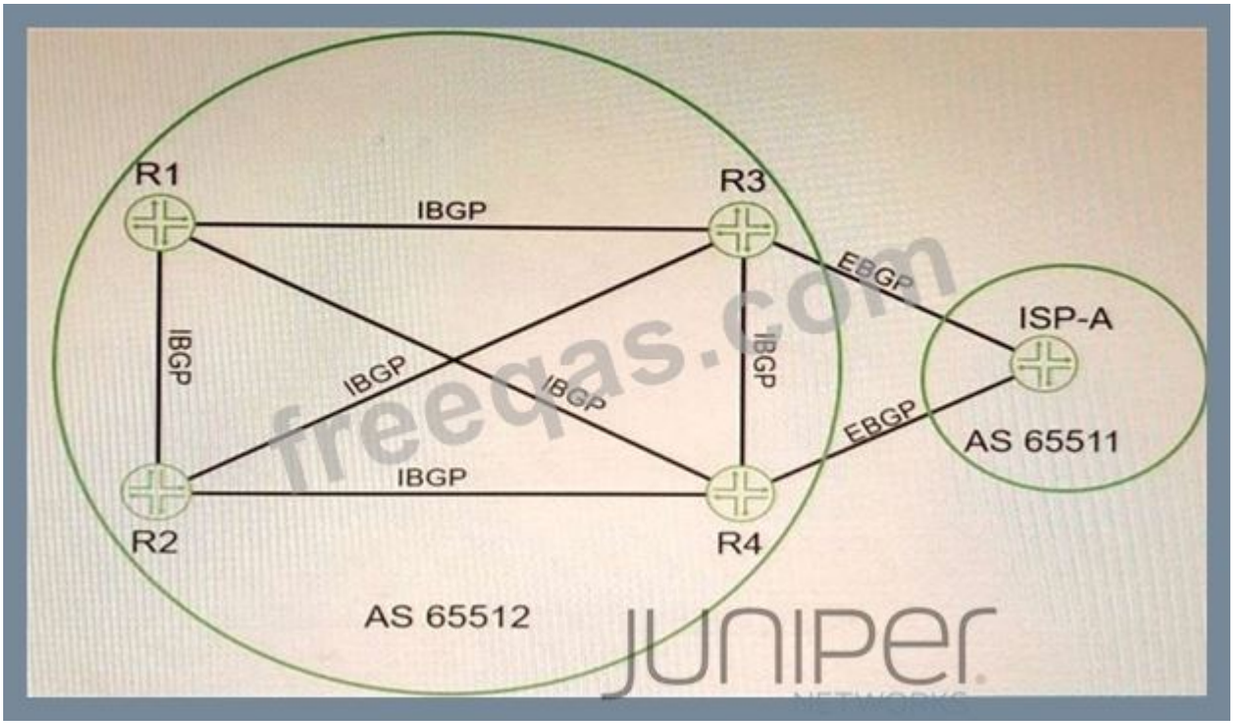


Click the Exhibit button-Referring to the exhibit, which two statements are correct about BGP routes on R3 that are learned from the ISP-A neighbor? (Choose two.)
**A.** By default, the next-hop value for these routes is not changed by ISP-A before being sent to R3.
**B.** The BGP local-preference value 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: (SHOW ANSWER)**
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: 32**
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 (LEAVE A REPLY)**
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: 33**
Exhibit

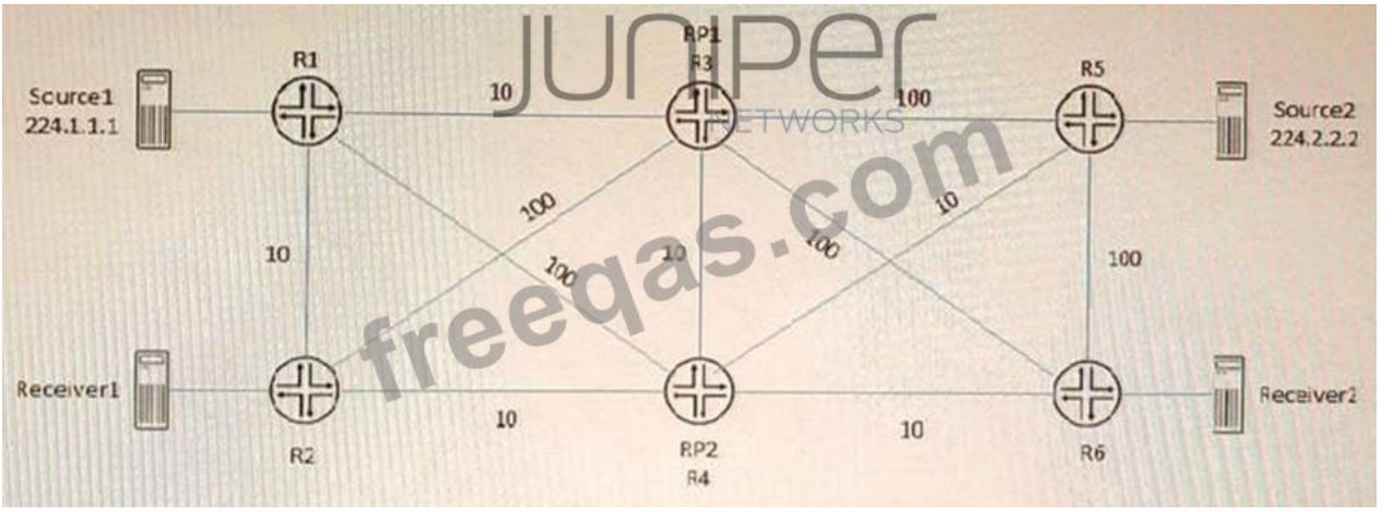


Referring to the exhibit, PIM-SM is configured on all routers, and Anycast-RP with Anycast-PIM is used for the discovery mechanism on RP1 and RP2. The interface metric values are shown for the OSPF area.

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: A,C (LEAVE A REPLY)**

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: 34**

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

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

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

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

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

**Answer: C,D (LEAVE A REPLY)**

The customer interface in an LDP-signaled pseudowire is the interface on the PE router that connects to the CE device. An LDP-signaled pseudowire is a type of Layer 2 circuit that uses LDP

to establish a point-to-point connection between two PE routers over an MPLS network. The customer interface can have different encapsulation types depending on the type of traffic that is carried over the pseudowire. The encapsulation types are ethernet-ccc, vlan-ccc, extended-vlan-ccc, atm-ccc, frame-relay-ccc, ppp-ccc, cisco-hdlc-ccc, and tcc-ccc. Depending on the encapsulation type, the customer interface can accept or reject tagged or untagged frames in the data plane, and include or exclude VLAN tags in the control plane LDP advertisement. The following table summarizes the behavior of different encapsulation types:

**NEW QUESTION: 35**
Exhibit

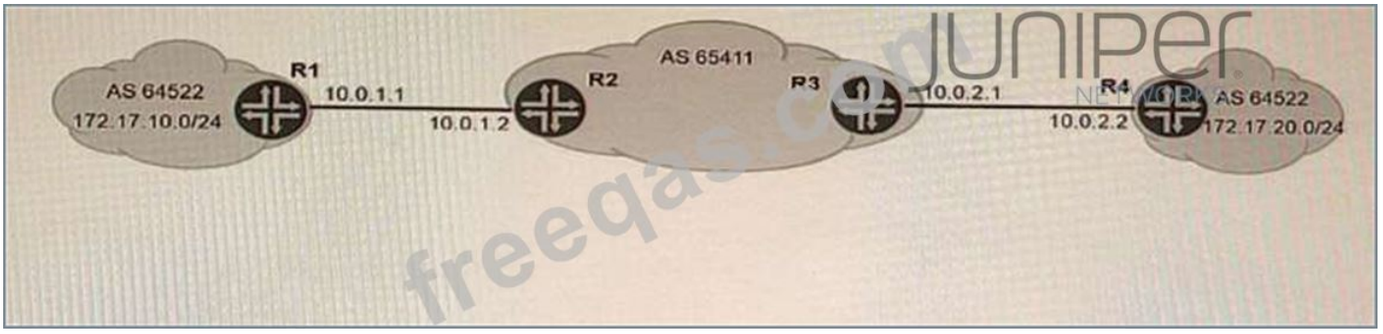


You are asked to exchange routes between R1 and R4 as shown in the exhibit. These two routers use the same AS number Which two steps will accomplish this task? (Choose two.)
**A.** Configure the BGP group with the advertise-peer-as parameter on R1 and R4.
**B.** Configure the BGP group with the as-override 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: A,B (LEAVE A REPLY)**
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: 36**
Exhibit

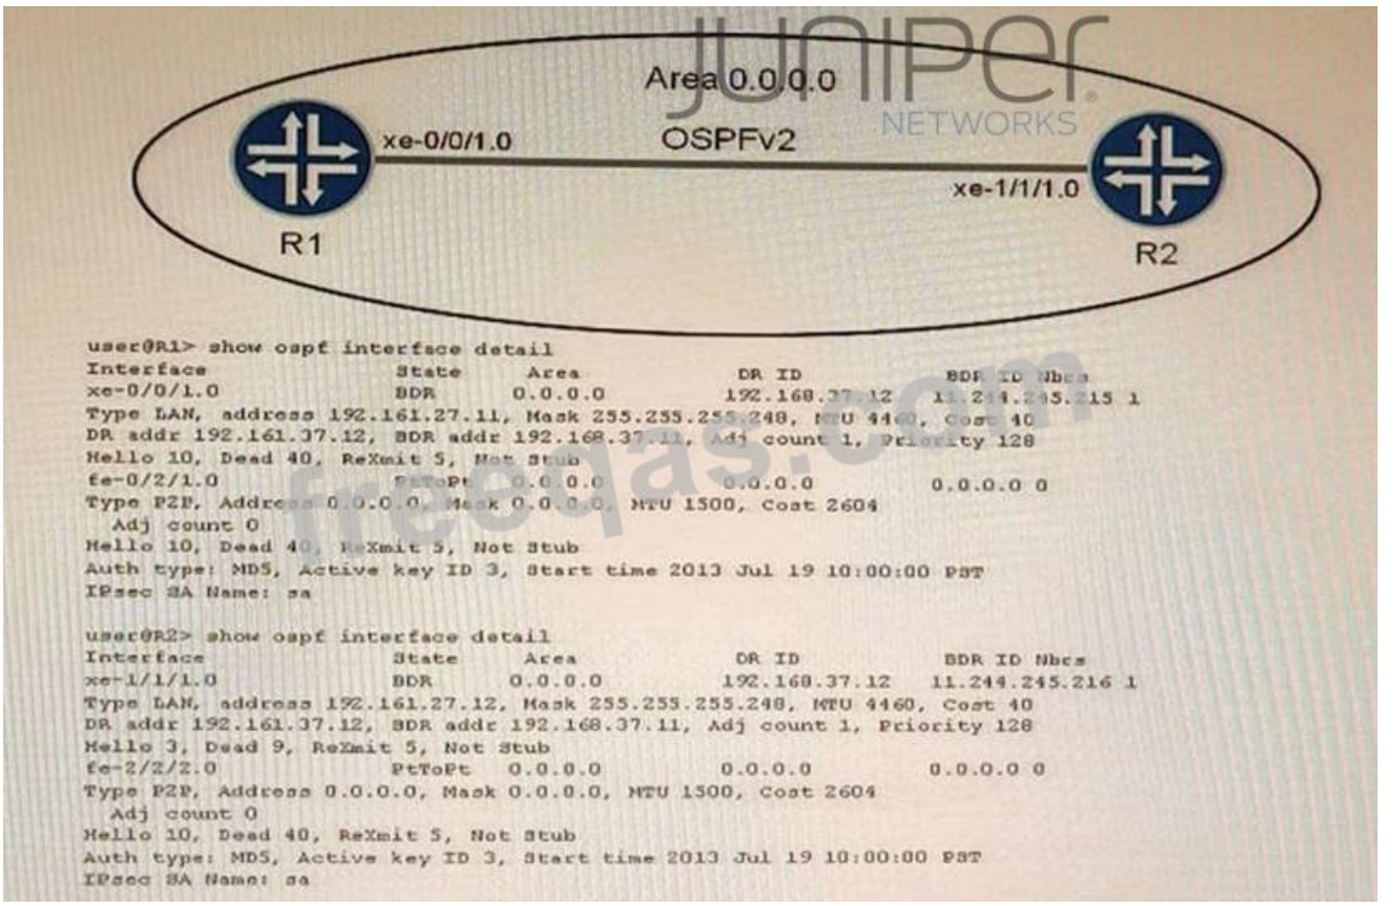


Which two statements are true about the OSPF adjacency displayed in the exhibit? (Choose two.)

**A.** There is a mismatch in the hello interval parameter between routers R1 and R2

**B.** There is a mismatch in the dead interval parameter between routers R1 and R2.

**C.** There is a mismatch in the OSPF hold timer parameter between routers R1 and R2.

**D.** There is a mismatch in the poll interval parameter between routers R1 and R2.

**Answer: A,B (LEAVE A REPLY)**

Explanation

The hello interval is the time interval between two consecutive hello packets sent by an OSPF router on an interface. The dead interval is the time interval after which a neighbor is declared down if no hello packets are received from it. These parameters must match between two OSPF routers for them to form an adjacency. In the exhibit, router R1 has a hello interval of 10 seconds and a dead interval of 40 seconds, while router R2 has a hello interval of 30 seconds and a dead interval of 120 seconds. This causes a mismatch and prevents them from becoming neighbors23.

**NEW QUESTION: 37**

Your organization manages a Layer 3 VPN for multiple customers To support advanced route than one BGP community on advertised VPN routes to remote PE routers.

Which routing-instance configuration parameter would support this requirement?

**A.** vrf-export

**B.** vrf-import

**C.** vrf-target export

**D.** vrf-target import

**Answer: (SHOW ANSWER)**

The vrf-target export parameter is used to specify one or more BGP extended community attributes that are attached to VPN routes when they are exported from a VRF routing instance to remote PE routers. This parameter allows you to control which VPN routes are accepted by remote PE routers based on their import policies. You can specify more than one vrf-target export value for a VRF routing instance to support advanced route filtering or route leaking scenarios.

**NEW QUESTION: 38**

An interface is configured with a behavior aggregate classifier and a multifield classifier How will the packet be processed when received on this interface?

**A.** The packet will be discarded.

**B.** The packet will be processed by the BA classifier first, then the MF classifier.

**C.** The packet will be forwarded with no classification changes.

**D.** The packet will be processed by the MF classifier first, then the BA classifier.

**Answer: C (LEAVE A REPLY)**

behavior aggregate (BA) classifiers and multifield (MF) classifiers are two types of classifiers that are used to assign packets to a forwarding class and a loss priority based on different criteria. The forwarding class determines the output queue for a packet. The loss priority is used by a scheduler to control packet discard during periods of congestion.

A BA classifier maps packets to a forwarding class and a loss priority based on a fixed-length field in the packet header, such as DSCP, IP precedence, MPLS EXP, or IEEE 802.1p CoS bits. A BA classifier is computationally efficient and suitable for core devices that handle high traffic volumes. A BA classifier is useful if the traffic comes from a trusted source and the CoS value in the packet header is trusted.

An MF classifier maps packets to a forwarding class and a loss priority based on multiple fields in the packet header, such as source address, destination address, protocol type, port number, or VLAN ID. An MF classifier is more flexible and granular than a BA classifier and can match packets based on complex filter rules. An MF classifier is suitable for edge devices that need to classify traffic from untrusted sources or rewrite packet headers.

You can configure both a BA classifier and an MF classifier on an interface. If you do this, the BA classification is performed first and then the MF classification. If the two classification results conflict, the MF classification result overrides the BA classification result.

Based on this information, we can infer the following statements:

The packet will be discarded. This is not correct because the packet will not be discarded by the classifiers unless it matches a filter rule that specifies discard as an action. The classifiers only assign packets to a forwarding class and a loss priority based on their match criteria.

The packet will be processed by the BA classifier first, then the MF classifier. This is correct because if both a BA classifier and an MF classifier are configured on an interface, the BA classification is performed first and then the MF classification. If they conflict, the MF classification result overrides the BA classification result.

The packet will be forwarded with no classification changes. This is not correct because the packet will be classified by both the BA classifier and the MF classifier if they are configured on an interface. The final classification result will determine which output queue and which discard policy will be applied to the packet.
The packet will be processed by the MF classifier first, then the BA classifier. This is not correct because if both a BA classifier and an MF classifier are configured on an interface, the BA classification is performed first and then the MF classification. If they conflict, the MF classification result overrides the BA classification result.

**NEW QUESTION: 39**
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: (SHOW ANSWER)**
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: 40**

Exhibit

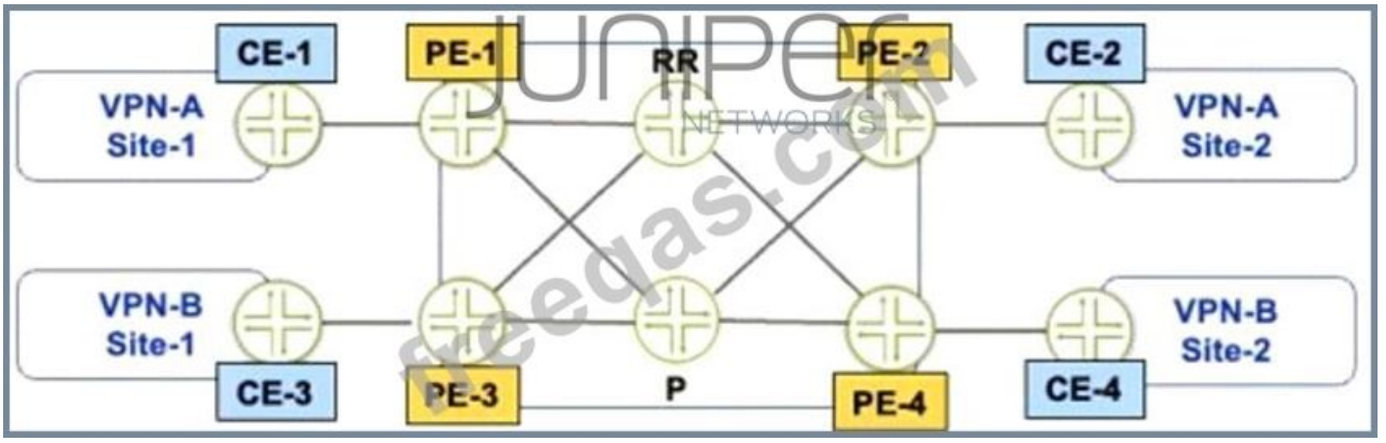


Referring to the exhibit, PE-1 and PE-2 are getting route updates for VPN-B when neither of them service that VPN Which two actions would optimize this process? (Choose two.)

**A.** Configure the resolution rib bgp.l3vpn.0 resolution-ribs inet.0 statement on the PEs.

**B.** Configure the family route-target statement on the RR.

**C.** Configure the resolution rib bgp.l3vpn.0 resolution-ribs inet.0 statement on the RR.

**D.** Configure the family route-target statement on the PEs.

**Answer: B,C (LEAVE A REPLY)**

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.

**NEW QUESTION: 41**

Your organization manages a Layer 3 VPN for multiple customers To support advanced route than one BGP community on advertised VPN routes to remote PE routers.

Which routing-instance configuration parameter would support this requirement?

**A.** vrf-export

**B.** vrf-import

**C.** vrf-target export

**D.** vrf-target import

**Answer: (SHOW ANSWER)**

Explanation

The vrf-target export parameter is used to specify one or more BGP extended community attributes that are attached to VPN routes when they are exported from a VRF routing instance to remote PE routers. This parameter allows you to control which VPN routes are accepted by

remote PE routers based on their import policies. You can specify more than one vrf-target export value for a VRF routing instance to support advanced route filtering or route leaking scenarios.

**NEW QUESTION: 42**
After a recent power outage, your manager asks you to investigate ways to automatically reduce the impact caused by suboptimal routing in your OSPF and OSPFv3 network after devices reboot.
Which three configuration statements accomplish this task? (Choose three.)
**A.** set protocols ospf overload timeout 900
**B.** set protocols ospf3 realm ipv4-unicast overload timeout 900
**C.** set protocols ospf overload
**D.** set protocols oapf3 overload timeout 900
**E.** set protocols ospf3 overload
**Answer: A,E (LEAVE A REPLY)**
Explanation
To reduce the impact of suboptimal routing in OSPF and OSPFv3 after devices reboot, you can use the overload feature to prevent a router from being used as a transit router for a specified period of time. This allows the router to stabilize its routing table before forwarding traffic for other routers. To enable the overload feature, you need to do the following:
* For OSPF, configure the overload statement under [edit protocols ospf] hierarchy level. You can also specify a timeout value in seconds to indicate how long the router should remain in overload state after it boots up. For example, set protocols ospf overload timeout 900 means that the router will be in overload state for 15 minutes after it boots up.
* For OSPFv3, configure the overload statement under [edit protocols ospf3] hierarchy level. You can also specify a realm (ipv4-unicast or ipv6-unicast) and a timeout value in seconds to indicate how long the router should remain in overload state after it boots up for each realm. For example, set protocols ospf3 realm ipv4-unicast overload timeout 900 means that the router will be in overload state for 15 minutes after it boots up for IPv4 unicast routing.

**NEW QUESTION: 43**
An interface is configured with a behavior aggregate classifier and a multifield classifier How will the packet be processed when received on this interface?
**A.** The packet will be discarded.
**B.** The packet will be processed by the BA classifier first, then the MF classifier.
**C.** The packet will be forwarded with no classification changes.
**D.** The packet will be processed by the MF classifier first, then the BA classifier.
**Answer: (SHOW ANSWER)**
Explanation
behavior aggregate (BA) classifiers and multifield (MF) classifiers are two types of classifiers that are used to assign packets to a forwarding class and a loss priority based on different criteria.

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

**NEW QUESTION: 44**
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 local address on R2 to be higher than R3.

**B.** Change the bootstrap priority on R2 to be higher than R3

**C.** Change the default route in inet.2 on R4 from R3 as the next hop to R2

**D.** Change the group-range to be more specific on R2 than R3.

**Answer: (SHOW ANSWER)**

**NEW QUESTION: 45**

Exhibit

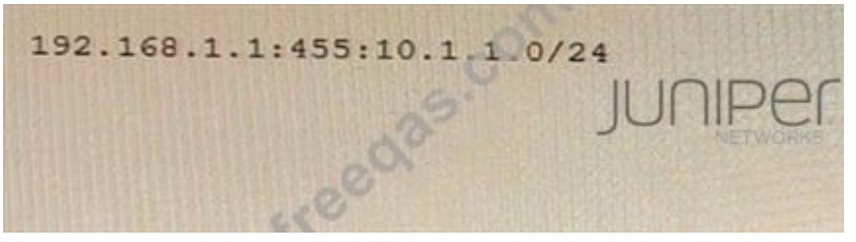


You are examining an L3VPN route that includes the information shown in the exhibit Which statement is correct in this scenario?

**A.** The information shows a Type 1 route distinguisher.

**B.** The information shows a Type 0 route distinguisher

**C.** The information shows a Type 2 route distinguisher.

**D.** The information shows a route target

**Answer: (SHOW ANSWER)**

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

**NEW QUESTION: 46**

Exhibit

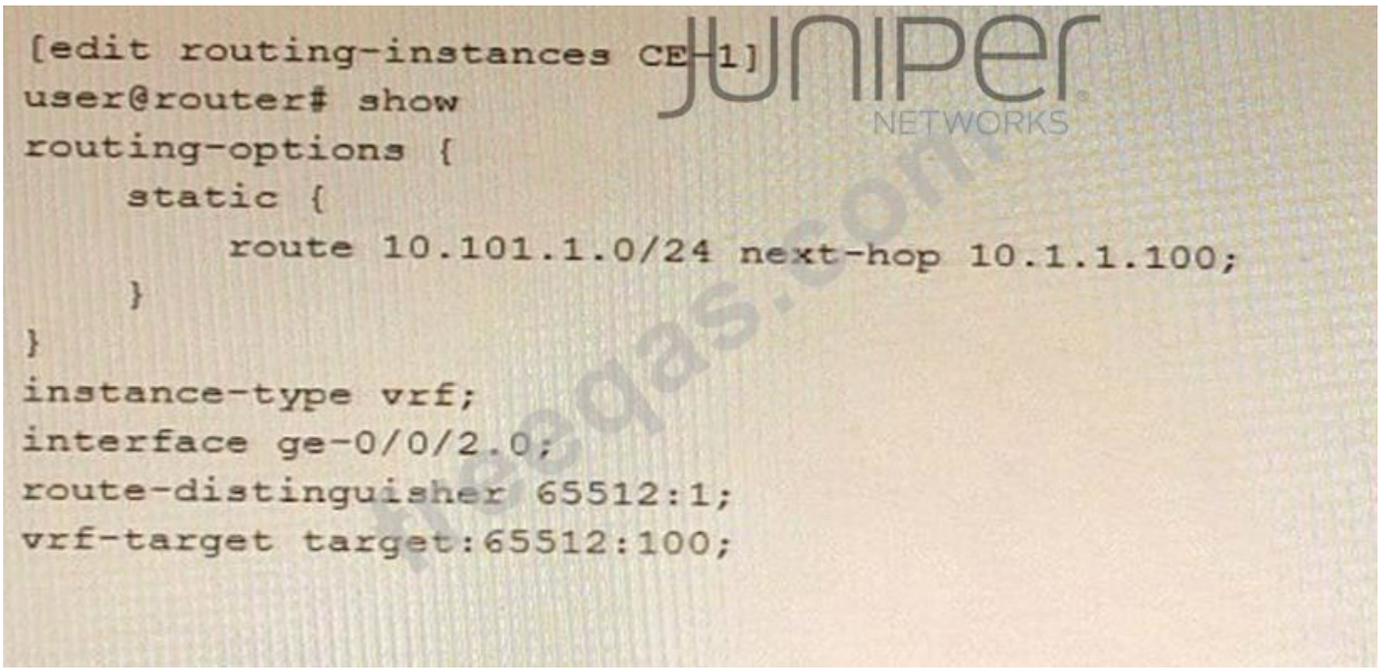


Referring to the exhibit, which statement is true?

**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 (LEAVE A REPLY)**

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: 47**

Which two statements are correct about reflecting inet-vpn unicast prefixes in BGP route reflection? (Choose two.)

**A.** Route reflectors do not change any existing BGP attributes by default when advertising routes.

**B.** A BGP peer does not require any configuration changes to become a route reflector client.
**C.** Clients add their originator ID when advertising routes to their route reflector
**D.** Route reflectors add their cluster ID to the AS path when readvertising client routes.
**Answer: A,B (LEAVE A REPLY)**
Explanation
Route reflection is a BGP feature that allows a router to reflect routes learned from one IBGP peer to another IBGP peer, without requiring a full-mesh IBGP topology. Route reflectors do not change any existing BGP attributes by default when advertising routes, unless explicitly configured to do so. A BGP peer does not require any configuration changes to become a route reflector client, only the route reflector needs to be configured with the client parameter under [edit protocols bgp group group-name neighbor neighbor-address] hierarchy level.

**NEW QUESTION: 48**
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: A,D (LEAVE A REPLY)**
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.