Juniper.JN0-664.v2024-08-24.q77

Exam Code:	JN0-664
Exam Name:	Service Provider, Professional (JNCIP-SP)
Certification Provider:	Juniper
Free Question Number:	77
Version:	v2024-08-24
# of views:	557
# of Questions views:	770
https://www.freepdfdumps.com/Juniper.JN0-664.v2024-08-24.q77.html	

NEW QUESTION: 1
Exhibit

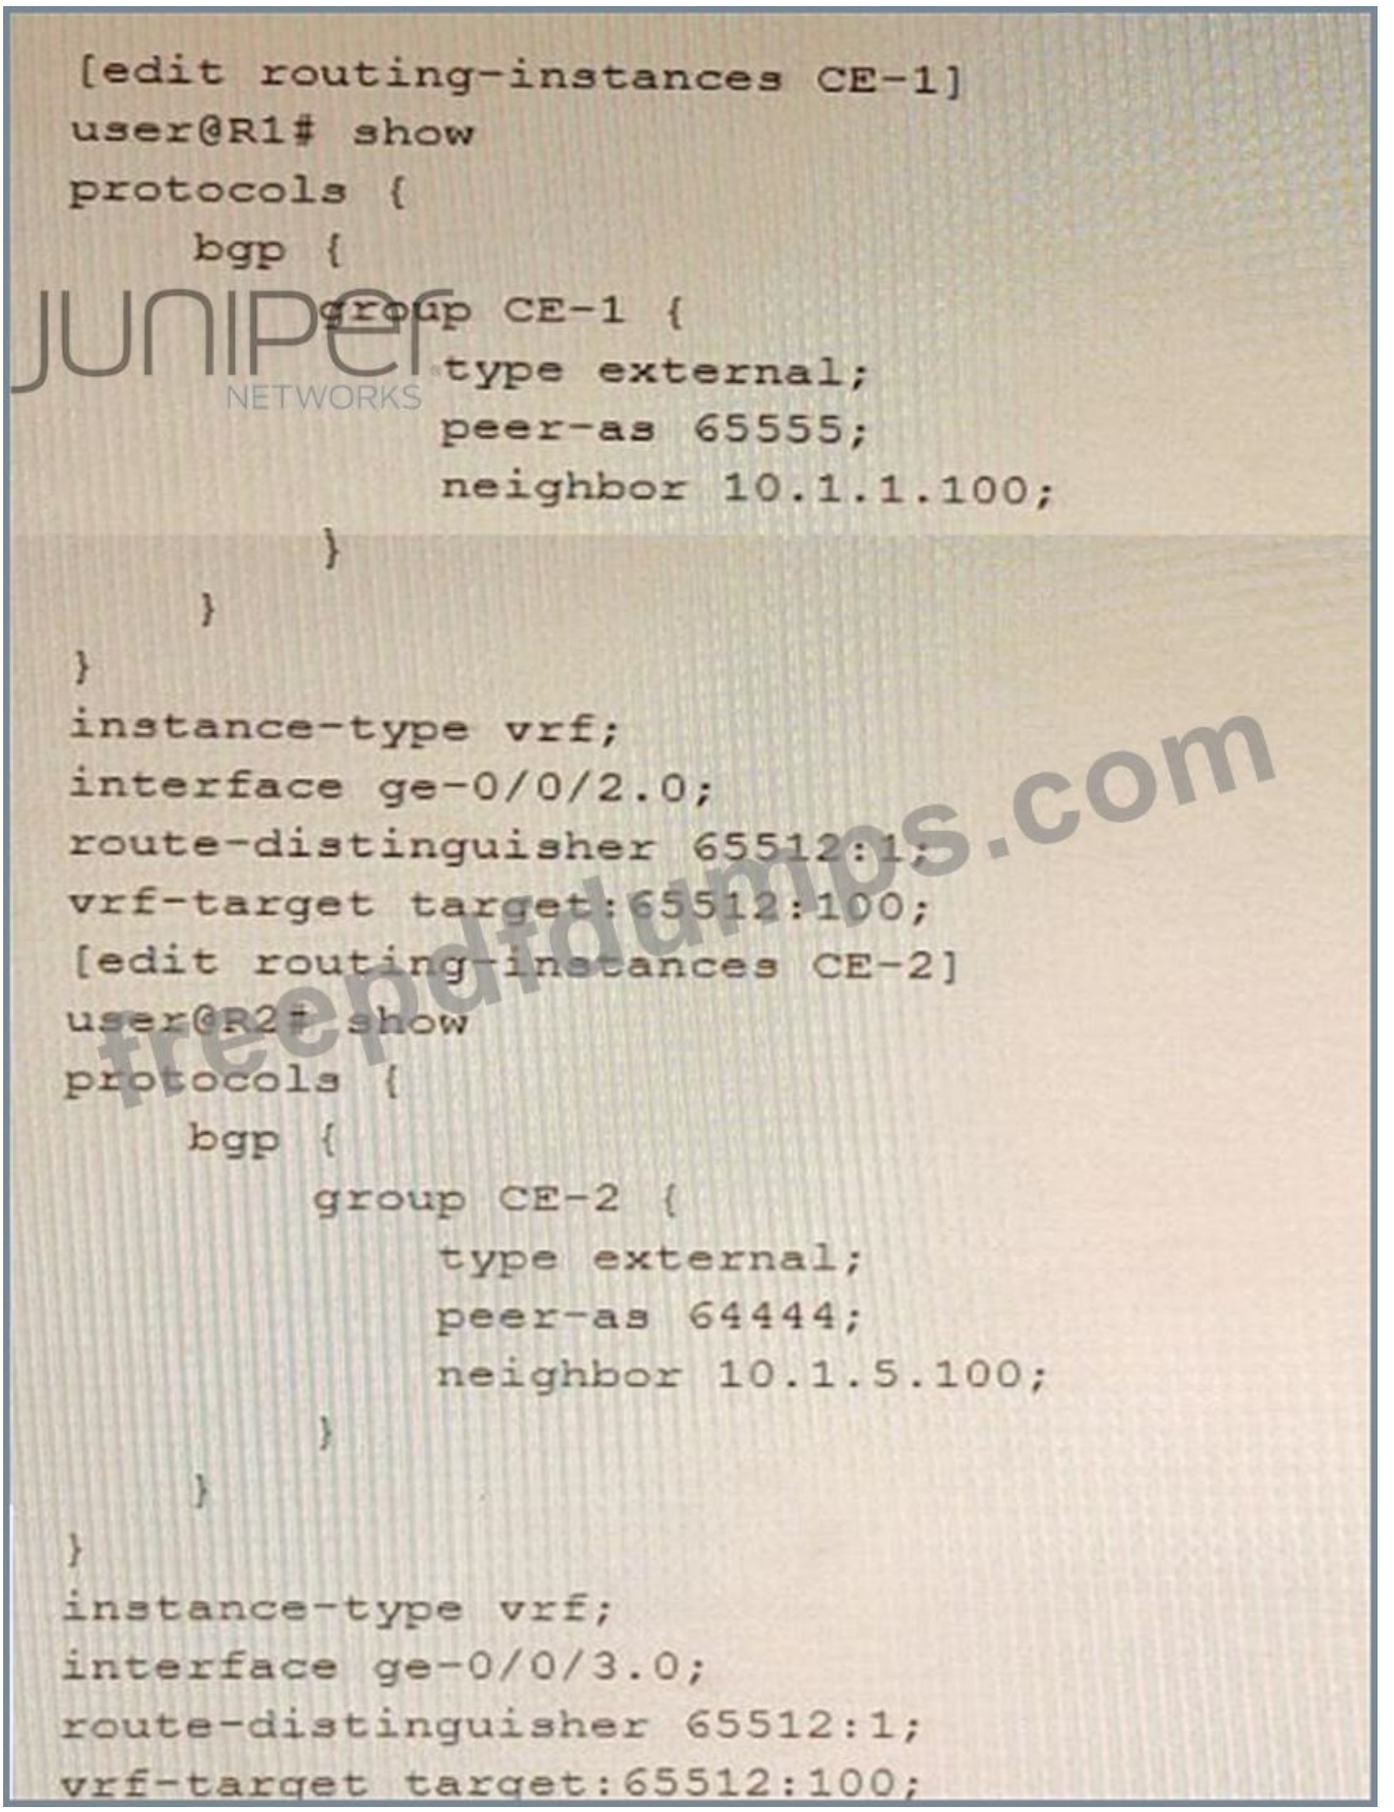

Referring to the exhibit, which statement is correct?

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: 2
Exhibit

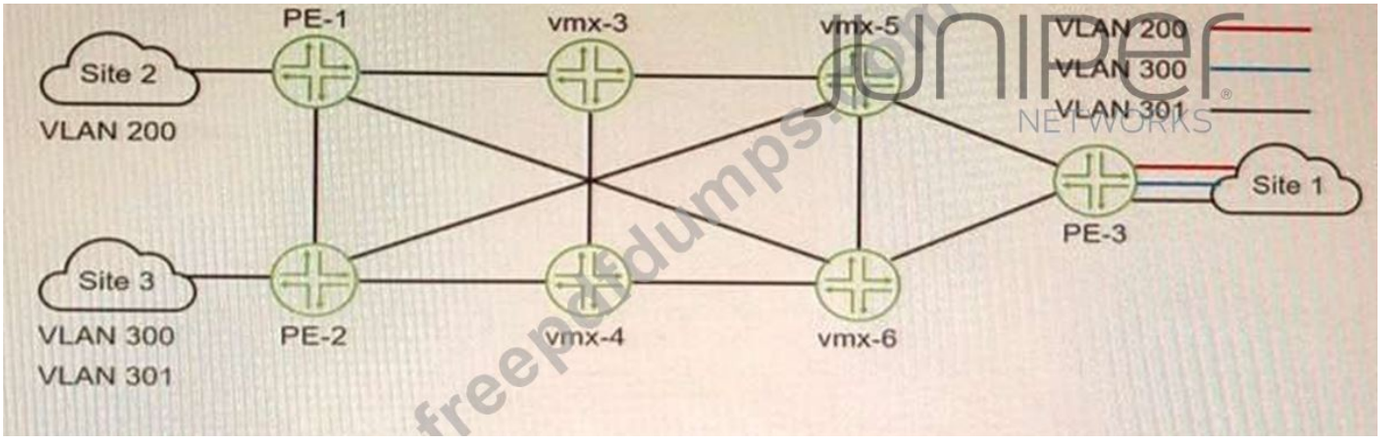

You want Site 1 to access three VLANs that are located in Site 2 and Site 3 The customer-facing interface on the PE-1 router is configured for Ethernet-VLAN encapsulation.
What is the minimum number of L2VPN routing instances to be configured to accomplish this task?
A. 1
B. 3
C. 2
D. 6
Answer: B (LEAVE 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: 3
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.

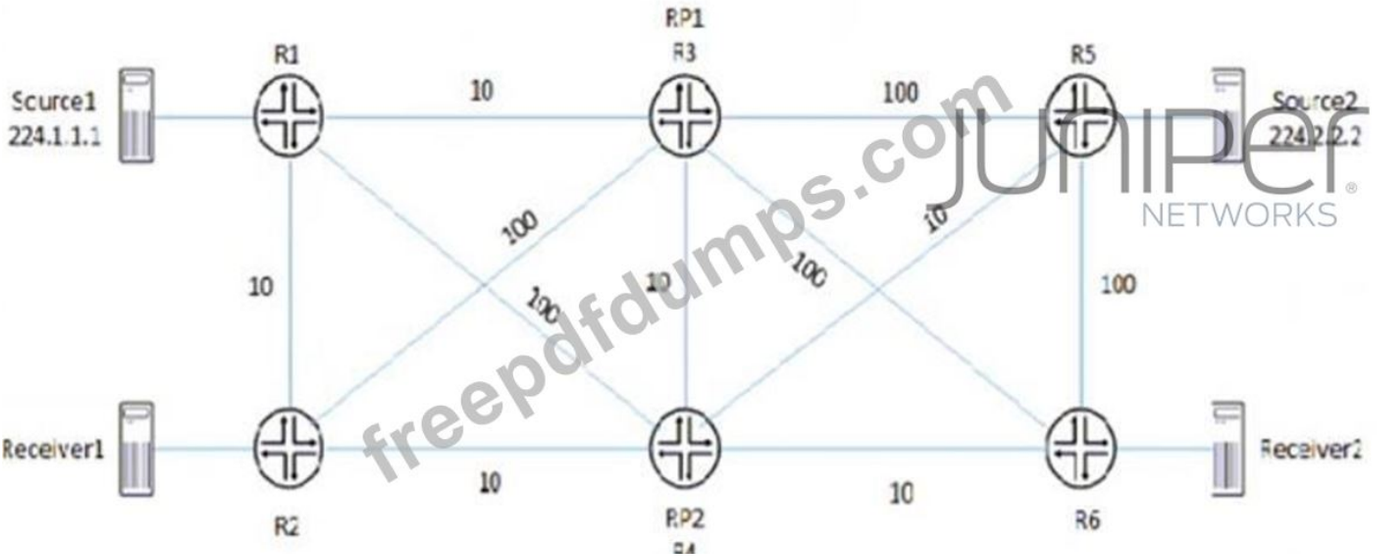

In this scenario, which two statements are correct about which RP is used? (Choose two.)
A. Source1 will use RP1 and Receiver1 will use RP1 for group 224.1.1.1.
B. Source1 will use RP1 and Receiver1 will use RP2 for group 224.1.1.1.
C. Source2 will use RP1 and Receiver2 will use RP1 for group 224.2.2.2.
D. Source2 will use RP2 and Receiver2 will use RP2 for group 224.2.2.2.
Answer: A,C (LEAVE A REPLY)

NEW QUESTION: 4
Which two statements describe PIM-SM? (Choose two)
A. Routers with receivers send join messages to their upstream neighbors.
B. Routers without receivers must periodically prune themselves from the SPT.
C. Traffic is initially flooded to all routers and an S,G is maintained for each group
D. Traffic is only forwarded to routers that request to join the distribution tree.
Answer: (SHOW ANSWER)
PIM sparse mode (PIM-SM) is a multicast routing protocol that uses a pull model to deliver multicast traffic.
In PIM-SM, routers with receivers send join messages to their upstream neighbors toward a rendezvous point (RP) or a source-specific tree (SPT). The RP or SPT acts as the root of a shared distribution tree for a multicast group. Traffic is only forwarded to routers that request to join the distribution tree by sending join messages.
PIM-SM does not flood traffic to all routers or prune routers without receivers, as PIM dense mode does.

NEW QUESTION: 5

Exhibit

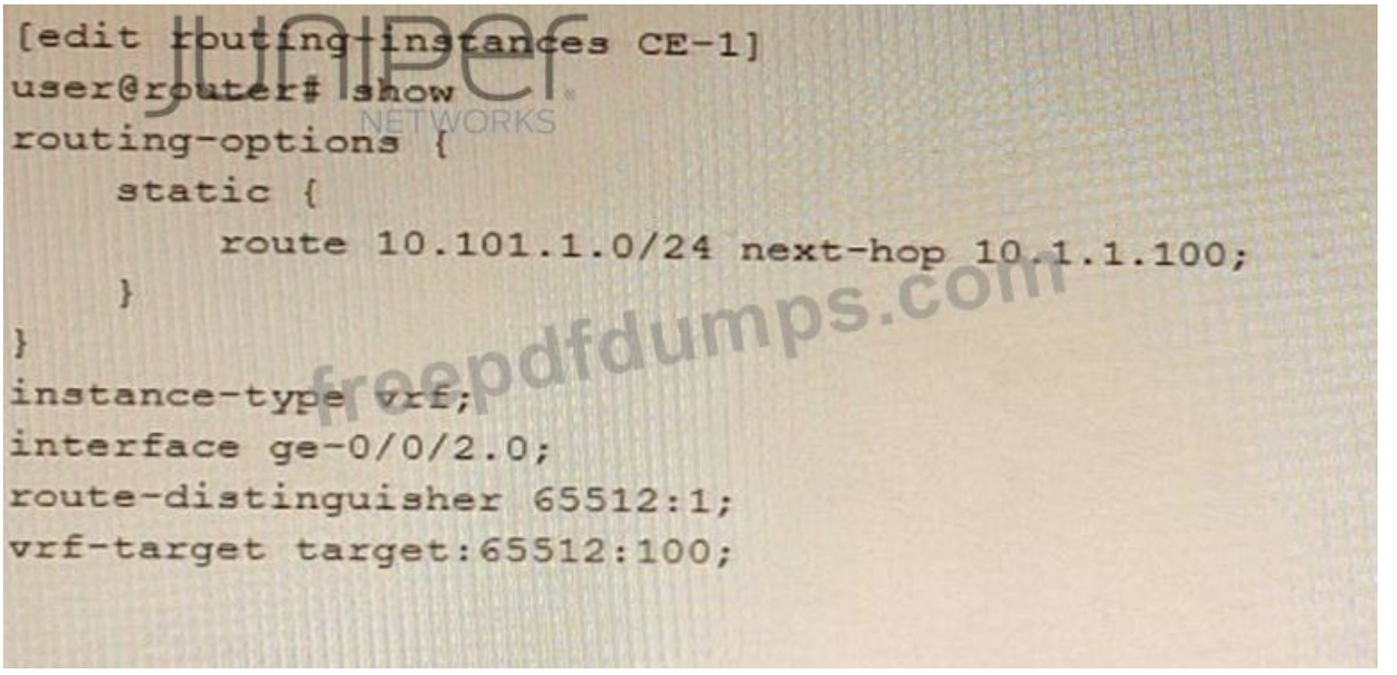

Referring to the exhibit, which statement is true?

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: (SHOW ANSWER)

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: 6

Exhibit

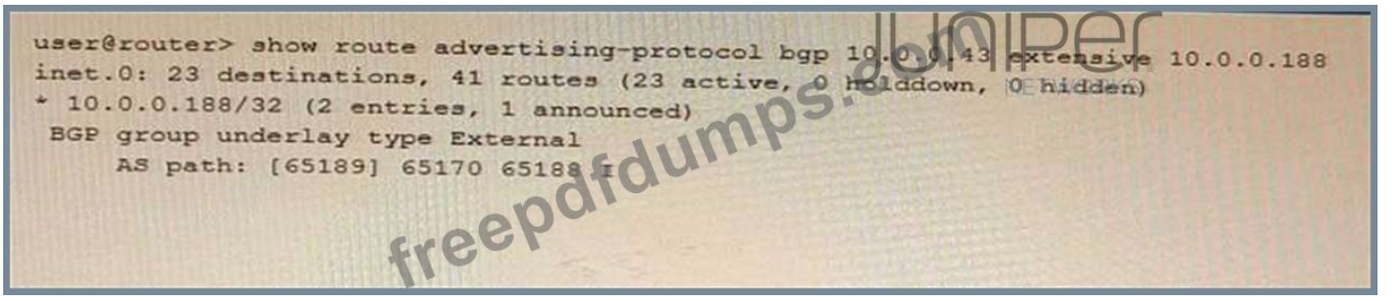

Referring to the exhibit, what do the brackets [ ] in the AS path identify?
A. They identify the local AS number associated with the AS path if configured on the router, or if AS path prepending is configured
B. They identify an AS set, which are groups of AS numbers in which the order does not matter
C. They identify that the autonomous system number is incomplete and awaiting more information from the BGP protocol.
D. They identify that a BGP confederation is being used to ensure that there are no routing loops.
Answer: B (LEAVE A REPLY)
The brackets [ ] in the AS path identify an AS set, which are groups of AS numbers in which the order does not matter. An AS set is used when BGP aggregates routes from different ASs into a single prefix. For example, if BGP aggregates routes 10.0.0.0/16 and 10.1.0.0/16 from AS 100 and AS 200, respectively, into a single prefix 10.0.0.0/15, then the AS path for this prefix will be [100 200]. An AS set reduces the length of the AS path and prevents routing loops.

NEW QUESTION: 7
Which two statements are correct about VPLS tunnels? (Choose two.)
A. BGP-signaled VPLS tunnels can use either RSVP or LDP between the PE routers.
B. BGP-signaled VPLS tunnels require manual provisioning of sites.
C. LDP-signaled VPLS tunnels use auto-discovery to provision sites.
D. LDP-signaled VPLS tunnels only support control bit 0.
Answer: A,D (LEAVE A REPLY)

NEW QUESTION: 8
Exhibit
user@RI show configuration interpolated-profile { interpolate {
fill-level [50 75 drop-probability [> }
class-of-service drop-profiles
];
20 60];
Which two statements are correct about the class-of-service configuration shown in the exhibit? (Choose two.)
A. The drop probability jumps immediately from 20% to 60% when the queue level reaches 75% full.
B. The drop probability gradually increases from 20% to 60% as the queue level increases from 50% full to
75% full
C. To use this drop profile, you reference it in a scheduler.
D. To use this drop profile, you apply it directly to an interface.
Answer: B,C (LEAVE A REPLY)
Explanation

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

NEW QUESTION: 9
Exhibit

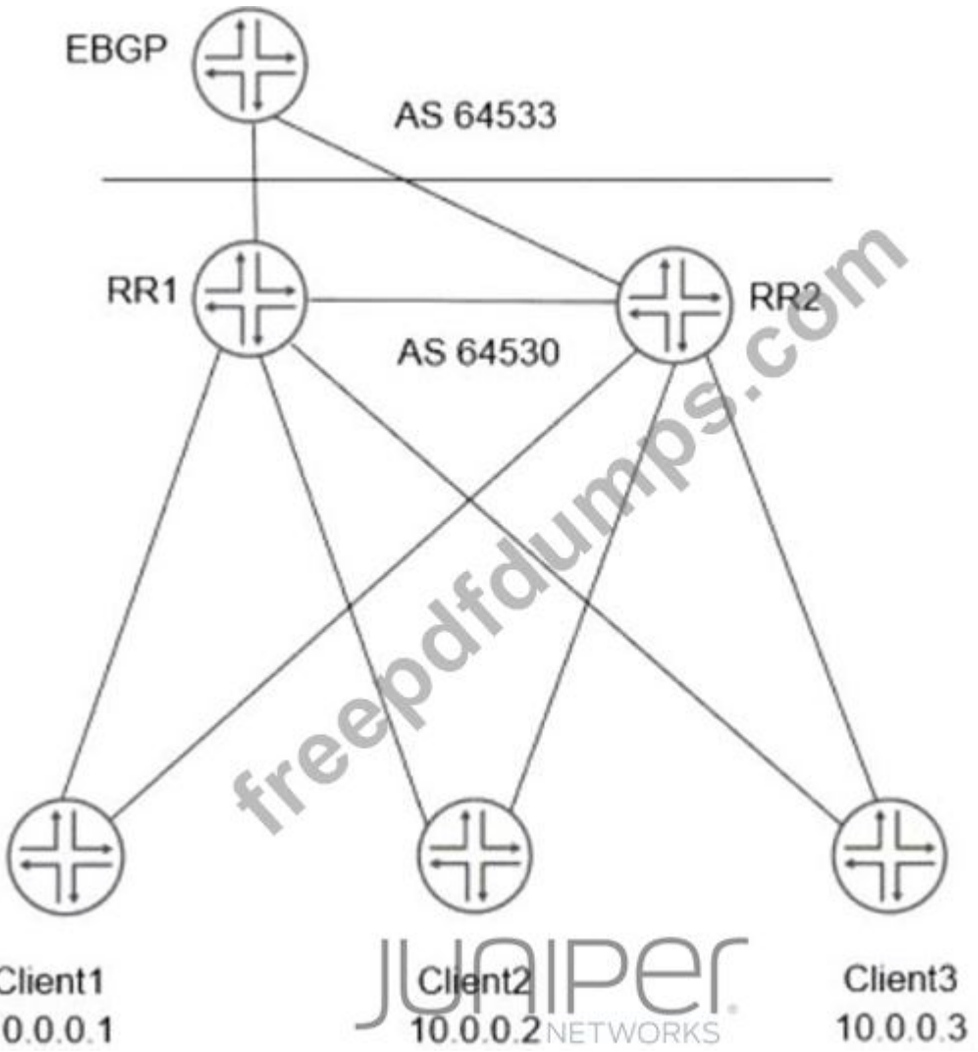

Referring to the exhibit, which two statements are correct about the dual route reflectors within a cluster?

(Choose two.)

A. RR1 and RR2 advertise routes learned from the clients to EBGP peers, using itself as the next hop.

B. RR1 and RR2 append the duster ID when advertising routes from dient to dient.

C. RR1 advertises routes from the client to RR2. using itself as the next hop.

D. RR1 and RR2 must have the same duster ID to exchange routes learned from the client.

Answer: (SHOW ANSWER)

NEW QUESTION: 10

Exhibit

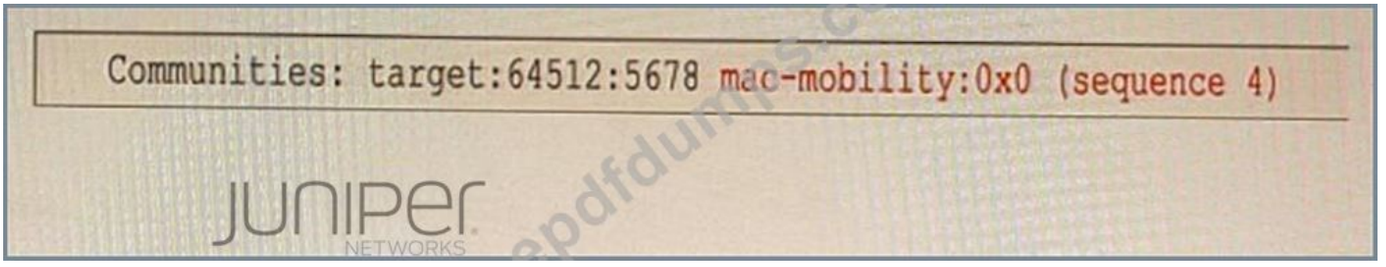

You have MAC addresses moving in your EVPN environment

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: 11

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: 12
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: (SHOW ANSWER)
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: 13
Exhibit

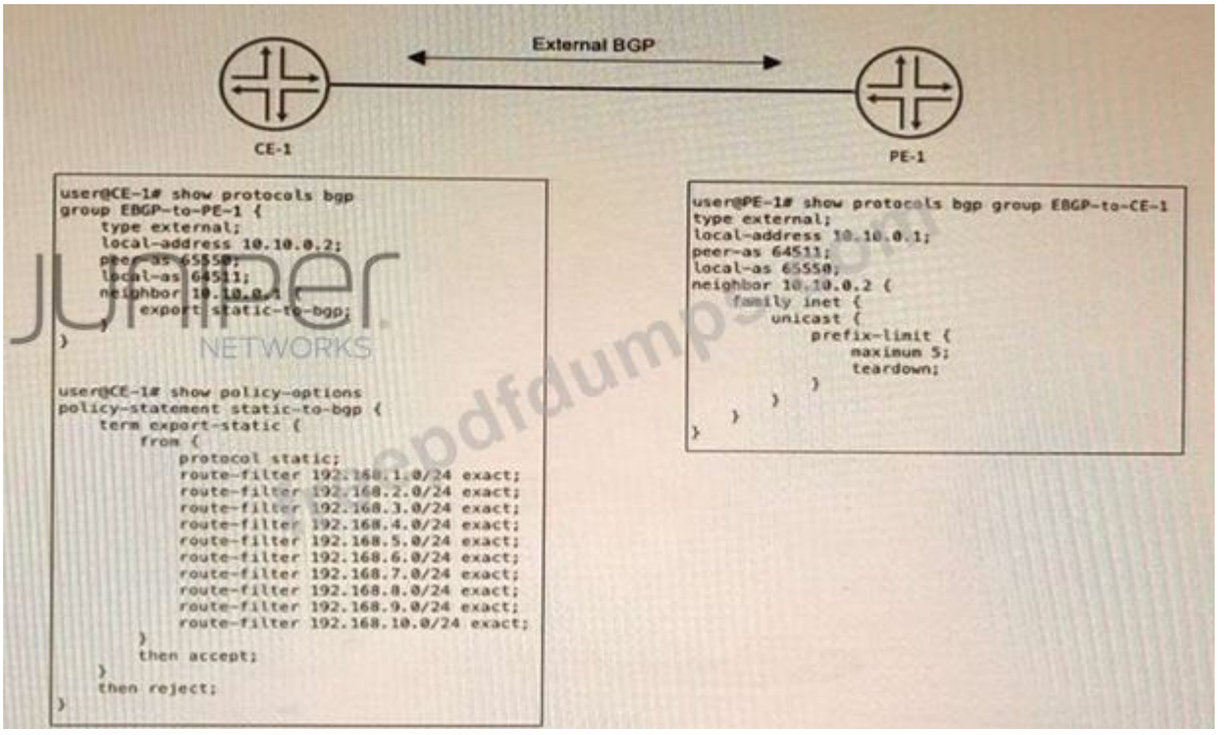

CE-1 must advertise ten subnets to PE-1 using BGP Once CE-1 starts advertising the subnets to PE-1, the BGP peering state changes to Active.

Referring to the CLI output shown in the exhibit, which statement is correct?

A. CE-1 is advertising its entire routing table.

B. CE-1 is configured with an incorrect peer AS

C. The prefix limit has been reached on PE-1

D. CE-1 is unreachable

Answer: (SHOW ANSWER)

Explanation

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

NEW QUESTION: 14

Exhibit

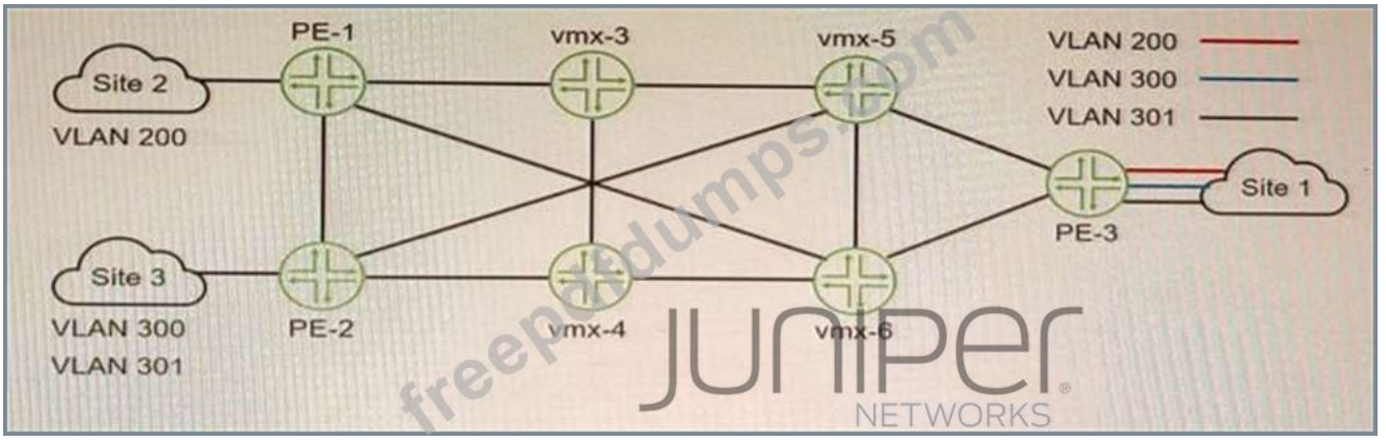

You want Site 1 to access three VLANs that are located in Site 2 and Site 3 The customer-facing interface on the PE-1 router is configured for Ethernet-VLAN encapsulation.

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)

Explanation

To allow Site 1 to access three VLANs that are located in Site 2 and Site 3, you need to configure three L2VPN routing instances on PE-1, one for each VLAN. Each L2VPN routing instance will have a different VLAN ID and a different VNI for VXLAN encapsulation. Each L2VPN routing instance will also have a different vrf-target export value to identify which VPN routes belong to which VLAN. This way, PE-1 can forward traffic from Site 1 to Site 2 and Site 3 based on the VLAN tags and VNIs.

NEW QUESTION: 15

You have an L2VPN connecting two CEs across a provider network that runs OSPF. You have OSPF configured on both CEs.

Which two statements are correct in this scenario? (Choose two.)

A. The CE and PE OSPF areas must match.

B. OSPF neighborship is formed between the CEs and PEs.

C. The CE and PE OSPF areas can be different.

D. OSPF neighborship is formed between the two CEs.

Answer: (SHOW ANSWER)

NEW QUESTION: 16

Exhibit.

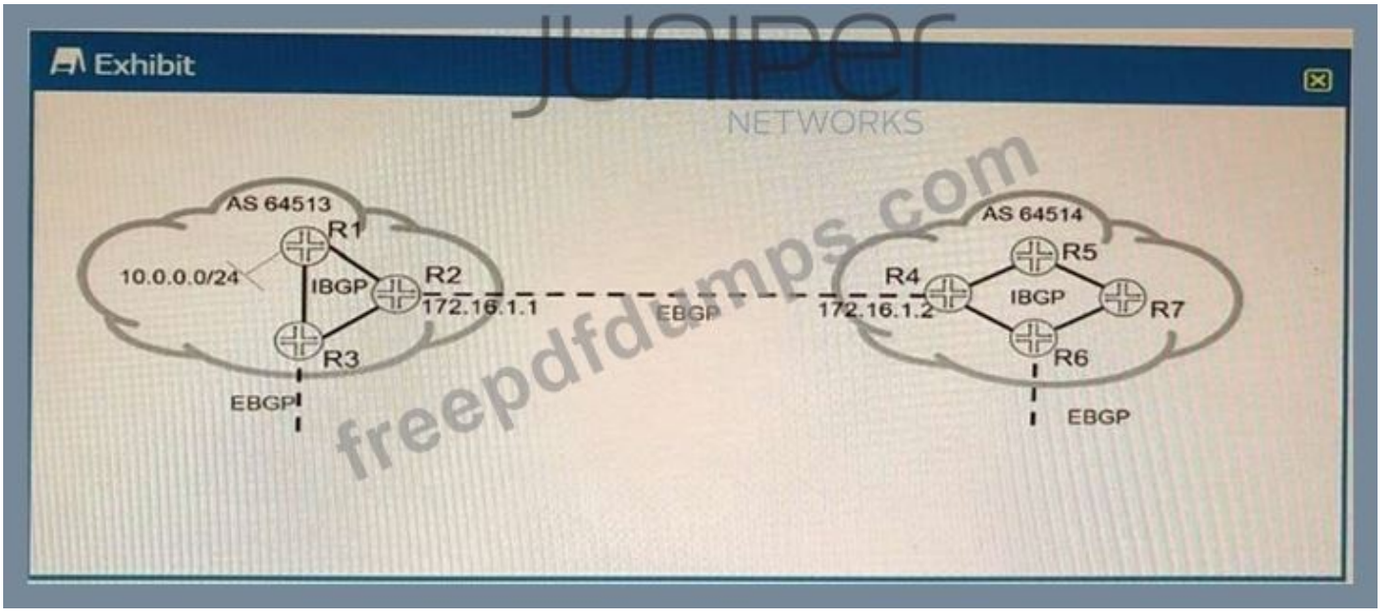

Referring to the exhibit; the 10.0.0.0/24 EBGP route is received on R5; however, the route is being hidden.

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: 17
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 (LEAVE A REPLY)
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
Which origin code is preferred by BGP?
A. Internal
B. External
C. Incomplete
D. Null
Answer: C (LEAVE A REPLY)
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: 19
Which two statements about the configuration shown in the exhibit are correct? (Choose two.)

```
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;
                }
            }
        }
    }
}
```

A. This VPN connects customer sites that use different AS numbers.

B. A Layer 2 VPN is configured.

C. This VPN connects customer sites that use the same AS number.

D. A Layer 3 VPN is configured.

Answer: C,D (LEAVE A REPLY)

NEW QUESTION: 20

Exhibit

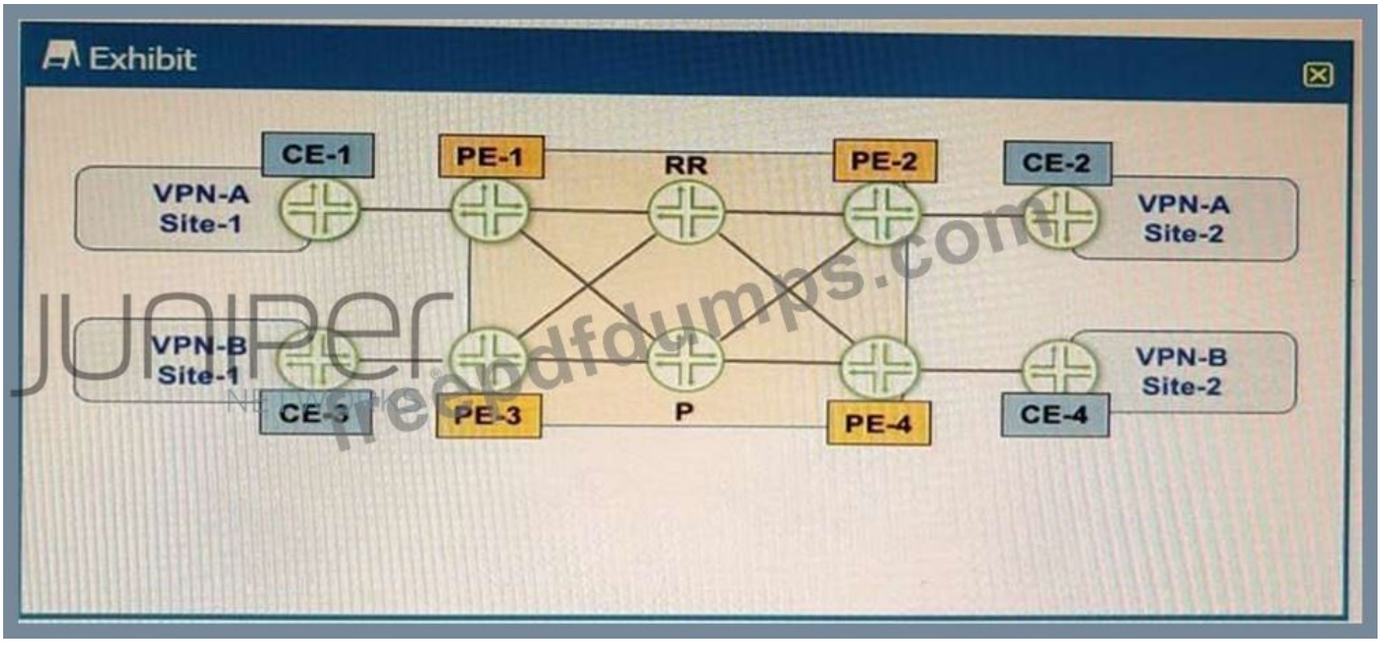

Referring to the exhibit, PE-1 and PE-2 are getting route updates for VPN-B when neither of them service that VPN Which two actions would optimize this process? (Choose two.)

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 inet. 0 Statement on the PEs.
D. Configure the resolution rib bgp.l3vpn.O resolution-ribs inet. 0 Statement on the RR
Answer: B,D (LEAVE A REPLY)
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: 21
Exhibit

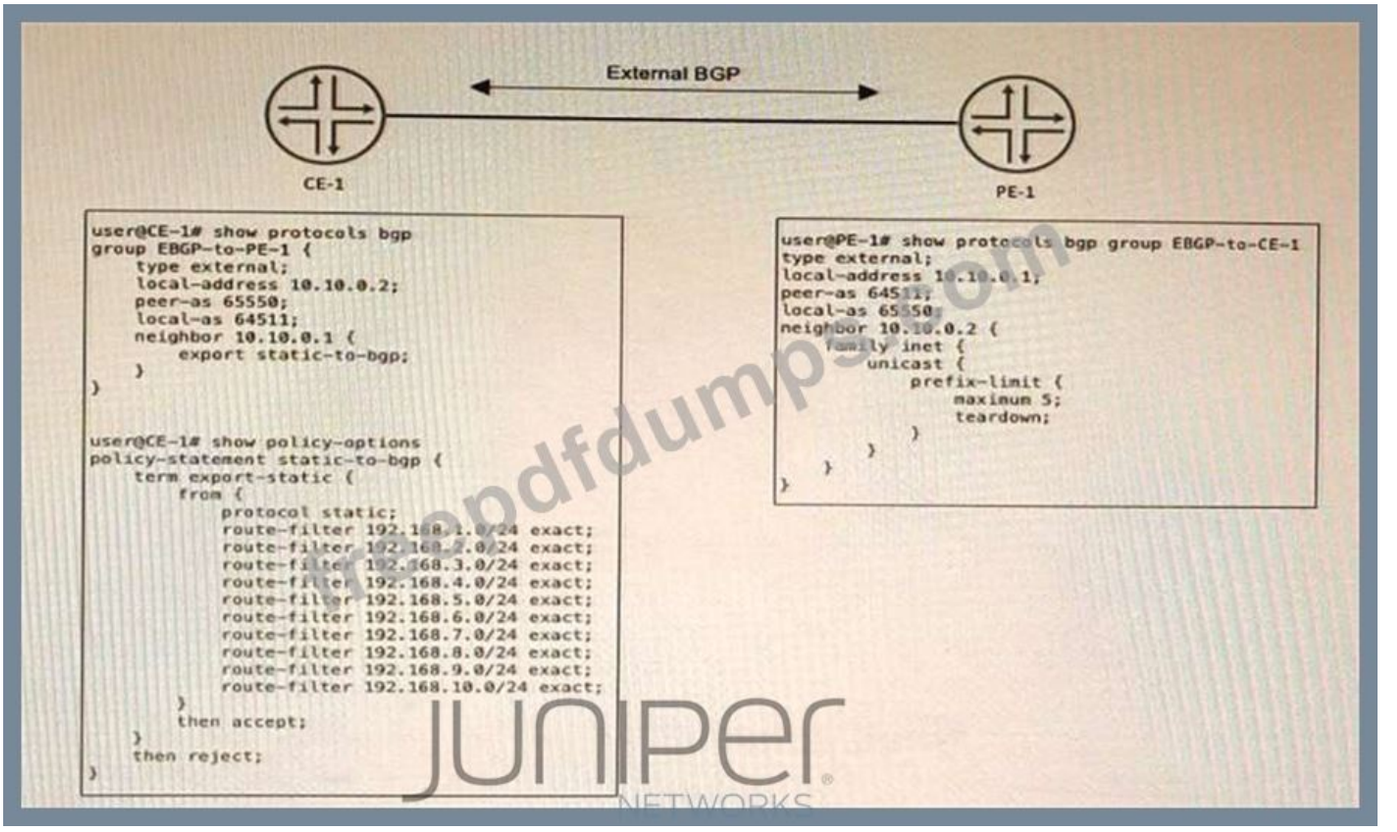

CE-1 must advertise ten subnets to PE-1 using BGP Once CE-1 starts advertising the subnets to PE-1, the BGP peering state changes to Active.

Referring to the CLI output shown in the exhibit, which statement is correct?

A. CE-1 is advertising its entire routing table.

B. CE-1 is configured with an incorrect peer AS

C. The prefix limit has been reached on PE-1

D. CE-1 is unreachable

Answer: B (LEAVE A REPLY)

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

NEW QUESTION: 22

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)

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: 23

In IS-IS, which two statements are correct about the designated intermediate system (DIS) on a multi-access network segment? (Choose two)

A. A router with a priority of 10 wins the DIS election over a router with a priority of 1.

B. A router with a priority of 1 wins the DIS election over a router with a priority of 10.

C. On the multi-access network, each router forms an adjacency to every other router on the segment

D. On the multi-access network, each router only forms an adjacency to the DIS.

Answer: A,D (LEAVE A REPLY)

Explanation

In IS-IS, a designated intermediate system (DIS) is a router that is elected on a multi-access network segment (such as Ethernet) to perform some functions on behalf of other routers on the same segment. A DIS is responsible for sending network link-state advertisements (LSPs), which describe all the routers attached to the network. These LSPs are flooded throughout a single area. A DIS also generates pseudonode LSPs, which represent the multi-access network as a single node in the link-state database. A DIS election is based on the priority value configured on each router's interface connected to the multi-access network. The priority value ranges from 0 to

127, with higher values indicating higher priority. The router with the highest priority becomes the DIS for the area (Level 1, Level 2, or both). If routers have the same priority, then the router with the highest MAC address is elected as the DIS. By default, routers have a priority value of 64. On a multi-access network, each router only forms an adjacency to the DIS, not to every other router on the segment. This reduces the amount of hello packets and LSP

NEW QUESTION: 24

Exhibit

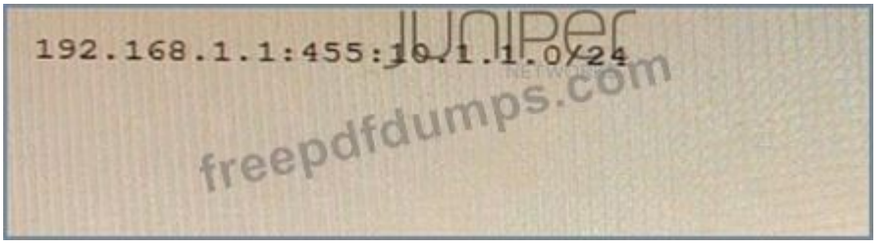

You are examining an L3VPN route that includes the information shown in the exhibit Which statement is correct in this scenario?

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: A (LEAVE A REPLY)

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: 25

Which origin code is preferred by BGP?

A. Internal

B. External

C. Incomplete

D. Null

Answer: (SHOW ANSWER)

Prefer the route with the lower origin code.

Routes learned from an IGP have a lower origin code than those learned from an exterior gateway protocol (EGP), and both have lower origin codes than incomplete routes (routes whose origin is unknown).

https://www.juniper.net/documentation/us/en/software/junos/vpn-l2/bgp/topics/concept/routing-protocols-address-representation.html

NEW QUESTION: 26

Which origin code is preferred by BGP?

A. Internal

B. External

C. Incomplete

D. Null

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: 27

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: 28
Which two statements about the output shown in the exhibit are correct? (Choose two.)

A. The PE router has the capability to pop flow labels.
B. There has been a VLAN ID mismatch.
C. The PE is attached to a single local site.

D. The connection has not flapped since it was initiated.

Answer: C,D (LEAVE A REPLY)

NEW QUESTION: 29

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

What should you do to solve the problem?

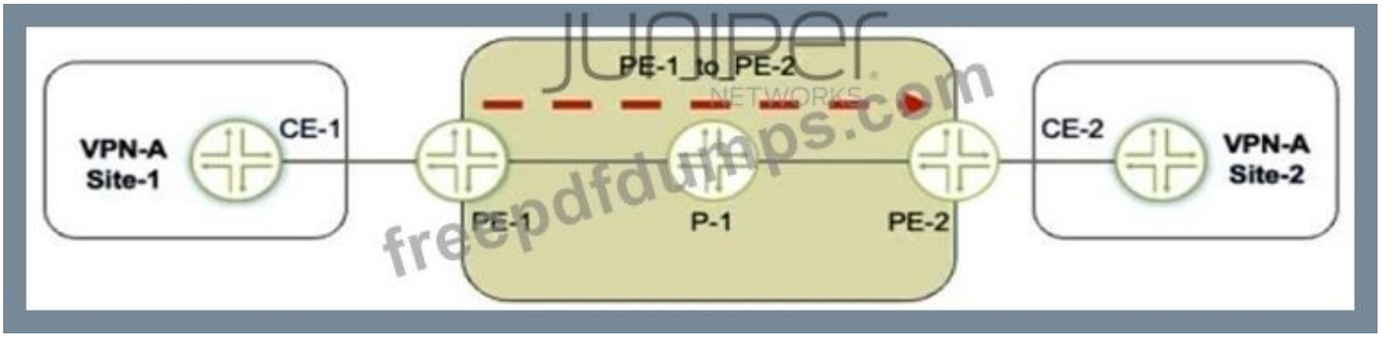

A. Set a static CoS value for the PE-1_to_PE-2 LSP.

B. Configure the explicit-null statement on PE-1.

C. Configure VPN prefix mapping for the PE-1_to_PE-2 LSP.

D. Configure the explicit-null statement on PE-2.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 30

Based on the configuration contents shown in the exhibit, which statement is true?

```
[edit policy-options]
user@router# show
policy-statement block-igmp {
    term 1 {
        from {
            route-filter 224.7.7.7/32 exact;
            source-address-filter 192.168.100.10/32 exact;
        }
        then reject;
    }
}
[edit protocols igmp]
user@router# show
interface ge-0/0/0.0 {
    group-policy block-igmp;
    group-limit 25;
}
```

A. Joins for group 224.7.7.7 are rejected if the source address is 192.168.100.10.

B. Joins for any group are accepted if the group count value is less than 25.

C. Joins for group 224.7.7.7 are accepted if the group count is less than 25.

D. Joins for group 224.7.7.7 are always rejected, regardless of the group count.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 31
Exhibit

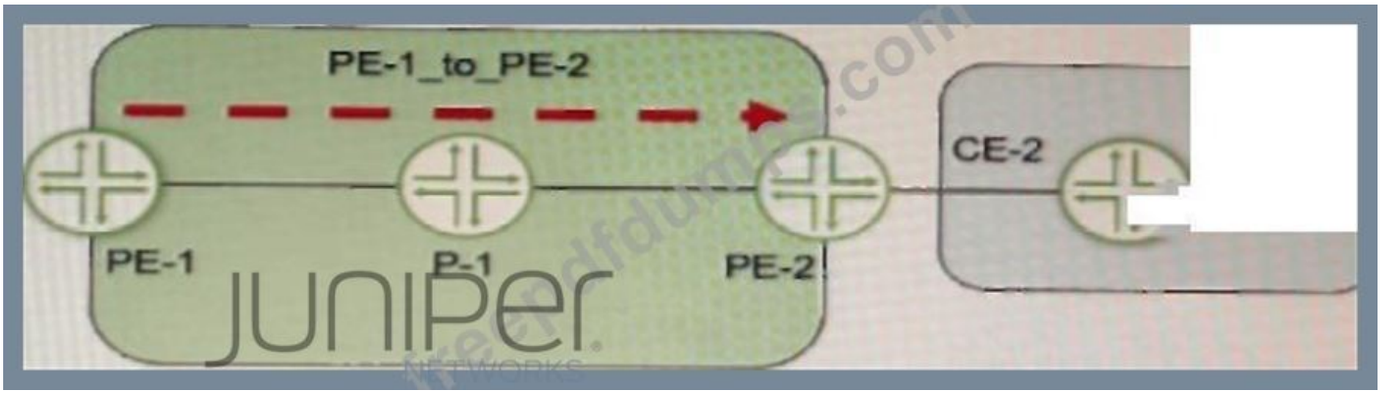

Referring to the exhibit, a working L3VPN exists that connects VPN-A sites CoS is configured correctly to match on the MPLS EXP bits of the LSP, but when traffic is sent from Site-1 to Site-2, PE-2 is not classifying the traffic correctly What should you do to solve the problem?
A. Configure the explicit-null statement on PE-1.
B. Configure the explicit-null statement on PE-2
C. Configure VPN prefix mapping for the PE-1_to_PE-2 LSP
D. Set a static CoS value for the PE-1_to_PE-2 LSP
Answer: A (LEAVE A REPLY)
Explanation
The explicit-null statement enables the PE router to send an MPLS label with a value of 0 (explicit null) instead of an IP header for packets destined to the VPN customer sites. This allows the penultimate hop router (the router before the egress PE router) to preserve the EXP bits of the MPLS label and pass them to the egress PE router. The egress PE router can then use these EXP bits to classify the traffic according to the CoS policy2
. In this example, PE-1 should configure the explicit-null statement under [edit protocols mpls label-switched-path PE-1_to_PE-2] hierarchy level.

NEW QUESTION: 32

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

NEW QUESTION: 33
Exhibit

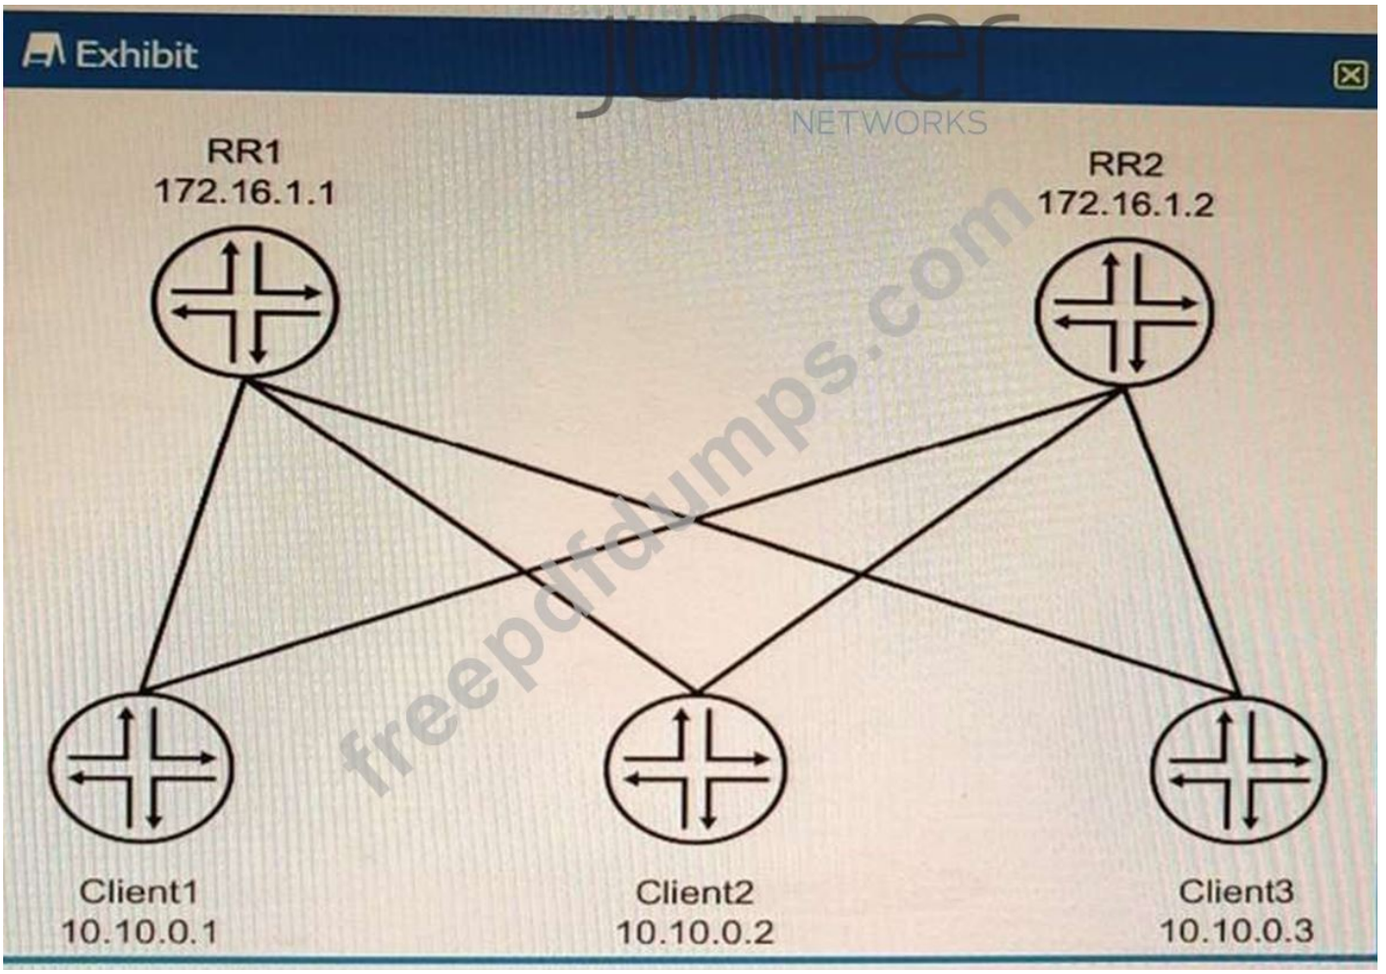

The environment is using BGP All devices are in the same AS with reachability redundancy
Referring to the exhibit, which statement is correct?

A. RR1 is peered to Client2 and RR2

B. RR2 is in an OpenConfirm State until RR1 becomes unreachable.

C. Client1 is peered to Client2 and Client3.

D. Peering is dynamically discovered between all devices.

Answer: A (LEAVE A REPLY)

Explanation

BGP route reflectors are BGP routers that are allowed to ignore the IBGP loop avoidance rule and advertise IBGP learned routes to other IBGP peers under specific conditions. BGP route reflectors can reduce the number of IBGP sessions and updates in a network by eliminating the need for a full mesh of IBGP peers.

BGP route reflectors can have three types of peerings:

* EBGP neighbor: A BGP router that belongs to a different autonomous system (AS) than the route reflector.

* IBGP client neighbor: An IBGP router that receives reflected routes from the route reflector. A client does not need to peer with other clients or non-clients.

* IBGP non-client neighbor: An IBGP router that does not receive reflected routes from the route reflector. A non-client needs to peer with other non-clients and the route reflector.

In the exhibit, we can see that RR1 and RR2 are route reflectors in the same AS with reachability redundancy.

They have two types of peerings: EBGP neighbors (R1 and R4) and IBGP client neighbors (Client1, Client2, and Client3). RR1 and RR2 are also peered with each other as IBGP non-client neighbors.

NEW QUESTION: 34

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

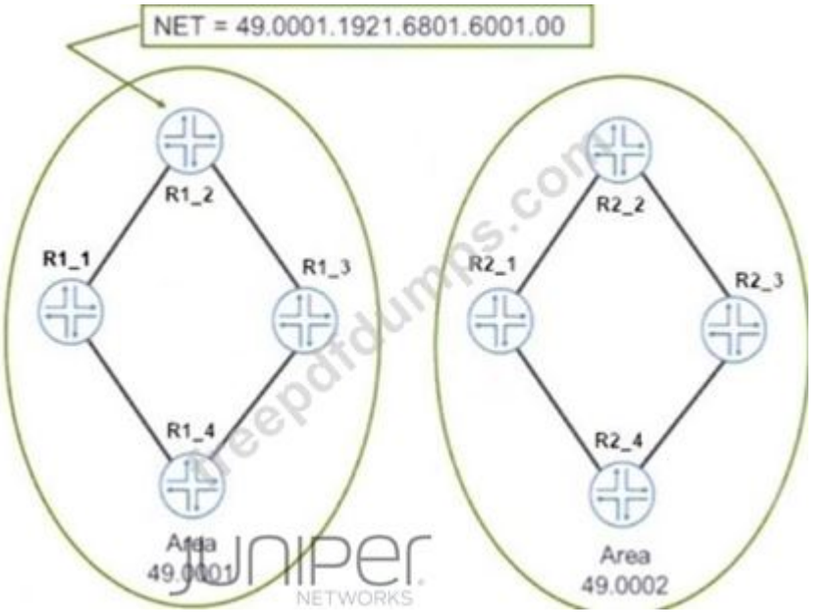

Which statement is correct in this scenario?

A. The routers are using unnumbered interfaces.
B. The system ID of R1_2 is 192.168.16.1.
C. The NSEL byte for Area 0001 is 00.
D. The area address is two bytes.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 35
Exhibit

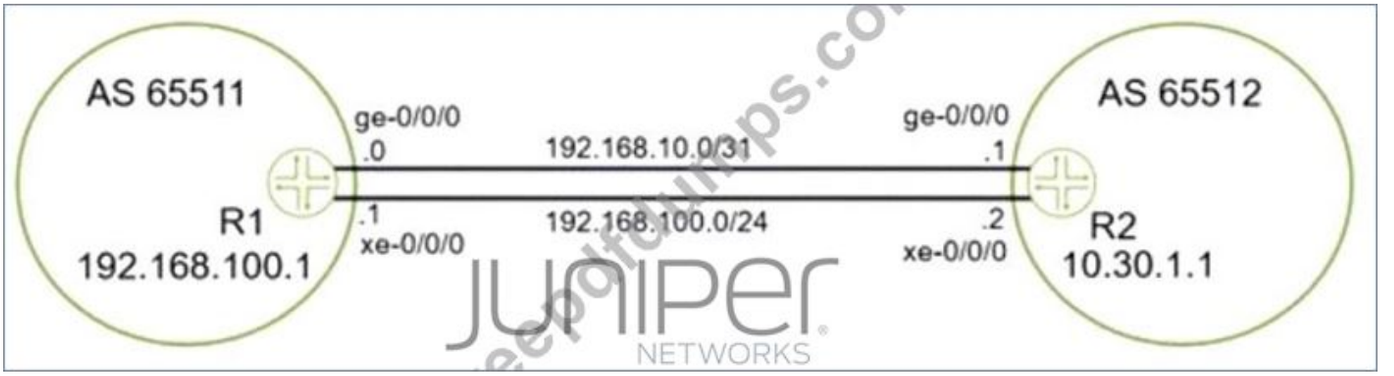

You want to use both links between R1 and R2 Because of the bandwidth difference between the two links, you must ensure that the links are used as much as possible.
Which action will accomplish this goal?
A. Define a policy to tag routes with the appropriate bandwidth community.
B. Disable multipath.
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: A (LEAVE A REPLY)
https://www.juniper.net/documentation/us/en/software/junos/sampling-forwarding-monitoring/bgp/topics/concep

NEW QUESTION: 36
Exhibit

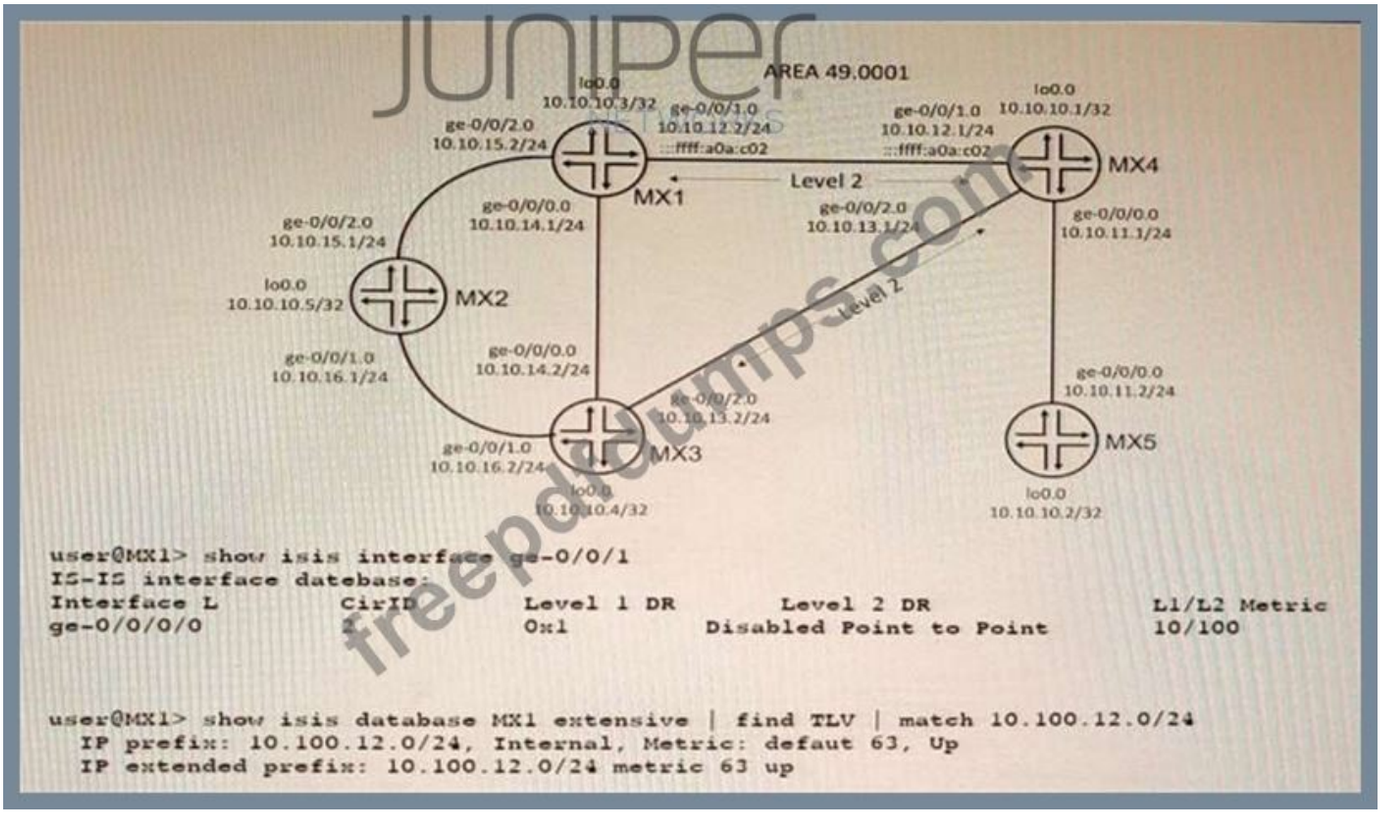

A network is using IS-IS for routing.

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 (LEAVE A REPLY)

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: 37

You are using a Layer 3 VPN to connect two customer sites. The VPN routes for the customer networks appear as hidden in the bgp.l3vpn.0 routing table on the PE routers.

What is causing this problem?

A. Route targets are not configured.

B. There is a routing loop in the service provider backbone.

C. The routes use overlapping IP addresses.

D. There is not an established MPLS LSP between the two PE routers.

Answer: (SHOW ANSWER)

NEW QUESTION: 38
A router running IS-IS is configured with an ISO address of 49.0001.00a0.c96b.c490.00.
Which part of this address is the system ID?
A. 00a0.c96b.c490 is the system identifier.
B. c96b.c490 is the system identifier.
C. 0001.00a0.c96b.c490 is the system identifier.
D. c490 is the system identifier.
Answer: (SHOW ANSWER)

NEW QUESTION: 39
Which statement is correct about IS-IS when it performs the Dijkstra algorithm?
A. The algorithm will stop processing once the tree database is empty.
B. When a new neighbor ID in the tree database matches a router ID in the LSDthe neighbor ID is moved to the candidate database.
C. The local router moves its own local tuples into the candidate database.
D. Tuples with the lowest cost are moved from the tree database to the LSDB.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 40
Exhibit

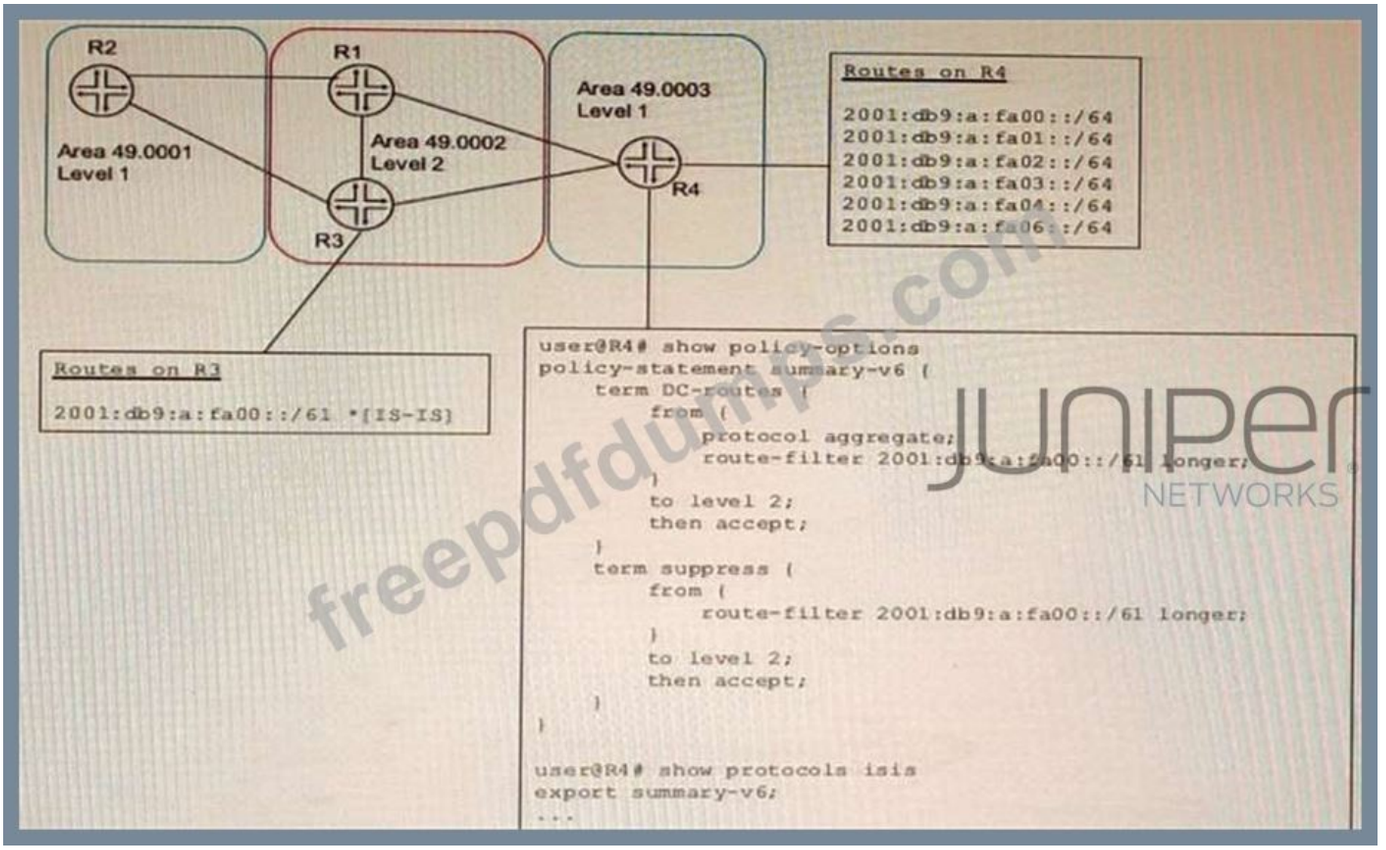

A network designer would like to create a summary route as shown in the exhibit, but the configuration is not working.
Which three configuration changes will create a summary route? (Choose three.)
A. set policy-options policy-statement leak-v6 term DC-routes then reject
B. delete policy-options policy-statement leak-v6 term DC-routes from route-filter 2001: db9 :a: fa00 : :/6l longer
C. set policy-options policy-statement leak-v6 term DC-routes from route-filter 2001:db9:a:faOO::/61 exact
D. delete protocols isis export summary-v6
E. set protocols isis import summary-v6
Answer: B,C,D (LEAVE A REPLY)
Explanation
To create a summary route for IS-IS, you need to configure a policy statement that matches the prefixes to be summarized and sets the next-hop to discard. You also need to configure a summary-address statement under the IS-IS protocol hierarchy that references the policy statement. In this case, the policy statement leak-v6 is trying to match the prefix 2001:db9:a:fa00::/61 exactly, but this prefix is not advertised by any router in the network. Therefore, no summary route is created. To fix this, you need to delete the longer keyword from the route-filter term and change the prefix length to /61 exact. This will match any prefix that falls within the /61 range. You also need to delete the export statement under protocols isis, because this will export all routes that match the policy statement to other IS-IS routers, which is not desired for a summary route.

NEW QUESTION: 41
CE-1 must advertise ten subnets to PE-1 using BGP. Once CE-1 starts advertising the subnets to PE-1, the BGP peering state changes to Active.
Referring to the CLI output shown in the exhibit, which statement is correct?

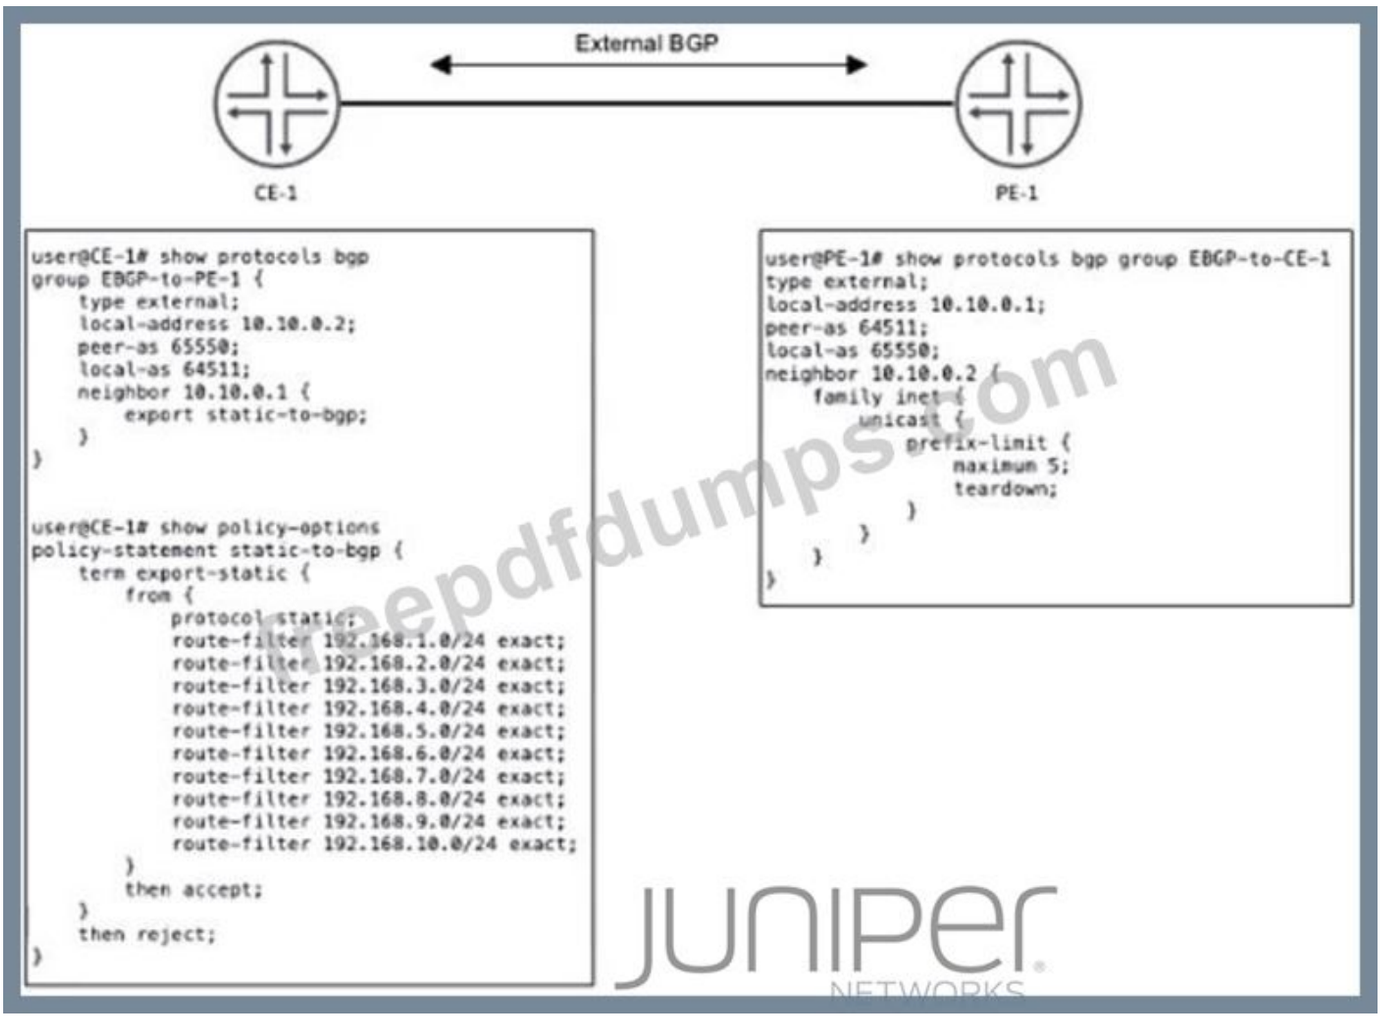

A. The prefix limit has been reached on PE-1.

B. CE-1 is configured with an incorrect peer AS.

C. CE-1 is advertising its entire routing table.

D. CE-1 is unreachable.

Answer: A (LEAVE A REPLY)

NEW QUESTION: 42

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

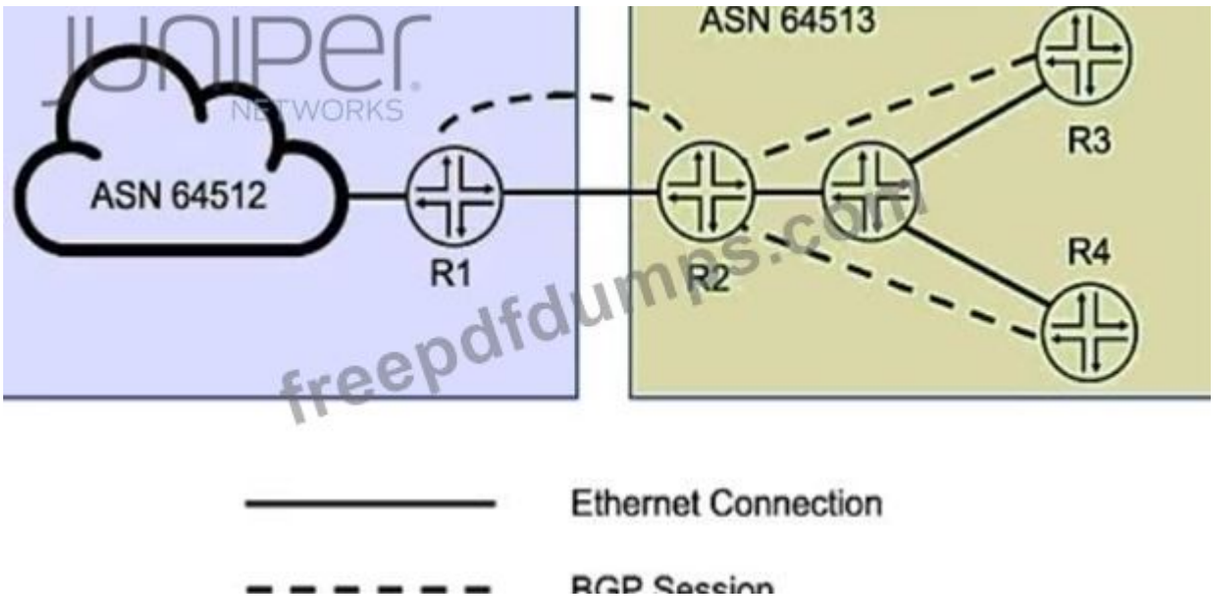

A. BGP GTSM requires a TTL of 1 to be configured between neighbors.
B. You can implement BGP GTSM between R2, R3, and R4.
C. You can implement BGP GTSM between R2 and R1.
D. BGP GTSM requires a firewall filter to discard packets with incorrect TTL.
E. BGP GTSM requires a TTL of 255 to be configured between neighbors
Answer: C,D,E (LEAVE A REPLY)
https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/ref/statement/multihop-edit-protocols-bgp.html

NEW QUESTION: 43
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?

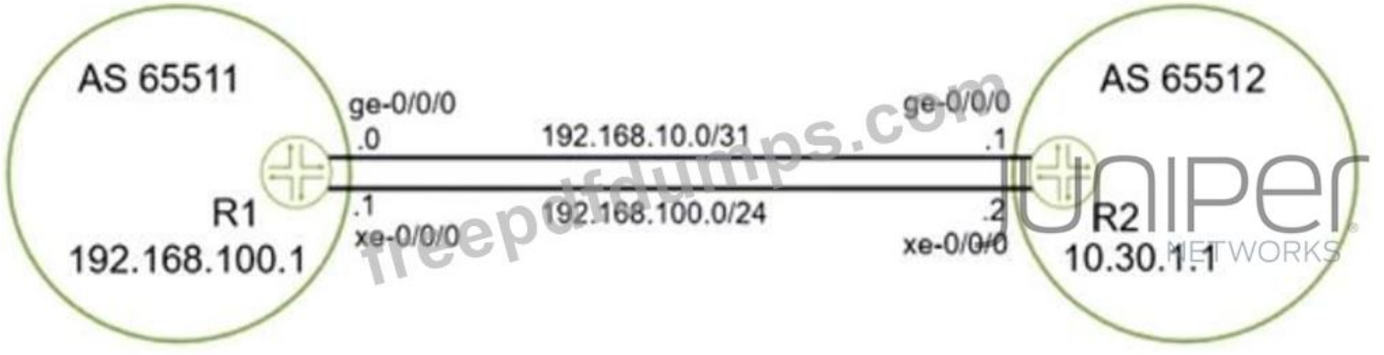

A. Ensure that the metric-out parameter on the Gigabit Ethernet interface is higher than the 10 Gigibit Ethernet interface.
B. Define a policy to tag routes with the appropriate bandwidth community.
C. Enable per-prefix load balancing.
D. Disable multipath.
Answer: B (LEAVE A REPLY)

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

NEW QUESTION: 44
Exhibit

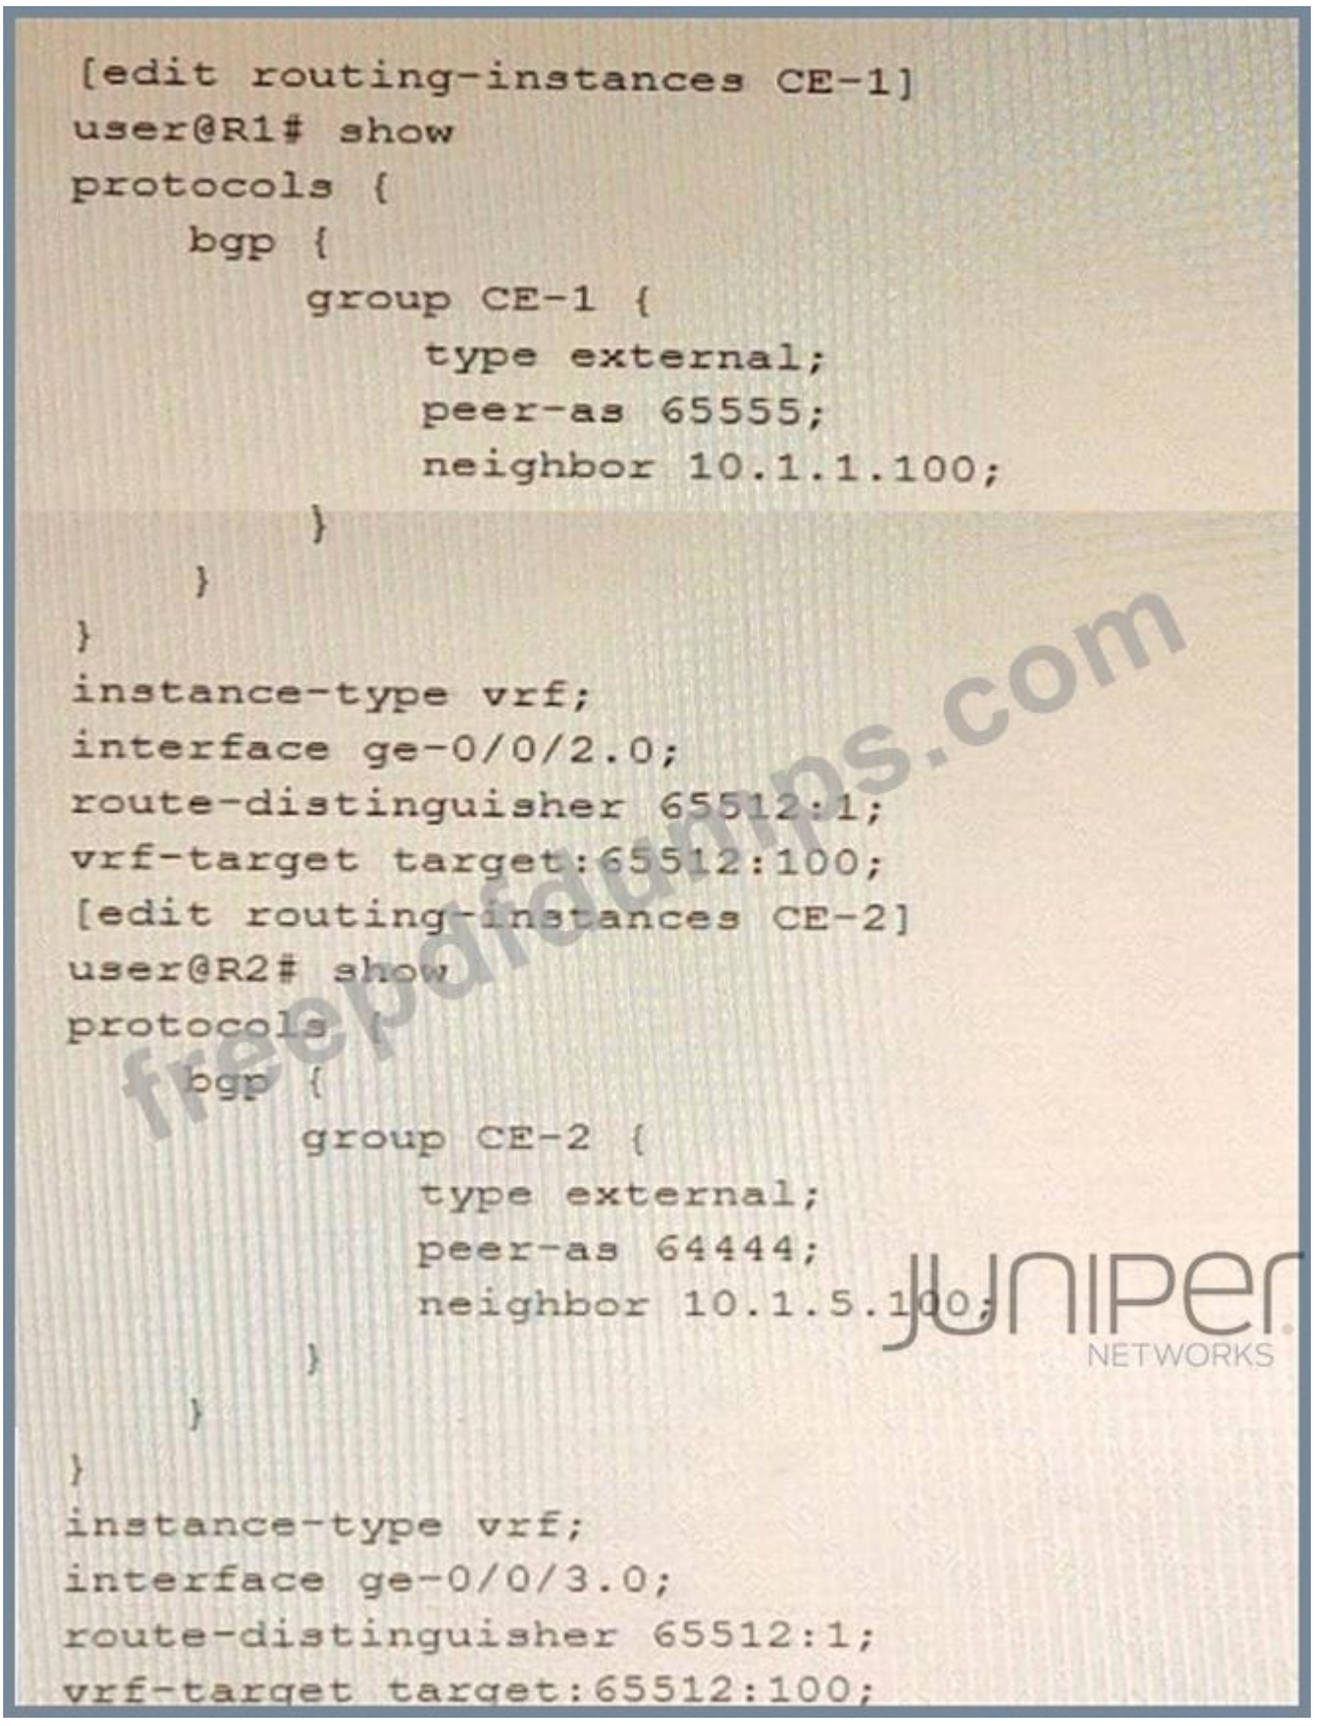

```
[edit routing-instances CE-1]
user@R1# show
protocols {
    bgp {
        group CE-1 {
            type external;
            peer-as 65555;
            neighbor 10.1.1.100;
        }
    }
}
instance-type vrf;
interface ge-0/0/2.0;
route-distinguisher 65512:1;
vrf-target target:65512:100;
[edit 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)
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: 45
Referring to the exhibit, which statement is true?

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

NEW QUESTION: 46
You are troubleshooting an issue for a customer site that uses 10.10.0.0/24 in AS 65224, but you see another AS in the AS path.
Referring to the exhibit, what is the cause of the problem?

```
net.0: 233 destinations, 233 routes (233 active, 0 holddown, 0 hidden)
 10.10.0.0/16 (1 entry, 1 announced)
    Accepted
    Nexthop: 10.16.40.1
    AS path: 65000 {65137 65224} I
    Aggregator: 65000 10.11.11.11
ser@router> show route 10.10.0.0/24
net.0: 233 destinations, 233 routes (233 active, 0 holddown, 0 hidden)
 = Active Route, - = Last Active, * = Both
0.10.0.0/24        *[BGP/170] 00:12:17, localpref 100
                      AS path: 65000 {65137 65224} I, validation-state: unverified
                    > to 10.16.40.1 via ge-0/0/2.0
```


A. AS 65000 is pre-pending AS 65137 to route advertisements.
B. The local AS is receiving two equal cost routes to 10.10.0.0/24.
C. AS 65137 is advertising the 10.10.0.0/24 prefix.
D. The local AS is in the process of withdrawing the route from AS 65137.
Answer: (SHOW ANSWER)

NEW QUESTION: 47
You have an OSPF environment. You have recently added a router called R4 that is directly connected to R1 and R2. You discover that R4 is only peering with R2.

```
root@R1> show ospf interface extensive
Interface           State   Area            DR ID           BDR ID          Nbrs
et-0/0/33.0         DR      0.0.0.0         192.168.252.0   0.0.0.0            0
  Type: LAN, Address: 192.168.254.0, Mask: 255.255.255.254, MTU: 9202, Cost: 1
  DR addr: 192.168.254.0, Priority: 128
  Adj count: 0
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
root@R4> show ospf interface extensive
Interface           State   Area            DR ID           BDR ID          Nbrs
et-0/0/48.0         Waiting 0.0.0.0         0.0.0.0         0.0.0.0            0
  Type: LAN, Address: 192.168.254.1, Mask: 255.255.255.254, MTU: 9202, Cost: 1
  Priority: 128
    Adj count: 0
  Hello: 5, Dead: 20, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
et-0/0/49.0         DR      0.0.0.0         192.168.253.0   192.168.252.1      1
  Type: LAN, Address: 192.168.254.9, Mask: 255.255.255.254, MTU: 9202, Cost: 1
  DR addr: 192.168.254.9, BDR addr: 192.168.254.8, Priority: 128
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
root@R2> show ospf interface et-0/0/33.0 extensive
Interface           State   Area            DR ID           BDR ID          Nbrs
et-0/0/33.0         BDR     0.0.0.0         192.168.253.0   192.168.252.1      1
  Type: LAN, Address: 192.168.254.8, Mask: 255.255.255.254, MTU: 9202, Cost: 1
  DR addr: 192.168.254.9, BDR addr: 192.168.254.8, Priority: 128
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```


Referring to the exhibit, how would you correct the peering?

A. Adjust the Hello Interval on R1 and R2 to match the Hello Interval on R4.

B. Change the MTU size on R1 and R2 to be 22 bytes higher than R4's MTU size.

C. Adjust the Priority on R1 to be lower than the Priority on R4.

D. Adjust the Dead Interval on R4 to match the Dead Interval on R1 and R2.

Answer: A (LEAVE A REPLY)

NEW QUESTION: 48

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: (SHOW ANSWER)
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: 49
Which statement is true regarding BGP FlowSpec?
A. It is used to protect a network from denial-of-service attacks dynamically.
B. It verifies that the source IP of the incoming packet has a resolvable route in the routing table.
C. It uses a remote triggered black hole to protect a network from a denial-of-service attack.
D. It uses dynamically created routing policies to protect a network from denial-of-service attacks.
Answer: (SHOW ANSWER)

NEW QUESTION: 50
Exhibit

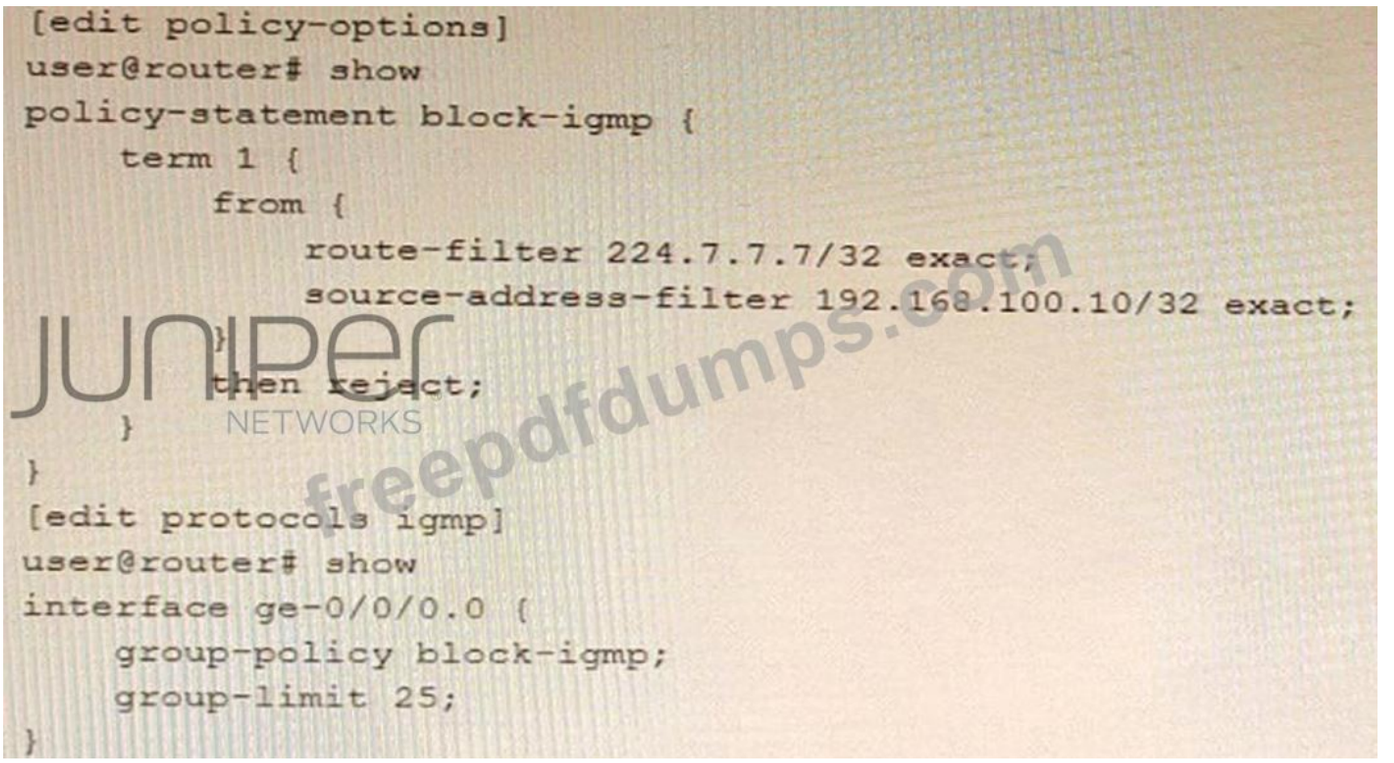

Based on the configuration contents shown in the exhibit, which statement is true?

A. Joins for group 224.7.7.7 are rejected if the source address is 192.168.100.10

B. Joins for any group are accepted if the group count value is less than 25.

C. Joins for group 224.7.7.7 are always rejected, regardless of the group count.

D. Joins for group 224.7.7.7 are accepted if the group count is less than 25

Answer: D (LEAVE A REPLY)

BGP policy framework is a set of tools that allows you to control the flow of routing information and apply routing policies based on various criteria. BGP policy framework consists of several components, such as route maps, prefix lists, community lists, AS path lists, and route filters. Route maps are used to define routing policies by matching certain conditions and applying certain actions. Prefix lists are used to filter routes based on their prefixes. Community lists are used to filter routes based on their community attributes. AS path lists are used to filter routes based on their AS path attributes. Route filters are used to filter routes based on their prefix length or range3. In this question, we have a route map named ISP-A that has two clauses: clause 10 and clause 20. Clause 10 matches any route with a prefix length between 8 and 24 bits and sets the local preference to 200. Clause 20 matches any route with a prefix of 224.7.7.7/32 and rejects it. The route map is applied inbound on the BGP neighborship with ISP-A. Based on this configuration, the correct statement is that joins for group 224.7.7.7 are always rejected, regardless of the group count. This is because clause 20 explicitly denies any route with a prefix of 224.7.7.7/32, which corresponds to the multicast group 224.7.7.7.

NEW QUESTION: 51

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-unicast configured.

Which parameter must you configure to solve this problem?

A. Under the family inet labeled-unicast hierarchy, add the resolve-vpn parameter.
B. Under the family inet labeled-unicast hierarchy, add the explicit null parameter.
C. Under the protocols mpls hierarchy, add the traffic-engineering parameter.
D. Under the protocols ospf hierarchy, add the traffic-engineering parameter.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 52
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 helps the local PE to identify the latest advertisement.
B. It identifies MAC addresses that should be discarded.
C. It resolves conflicting MAC address ownership claims.
D. It is advertised using a Type 2 message.
Answer: A,C (LEAVE A REPLY)

NEW QUESTION: 53
You are configuring schedulers to define the class-of-service properties of output queues. You want to control packet drops during periods of congestion.
In this scenario, which CoS configuration parameter would be used to accomplish this task?
A. drop profile
B. shaping rate
C. priority
D. buffer size
Answer: (SHOW ANSWER)

NEW QUESTION: 54
When using OSPFv3 for an IPv4 environment, which statement is correct?
A. OSPFv3 only supports IPv4.
B. OSPFv3 supports both IPv6 and IPv4, but not in the same routing instance.
C. OSPFv3 is not backward compatible with IPv4
D. OSPFv3 supports IPv4 only on interfaces with family inet6 defined
Answer: C (LEAVE A REPLY)
Explanation
OSPFv3 is an extension of OSPFv2 that supports IPv6 routing and addressing. OSPFv3 is not backward compatible with IPv4 because it uses a different packet format and a different link-state advertisement (LSA) structure than OSPFv2. OSPFv3 also uses IPv6 link-local addresses as router IDs and neighbor addresses, instead of IPv4 addresses. To use OSPFv3 for an IPv4

environment, you need to enable the IPv4 unicast address family under [edit protocols ospf3] hierarchy level and configure IPv4 addresses on the interfaces.

NEW QUESTION: 55
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: 56
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)
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: 57
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: 58

Exhibit

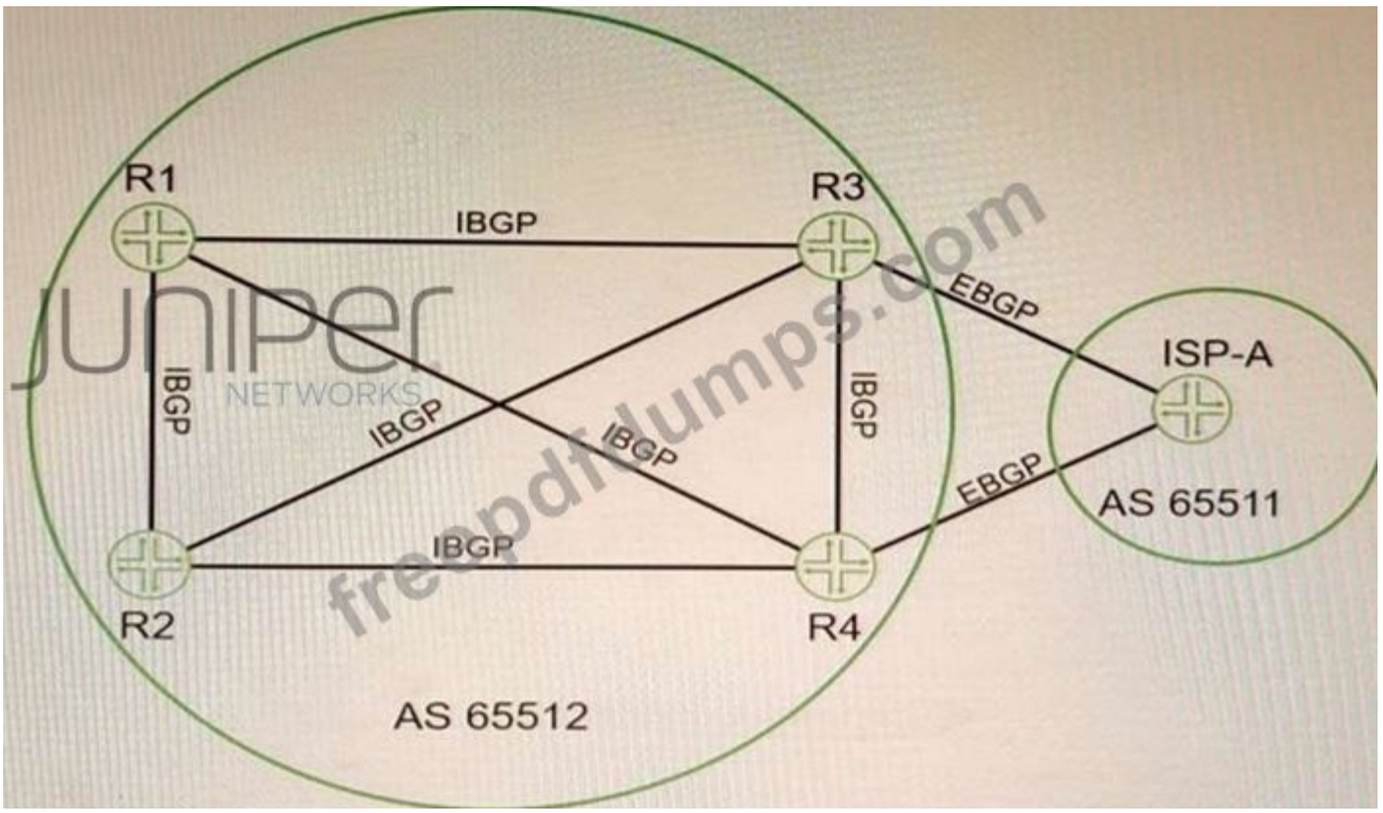

Click the Exhibit button-Referring to the exhibit, which two statements are correct about BGP routes on R3 that are learned from the ISP-A neighbor? (Choose two.)

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: A,B (LEAVE A REPLY)

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: 59
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: B (LEAVE A REPLY)
https://www.juniper.net/documentation/us/en/software/junos/cos/topics/concept/packet-flow-cos-process-cos-config-guide.html

NEW QUESTION: 60
Which two statements are correct regarding bootstrap messages that are forwarded within a PIM sparse mode domain? (Choose two.)

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

NEW QUESTION: 61
Exhibit

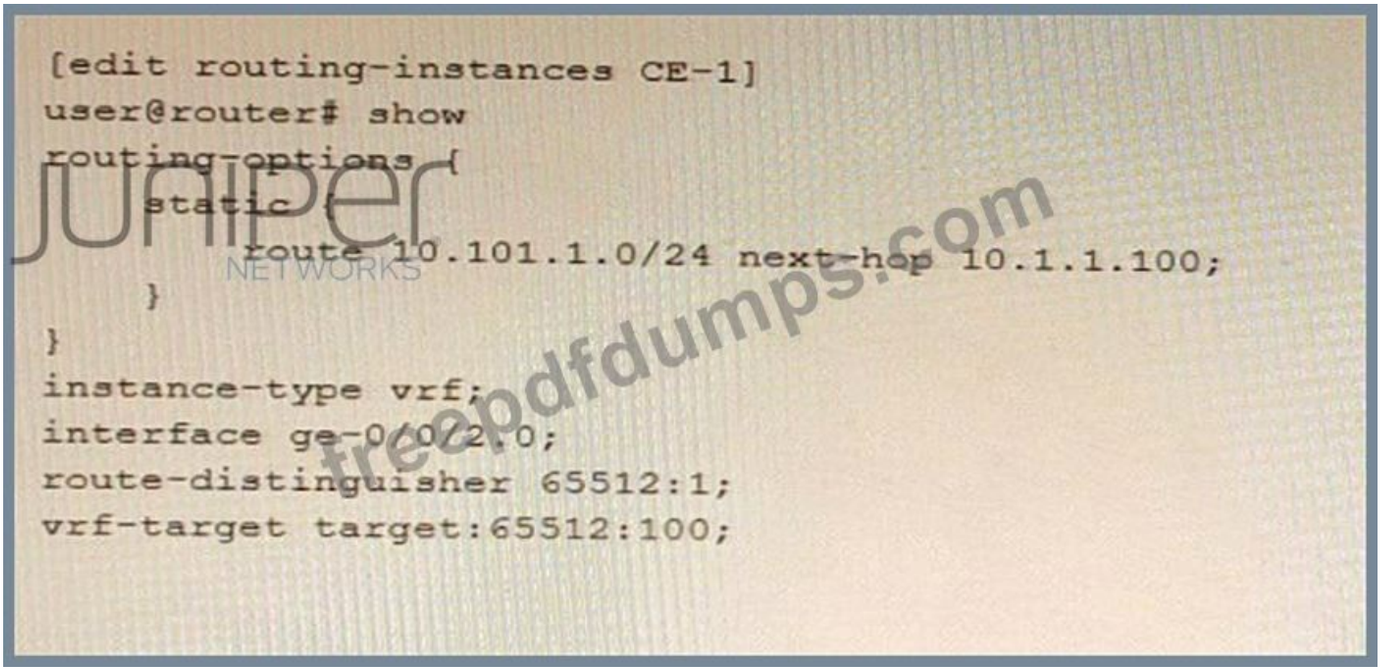

Referring to the exhibit, which statement is true?
A. The 10.101.1.0/24 route will be shared if the vrf-table-label parameter is configured.
B. The 10.101.1.0/24 route will only be shared if BGP is configured in the routing instance
C. The 10.101.1 0/24 route will be shared if there are other VRFs that 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: (SHOW ANSWER)
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: 62
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: 63
Referring to the exhibit, which two statements are true? (Choose two.)

```
user@R1> show route protocol bgp
inet.0: 8 destinations, 12 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
172.16.20.4/30     *[BGP/170] 00:49:55, localpref 100
                      AS path: 2 I, validation-state: unverified
                    > to 10.0.18.2 via ge-1/0/4.0
                      to 10.0.19.2 via ge-1/0/5.0
                    [BGP/170] 00:49:55, localpref 100
                      AS path: 2 I, validation-state: unverified
                    > to 10.0.19.2 via ge-1/0/5.0
```

A. The multihop configuration is used for load balancing.

B. This route is learned from two different AS numbers.
C. This route is learned from the same AS number.
D. The multipath configuration is used for load balancing.
Answer: C,D (LEAVE A REPLY)

NEW QUESTION: 64
Exhibit

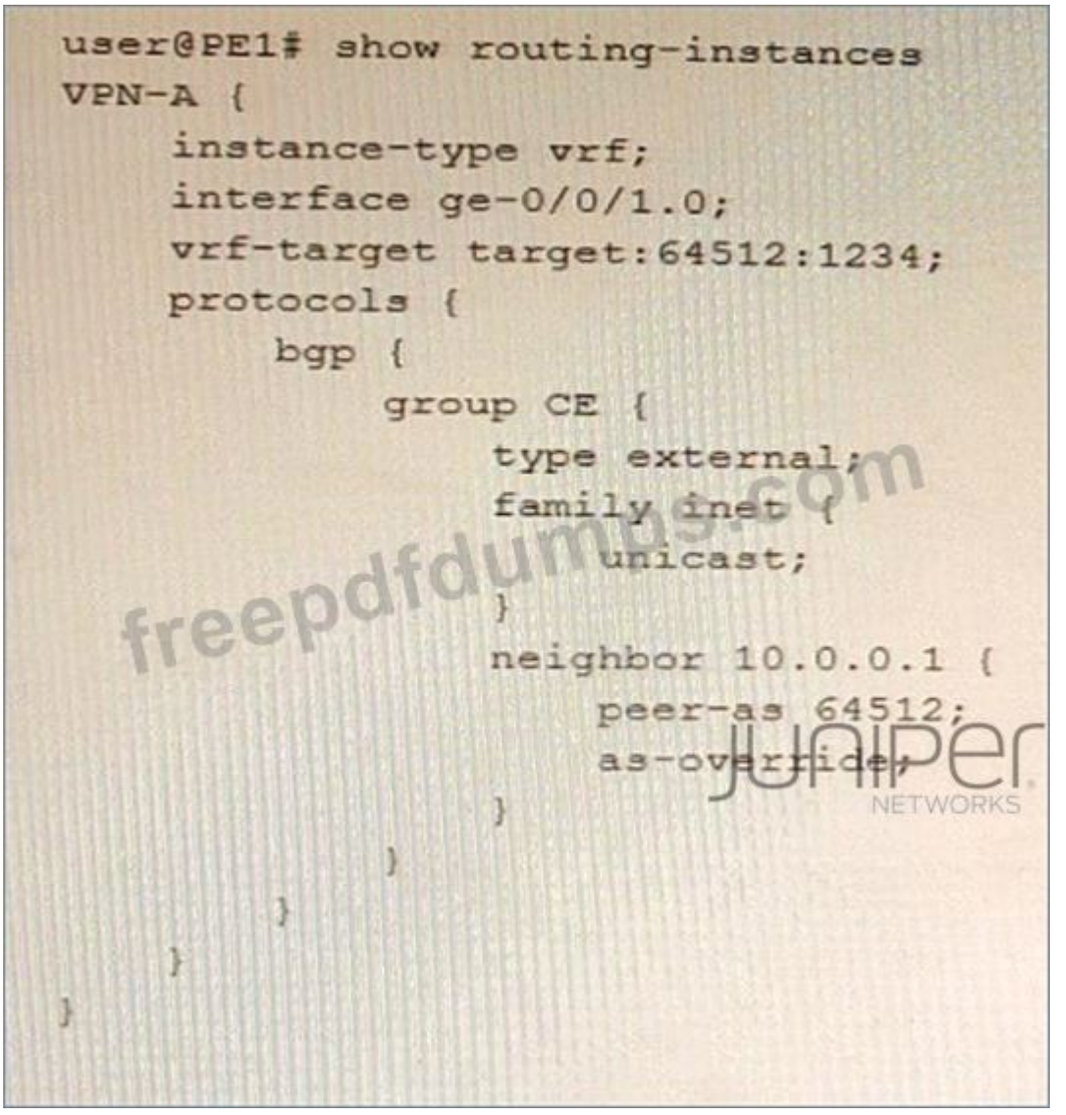

Which two statements about the configuration shown in the exhibit are correct? (Choose two.)
A. This VPN connects customer sites that use different AS numbers.
B. This VPN connects customer sites that use the same AS number
C. A Layer 2 VPN is configured.
D. A Layer 3 VPN is configured.
Answer: (SHOW ANSWER)
The configuration shown in the exhibit is for a Layer 3 VPN that connects customer sites that use different AS numbers. A Layer 3 VPN is a type of VPN that uses MPLS labels to forward packets across a provider network and BGP to exchange routing information between PE routers and CE routers. A Layer 3 VPN allows customers to use different routing protocols and AS numbers at

their sites, as long as they can peer with BGP at the PE-CE interface. In this example, CE-1 is using AS 65530 and CE-2 is using AS 65531, but they can still communicate through the VPN because they have BGP sessions with PE-1 and PE-2, respectively.

NEW QUESTION: 65
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: 66
Exhibit

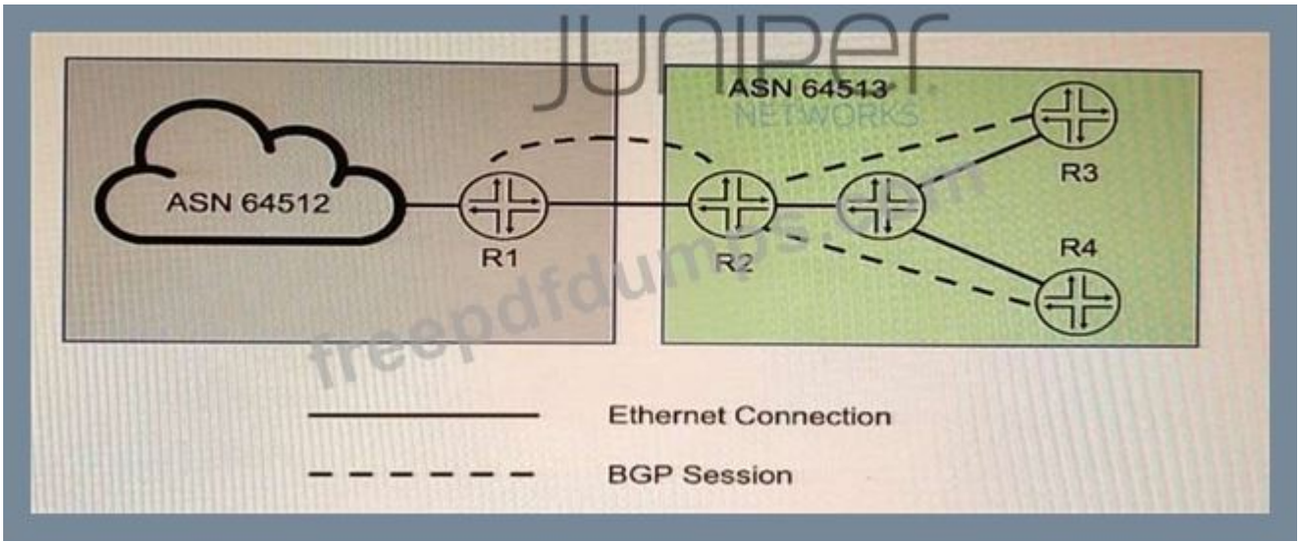

You want to implement the BGP Generalized TTL Security Mechanism (GTSM) on the network Which three statements are correct in this scenario? (Choose three)

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)

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: 67
Your organization manages a Layer 3 VPN for multiple customers. To support advanced route filtering on your PE routers, you must advertise more than one BGP community on advertised VPN routes to remote PE routers.
Which routing-instance configuration parameter would support this requirement?
A. vrf-import
B. vrf-target export
C. vrf-target import
D. vrf-export
Answer: D (LEAVE A REPLY)

NEW QUESTION: 68
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.

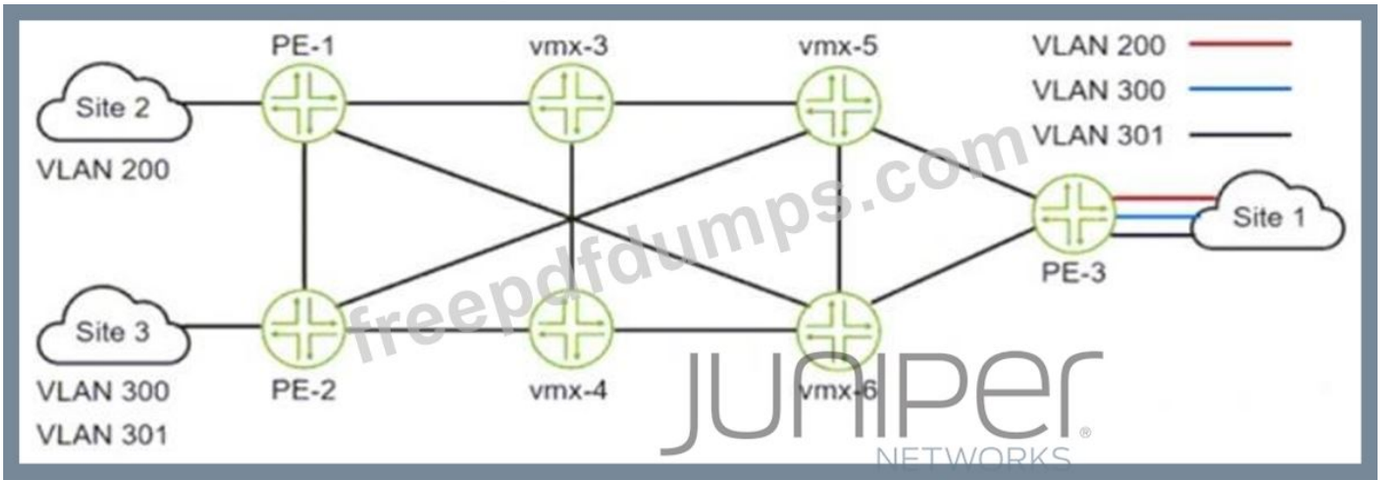

What is the minimum number of L2VPN routing instances to be configured to accomplish this task?
A. 3
B. 6
C. 2
D. 1
Answer: (SHOW ANSWER)

NEW QUESTION: 69
Which three mechanisms are used by Junos platforms to evaluate incoming traffic for CoS purposes? (Choose three.)
A. rewrite rules
B. multifield classifiers
C. behavior aggregate classifiers
D. traffic shapers
E. fixed classifiers
Answer: B,C,E (LEAVE A REPLY)

NEW QUESTION: 70
Exhibit

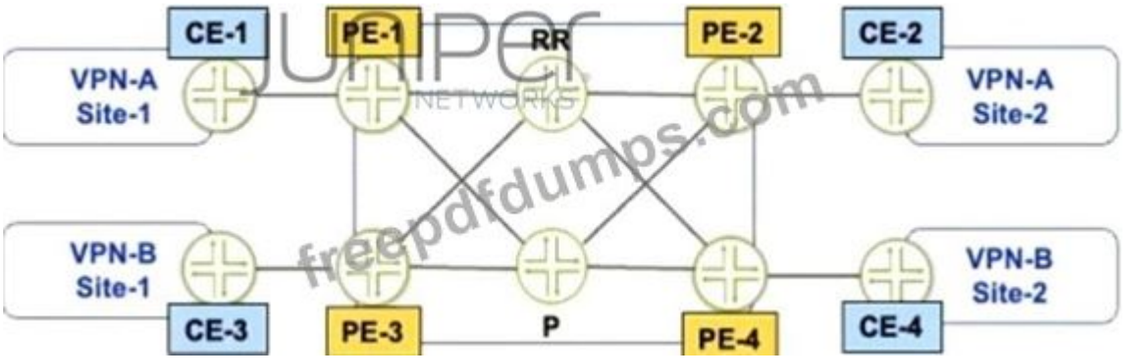

Referring to the exhibit, PE-1 and PE-2 are getting route updates for VPN-B when neither of them service that VPN Which two actions would optimize this process? (Choose two.)
A. Configure the resolution rib bgp.l3vpn.0 resolution-ribs inet.0 statement on the PEs.
B. Configure the family route-target 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: 71
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: 72
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. EVPN uses Type 2 routes to advertise MAC address and IP address pairs learned using ARP snooping.
B. VXLAN must be configured on all PE routers.
C. EVPN uses Type 3 routes to join a multicast tree to flood traffic.
D. Segment routing must be configured on all PE routers.
Answer: A,C (LEAVE A REPLY)

NEW QUESTION: 73
Exhibit

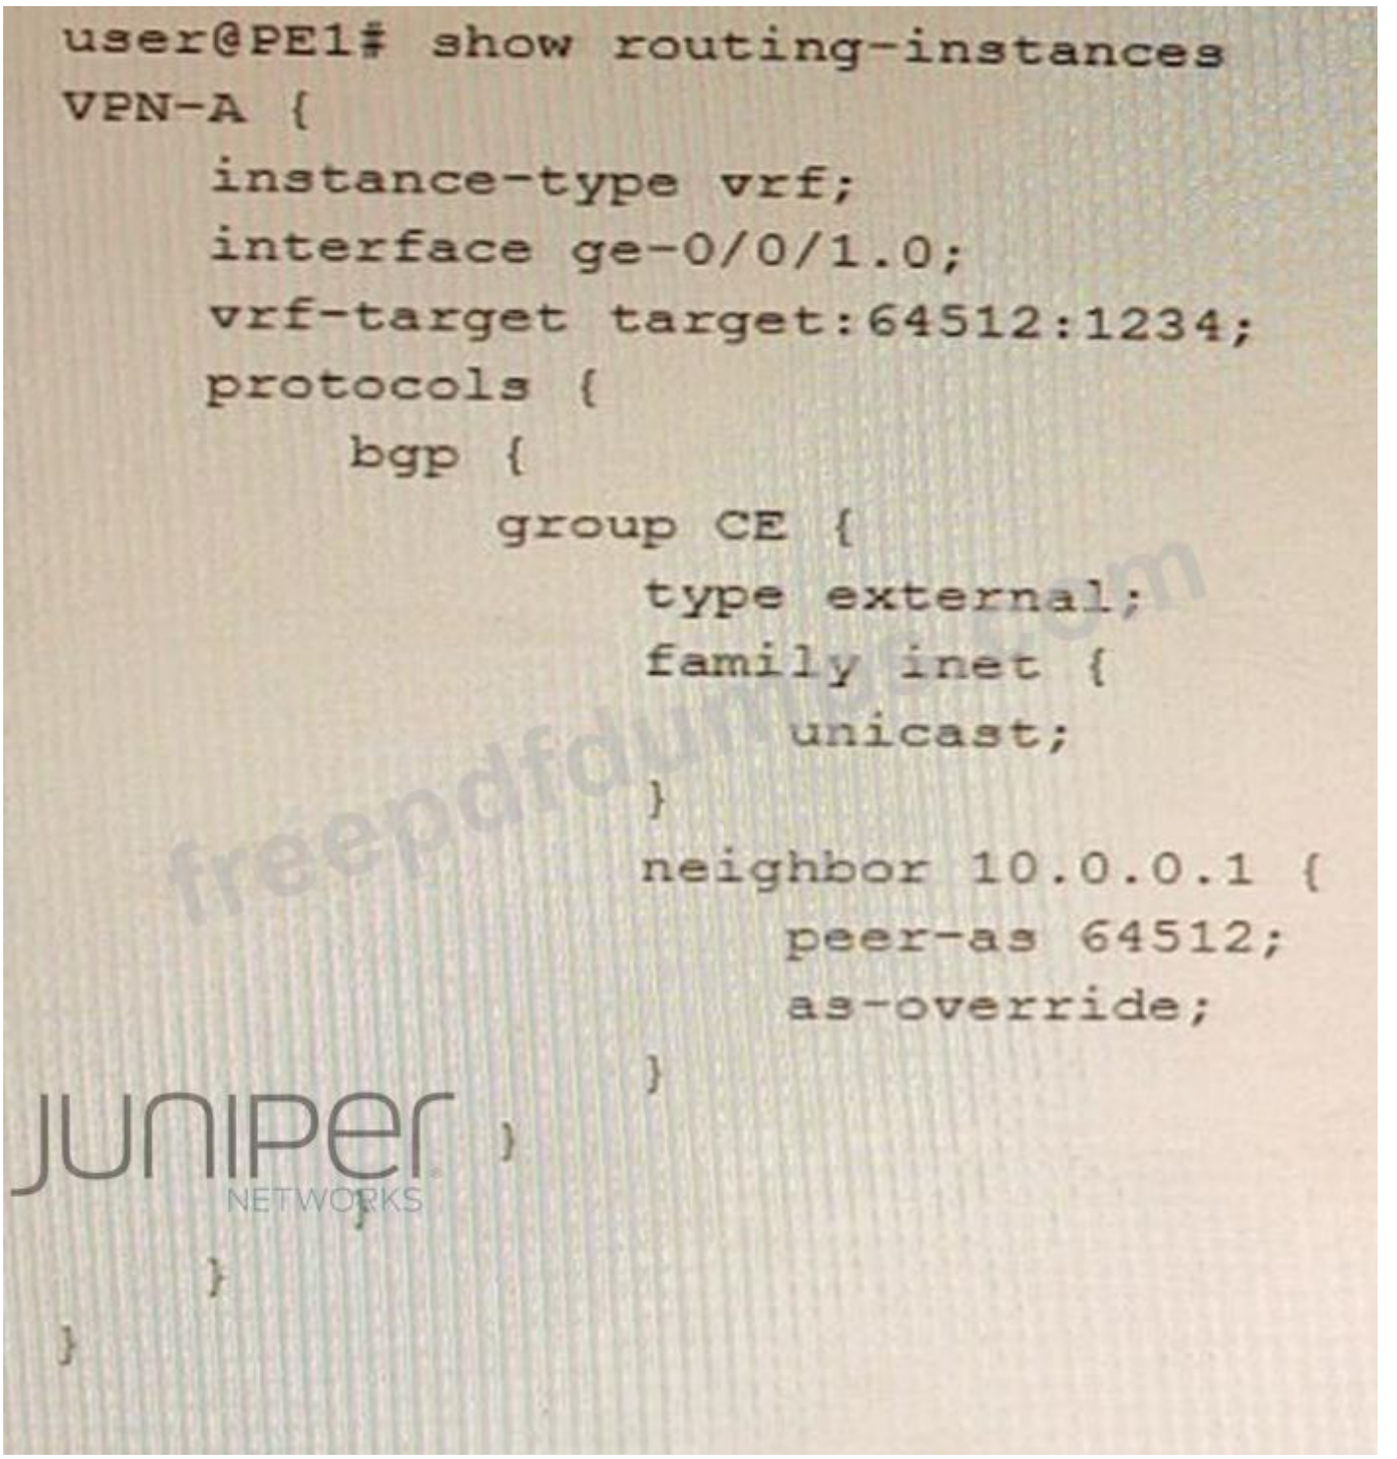

Which two statements about the configuration shown in the exhibit are correct? (Choose two.)

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: (SHOW ANSWER)

Explanation

The configuration shown in the exhibit is for a Layer 3 VPN that connects customer sites that use different AS numbers. A Layer 3 VPN is a type of VPN that uses MPLS labels to forward packets across a provider network and BGP to exchange routing information between PE routers and CE

routers. A Layer 3 VPN allows customers to use different routing protocols and AS numbers at their sites, as long as they can peer with BGP at the PE-CE interface. In this example, CE-1 is using AS 65530 and CE-2 is using AS 65531, but they can still communicate through the VPN because they have BGP sessions with PE-1 and PE-2, respectively.

NEW QUESTION: 74

Exhibit

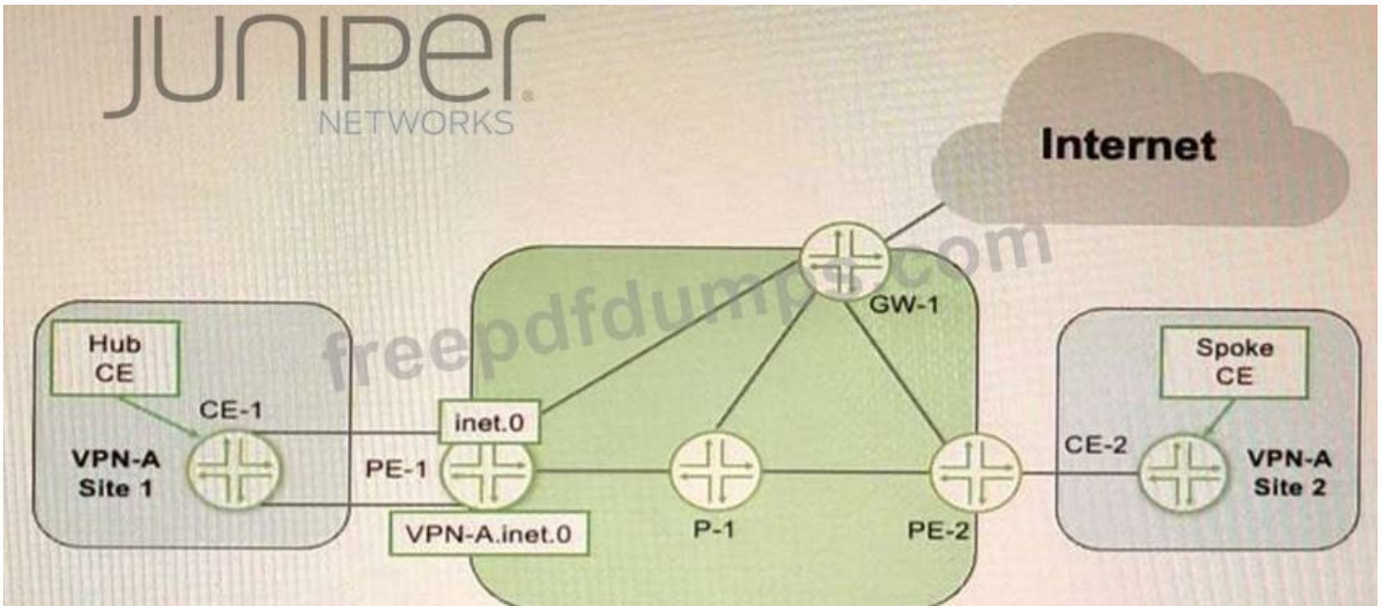

Referring to the exhibit, you must provide Internet access for VPN-A using CE-1 as the hub CE. Which two statements are correct in this situation? (Choose two.)

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: 75

Exhibit

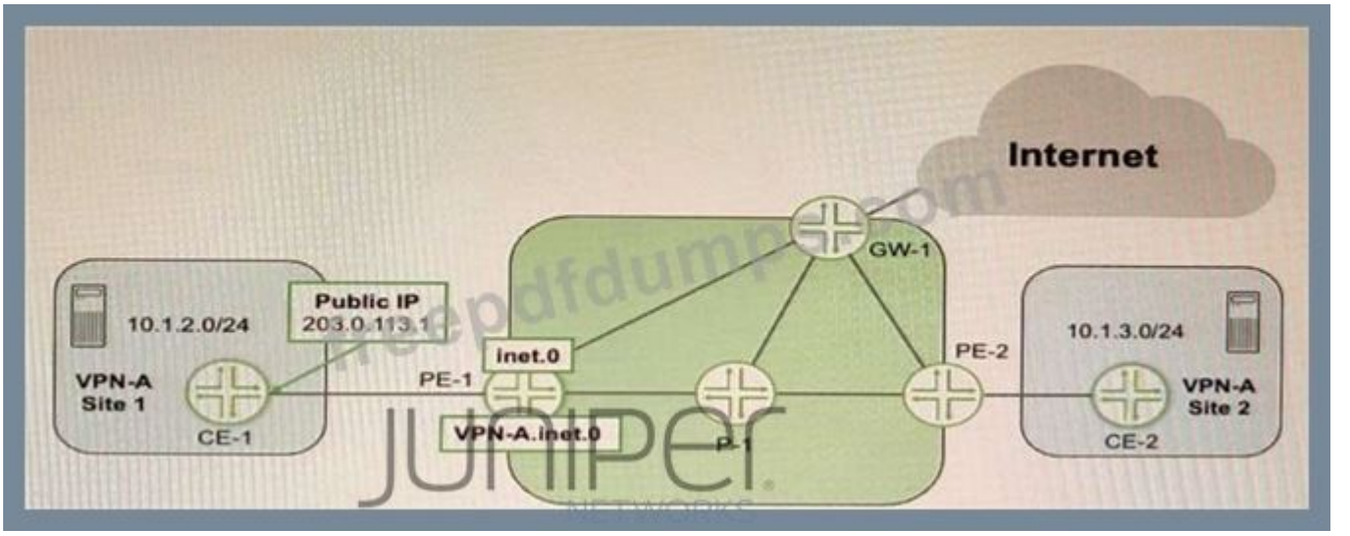

Referring to the exhibit, CE-1 is providing NAT services for the hosts at Site 1 and you must provide Internet access for those hosts Which two statements are correct in this scenario? (Choose two.)

A. You must configure a static route in the main routing instance for the 10 1 2.0/24 prefix that uses the VPN-A.inet.0 table as the next hop

B. You must configure a static route in the main routing instance for the 203.0.113.1/32 prefix that uses the VPN-A.inet.0 table as the next hop.

C. You must configure a RIB group on PE-1 to leak a default route from the inet.0 table to the VPN-A.inet.0 table.

D. You must configure a RIB group on PE-1 to leak the 10 1 2.0/24 prefix from the VPN-A.inet.0 table to the inet.0 table.

Answer: A,B (LEAVE A REPLY)

Explanation

To provide Internet access for the hosts at Site 1, you need to configure static routes in the main routing instance on PE-1 that point to the VPN-A.inet.0 table as the next hop. This allows PE-1 to forward traffic from the Internet to CE-1 using MPLS labels and vice versa. You need to configure two static routes: one for the

10.1.2.0/24 prefix that represents the private network of Site 1, and one for the 203.0.113.1/32 prefix that represents the public IP address of CE-1.

NEW QUESTION: 76

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: 77

Exhibit

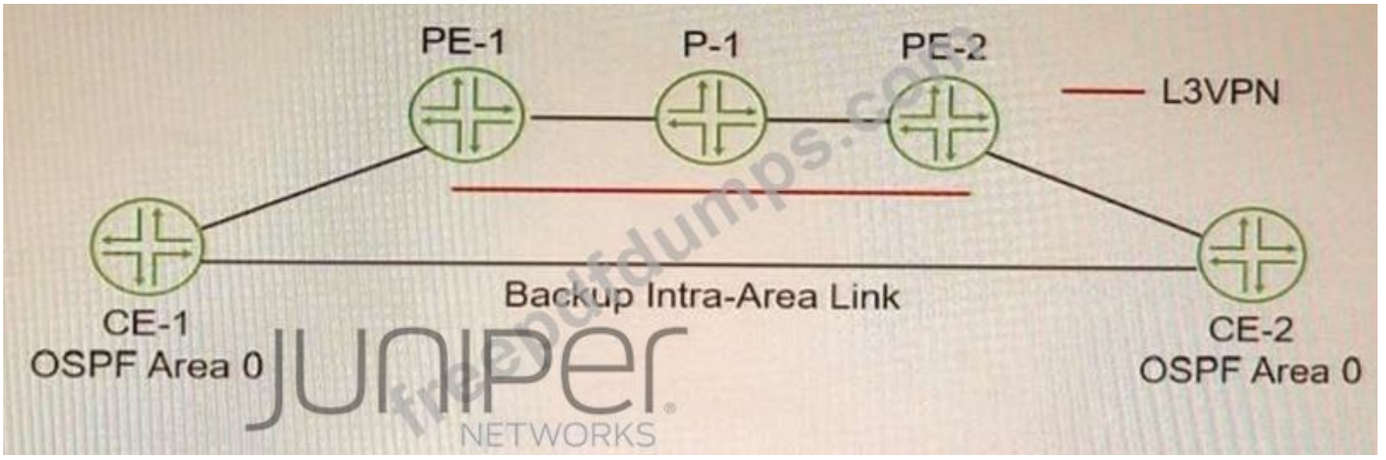

You must ensure that the VPN backbone is preferred over the back door intra-area link as long as the VPN is available. Referring to the exhibit, which action will accomplish this task?

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)

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.