DUMPSARENA

Nokia Virtual Private Routed Networks

Alcatel-Lucent 4A0-106

Version Demo

Total Demo Questions: 29

Total Premium Questions: 299

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Volume A	121
Topic 2, Volume B	158
Total	299

QUESTION NO: 1

For the data packets in a VPRN, how is the inner label in an MPLS packet signaled?

A. Signaled via MP-BGP and used to identify the egress PE.

B. Signaled via T-LDP and used to identify the customer VPRN.

C. Signaled via MP-BGP and used to identify the customer VPRN.

D. Signaled via RSVP and LDP and used to identify the egress PE.

E. Signaled via RSVP and LDP and used to identify the customer VPRN.

ANSWER: C

Explanation:

Signaled via MP-BGP and used to identify the customer VPRN is correct. In an MPLS Layer 3 VPN, commonly implemented as a VPRN on Nokia 7750 SR routers, a provider-edge router imposes a two-label stack on VPN traffic. The outer transport label provides the path across the provider MPLS network toward the egress provider-edge router. The inner VPN label is allocated by the egress provider-edge router and advertised to remote provider-edge routers with the VPN route using Multiprotocol BGP. The label is carried with the route's BGP label field, together with the route distinguisher and route-target-based VPN membership information. When the packet reaches the egress provider-edge router, that router removes the transport label and uses the inner VPN label to select the appropriate VPRN forwarding context, enabling the packet to be delivered using the correct customer VPN routing table. This separation allows identical customer prefixes to exist in different VPRNs while retaining unambiguous forwarding at the egress. The MPLS VPN label distribution and forwarding model is defined in RFC 4364, BGP/MPLS IP Virtual Private Networks; Nokia's VPRN implementation is documented in the Nokia 7750 SR VPRN overview.

QUESTION NO: 2

Which of the following commands may be used on an Alcatel-Lucent 7750 SR to display a network exported locally to BGP on the originating router? (Choose two)

```
community "Customer 1" members  "target:65100:2"
community "Customer 1 and 2" members  "target:65100:3"
policy-statement "Sample Policy"
                entry 10
                    action accept
                        community add "Customer 1 and 2"
                    exit
                exit
            exit
policy-statement "Sample Policy 3"
                entry 10
                    from
                        community "Customer 1 and 2"
                    exit
                 action accept
                 exit
```

A. Show router bgp routes

B. Show router bgp routes prefix

C. Show router bgp routes self-originated

D. Show router bgp routes hunt

E. Show router bgp neighbor advertised-routes

ANSWER: C E

Explanation:

On the originating 7750 SR, `Show router bgp routes <prefix> self-originated` can identify the specified BGP route as self-originated. The `self-originated` qualifier limits the displayed route information to routes that were originated by the local router, which is the relevant perspective when confirming that a locally available network has entered BGP through local export or origination processing. This provides a direct route-table view without requiring the route to be learned back from another BGP speaker.

`Show router bgp neighbor <neighbor ip=""> advertised-routes` provides the complementary export-verification view. It displays the routes that the local BGP speaker has advertised to a specified peer after BGP export policy processing. Therefore, it can be used on the originating router to confirm that the locally exported network is actually being sent to that BGP neighbor. Together, these commands validate both local origination and per-peer advertisement, which are distinct but useful checks when troubleshooting BGP route export. Nokia's SR OS documentation describes BGP route display and neighbor operational commands in the SR OS documentation portal; see also the 7750 SR BGP configuration guide.

QUESTION NO: 3

Which of the following about inter-AS model AVPRN is TRUE?

A. It is the most scalable of the Inter-AS models.

B. VRF configuration is not required on the ASBRs.

C. Most routers that support VPRN functionality do not support model A by default.

D. It is not necessary to run MPLS at the boundary between autonomous systems.

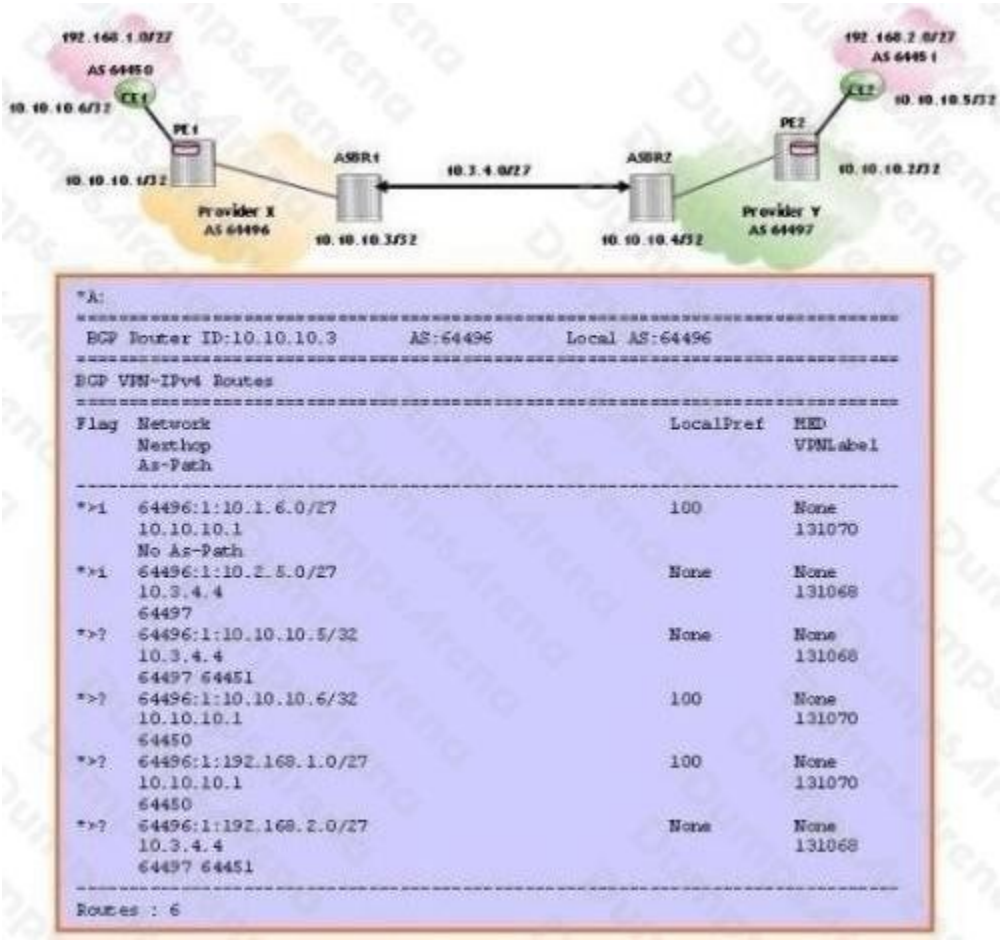
ANSWER: D

Explanation:

It is not necessary to run MPLS at the boundary between autonomous systems. In inter-AS VPRN model A, each autonomous system maintains its own VPN and MPLS domain. The autonomous-system boundary routers provide connectivity between corresponding customer VPN instances by placing interfaces toward the neighboring autonomous system in the applicable VPRN service and exchanging VPN customer routes, commonly through an external BGP peering. Consequently, labeled MPLS forwarding does not need to extend across the inter-AS link: packets can be forwarded as ordinary IP traffic over that boundary and then use the receiving autonomous system's local MPLS transport and VPRN label stack as needed. This model cleanly separates label-distribution and transport responsibilities between providers while still allowing end-to-end VPN reachability. RFC 4364 describes this approach as the inter-provider VPN method in which VRFs are connected by back-to-back interfaces and VPN routes are exchanged between those VRFs. Nokia's Service Router documentation provides the platform context for VPRN service operation and inter-provider deployments. See RFC 4364, Section 10a and Nokia Service Router documentation.

QUESTION NO: 4

Click the exhibit.



Which of the following commands produces the output for the inter-AS model B VPRN?

A. ASBR1# show router 10 bgp routes vpn-ipv4

B. ASBR1# show router bgp routes

C. ASBR1# show router bgp routes vpn-ipv4

D. ASBR1# show router route-table

ANSWER: C

Explanation:

`ASBR1# show router bgp routes vpn-ipv4` is the correct command because an inter-AS VPRN deployment using model B exchanges labeled VPN-IPv4 reachability directly between the autonomous-system border routers. The relevant routes are therefore held in the base router's BGP VPN-IPv4 address family, rather than being queried from an individual VPRN service router. This command displays the VPN-IPv4 BGP routing information, including the route distinguisher-qualified prefixes and associated VPN label information that permit the receiving border router to retain customer-VPN route context while forwarding traffic across the inter-AS boundary. On SR OS, the `show router bgp routes` command supports address-family-specific route display; selecting `vpn-ipv4` narrows the output to precisely the VPN route family used for this control-plane exchange. Nokia's SR OS documentation describes BGP route display commands in its command reference and routing documentation: Nokia SR OS Classic CLI Guide. The VPN-IPv4 route format and its use for Layer 3 VPN control-plane distribution are defined in RFC 4364.

QUESTION NO: 5

When does a router add a Route Target to a route?

A. When the P receives the route from the PE.

B. When the PE populates the route received from the CE into the VRF.

C. When the PE propagates the route to the CE.

D. When the PE exports the route from the VRF into the MP-BGP table.

ANSWER: D

Explanation:

A router adds a Route Target when the PE exports the route from the VRF into the MP-BGP table. In an RFC 4364 BGP/MPLS IP VPN, a customer route learned in a VRF is converted into a VPN-IPv4 route for distribution through MP-BGP. At this export stage, the PE attaches one or more Route Target extended communities according to the VRF's export policy. The Route Target identifies the VPN membership or distribution scope for that route; remote PEs use their VRFs' import Route Target policies to determine whether the received VPN route belongs in a particular VRF. The Route Target is therefore associated with VPN route distribution between PEs, rather than with ordinary packet forwarding in the provider core or the PE-to-CE routing exchange. This separation allows overlapping customer address spaces to be carried safely: the route distinguisher makes the VPN route unique, while the Route Target controls which VRFs may import it. RFC 4364 specifies that a PE attaches export Route Target attributes when it distributes VPN-IPv4 routes, and that receiving PEs use Route Target import attributes to select routes for VRFs. See RFC 4364, section 4.3.1 and Nokia SR OS Classic VPN documentation.

QUESTION NO: 6

In a VPRN, which of the following are supported as PE-CE routing methods on the Alcatel-Lucent 7750 SR? (Choose four).

A. Static

B. RIP

C. OSPF

D. IS-IS

E. MP-BGP

F. BGP

ANSWER: A B C F

Explanation:

On the 7750 SR, a VPRN supports the standard PE-to-CE routing choices of Static, RIP, OSPF, and BGP. Static routing allows the provider-edge router to install explicitly configured customer-prefix reachability in the VPRN routing table. It is appropriate where the CE topology is small, stable, or intentionally controlled without routing-protocol exchanges.

RIP and OSPF can operate in the VPRN context so that the PE and CE exchange customer routes dynamically. These IGP choices allow the CE-facing part of the VPN to use a routing protocol appropriate to the customer network while keeping each customer's routes isolated in the relevant VPRN. OSPF additionally supports the common requirement to exchange internal customer routes and externally learned VPN routes under defined OSPF policies.

BGP is also supported for PE-CE route exchange. In this role, BGP establishes a CE-facing peer relationship within the VPRN and advertises or receives customer VPN prefixes according to the configured import, export, and routing policies. This is distinct from the provider VPN route-distribution function, but it is still a supported customer-edge routing method. Nokia's SR OS service documentation describes VPRN IP routing and supported routing-protocol configuration: Nokia 7750 SR OS documentation and Nokia 7750 Service Router product documentation resources.

QUESTION NO: 7

What are the two main benefits of a hub and spoke VPRN design? (Choose two)

A. Reduced number of VPN tunnels that need to be managed

B. Centralized filtering policies with the introduction of a hub site

C. Optimal routing between all sites

D. Limited configuration required on the hub PE device

ANSWER: A B

Explanation:

A hub-and-spoke VPRN is designed so that spoke sites exchange traffic through a central hub rather than requiring direct connectivity between every pair of sites. Consequently, *Reduced number of VPN tunnels that need to be managed* is a principal benefit: as the number of branch sites grows, the design avoids the operational overhead associated with a full mesh of intersite transport adjacencies or logical paths. This provides a simpler and more scalable service topology for customers whose branch-to-branch traffic does not need direct forwarding.

Centralized filtering policies with the introduction of a hub site is also a key benefit. Because inter-spoke traffic traverses the hub, the hub becomes a natural enforcement point for security, Internet access, firewall service insertion, route control, and traffic-monitoring policies. Administrators can apply and maintain common policy controls at one central location while retaining the VPRN separation provided by the provider edge routers. This model is especially useful when a head-office or data-centre site must inspect, authorize, or log traffic from many remote locations. Nokia's Service Router documentation covers Layer 3 VPN service architectures, while RFC 4364 defines the BGP/MPLS IP VPN model underlying VPRN-style services: Nokia Service Router documentation and RFC 4364.

QUESTION NO: 10

Why would a selective export policy be applied to a customer VPRN service? (Choose three)

A. To limit the number of routes advertised from a customer site

B. To limit the number of routes learned from other customer sites

C. To limit connectivity to selected customer networks

D. To perform packet filtering

ANSWER: A B C

Explanation:

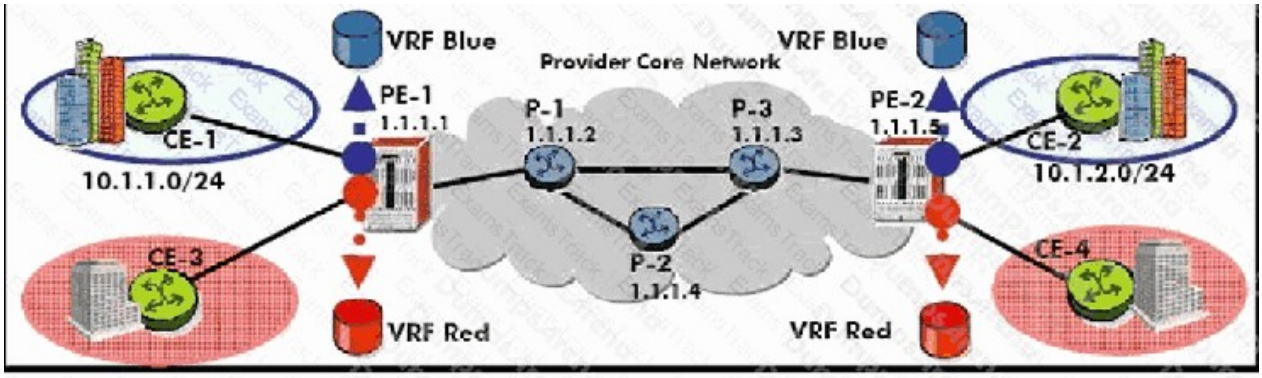
A selective export policy is used to control route distribution for a customer VPRN rather than indiscriminately making every customer route available throughout the VPN. It can limit the number of routes advertised from a customer site by matching only the required prefixes before they are exported into the provider VPN routing domain. This provides route-scale control and prevents unnecessary reachability information from being propagated.

Selective route-policy control also supports limiting the routes a site learns from other customer sites. By constraining the eligible customer prefixes for VPN route exchange, the provider can keep unnecessary routes out of the relevant customer routing context and reduce the effective routing-table requirements for that site. This is especially useful where customer locations need only partial reachability.

Finally, exporting only selected routes is a fundamental way to limit connectivity to selected customer networks. If reachability to a destination prefix is not distributed, traffic cannot use the VPRN control plane to reach that destination. This implements customer-specific topology and segmentation requirements using routing policy. The underlying BGP/MPLS VPN route-distribution model is described in RFC 4364; Nokia's service-routing documentation and configuration references are available through the Nokia SR OS documentation portal.

QUESTION NO: 13

When configuring the CE with the selected CE-PE routing protocol on an Nokia 7750 SR, which of the following statements is true? (Choose three.)




```
PE-1>config>service# VPRN 20 customer 101 create
PE-1>config>service>VPRN$ description "VPRN Service"
PE-1>config>service>VPRN$ route-distinguisher 65100:101
PE-1>config>service>VPRN$ vrf-target target:65100:101
PE-1>config>service>VPRN$ auto-bind ldp
PE-1>config>service>VPRN$ router-id 1.1.1.1
PE-1>config>service>VPRN$ exit
PE-1>config>service#
```

A. The protocol is configured under the config>service context

B. The protocol is configured under the global router context

C. The interface leading to the PE is defined as a protocol neighbor

D. The routing protocol used must be the same as the routing protocol used at other CE-PE locations in the same VPN.

E. A routing policy may or may not be required depending on the routing requirements and routing protocol

ANSWER: B C E

Explanation:

When a Nokia 7750 SR is used as the customer-edge router, customer-side routing is configured in the router's global routing context. This is where the router owns its customer-facing interfaces, routing-table behavior, and routing-protocol processes. The selected dynamic CE-to-PE protocol establishes reachability with the provider-edge router across the connecting interface. Consequently, the interface toward the provider edge is associated with the protocol's peer or neighbor relationship, allowing routing information to be exchanged over that attachment. The exact configuration syntax varies by protocol—for example, a peer relationship is explicitly configured for BGP, while an interface participates in an interface-based protocol such as OSPF or IS-IS—but the CE-to-PE attachment is the relevant protocol adjacency. A routing policy may also be needed to control route import, export, filtering, redistribution, or route preference. Whether it is necessary depends on the protocol selected and on the intended customer routing design; simple deployments may not need additional policy, whereas controlled route exchange commonly does. Nokia's SR OS documentation describes router-context routing configuration and the protocol-specific controls available on the platform. See the Nokia SR OS documentation portal and the SR OS router configuration guide.

QUESTION NO: 14

Which of the following statements is true regarding the Route Distinguisher? (Choose three)

A. It is an 8 byte value containing 3 fields

B. The Route Distinguisher must have the same value as the Route Target.

C. The Administrator field contains either an AS number or an IP number

D. To ensure uniqueness, Route Distinguishers are assigned by the IANA.

E. The Route Distinguisher is not used if the customer addresses do not overlap

F. The Assigned Number field must contain a public IP address

G. The Route Distinguisher is assigned by the service provider and combines with the IPv4 prefix to form a unique VPN-IPv4 address.

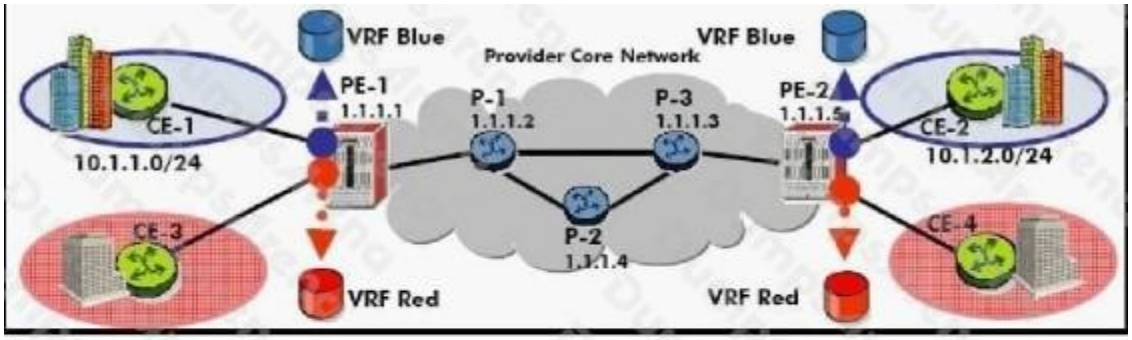
ANSWER: A C G

Explanation:

A Route Distinguisher is an 8-octet value prepended to an IPv4 prefix to form a VPN-IPv4 address. This allows identical customer IPv4 prefixes used in separate VPNs to be advertised and handled as distinct routes in Multiprotocol BGP. Its encoding comprises three logical components: a 2-octet type field, an administrator field, and an assigned-number field. The type determines the format and length of the remaining fields. In the standard formats, the administrator component identifies either an autonomous-system number or an IPv4 address, while the assigned-number component is chosen by the service provider to distinguish routes or VPN contexts under that administrator value. Route Distinguishers are provider-selected values, rather than values allocated by IANA, and their purpose is route uniqueness; they are separate from Route Targets, which control VPN route import and export policy. The statement that the Route Distinguisher is assigned by the service provider and combines with the IPv4 prefix to form a unique VPN-IPv4 address captures this required operational role. See RFC 4364, Section 4.2 for the Route Distinguisher format and VPN-IPv4 construction, and Nokia's Service Router Operating System documentation portal for SR OS service documentation.

QUESTION NO: 15

When configuring the CE with the selected CE-PE routing protocol on an Alcatel-Lucent 7750 SR, which of the following statements is true? (Choose three.)




```
PE-1>config>service# VPRN 20 customer 101 create
PE-1>config>service>VPRN$ description "VPRN Service"
PE-1>config>service>VPRN$ route-distinguisher 65100:101
PE-1>config>service>VPRN$ vrf-target target:65100:101
PE-1>config>service>VPRN$ auto-bind ldp
PE-1>config>service>VPRN$ router-id 1.1.1.1
PE-1>config>service>VPRN$ exit
PE-1>config>service#
```

A. The protocol is configured under the config>service context

B. The protocol is configured under the global router context

C. The interface leading to the PE is defined as a protocol neighbor

D. The routing protocol used must be the same as the routing protocol used at other CE-PE locations in the same VPN.

E. A routing policy may or may not be required depending on the routing requirements and routing protocol

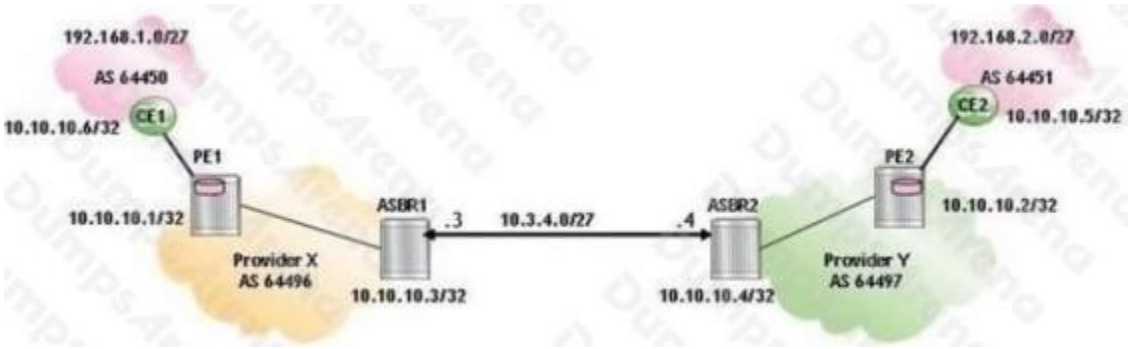
ANSWER: B C E

Explanation:

On a router operating as the customer-edge device, the CE–PE routing protocol belongs to the router's global routing-instance configuration. This is the appropriate context for the CE's own routing table and for protocol operation toward the provider edge; it is distinct from the service-based virtual-router context used when the 7750 SR itself is providing a customer VPRN service. The interface or peer endpoint toward the provider edge must be associated with the selected routing protocol as a protocol neighbor or adjacency, allowing routing information to be exchanged across the CE–PE link. The exact syntax differs by protocol, but the essential requirement is to identify and activate the provider-edge peer through that link. Routing policy is conditional rather than universally mandatory: it is applied when route import, export, redistribution, filtering, attribute control, or route-selection requirements call for it. A simple deployment can exchange the necessary routes without an additional policy, whereas a deployment requiring controlled route propagation uses one. This separation of customer routing from provider VPN routing aligns with the PE/CE model described in RFC 4364 and Nokia's 7750 SR documentation portal.

QUESTION NO: 16

Click the exhibit.



For the inter-AS model B VPRN, which of the following is the correct BGP configuration on ASBR1?

A.
```
ASBR1>config>router>bgp# info
----------------------------------------------
        enable-inter-as-vpn
        group "ASBR"
            family ipv4
            neighbor 10.10.10.4
                type external
                peer-as 64497
            exit
```

B.
```
ASBR1>config>router>bgp# info
----------------------------------------------
        enable-inter-as-vpn
        group "ASBR"
            family vpn-ipv4
            neighbor 10.10.10.4
                type external
                peer-as 64497
            exit
```

C.
```
ASBR1>config>router>bgp# info
----------------------------------------------
        enable-inter-as-vpn
        group "ASBR"
            family vpn-ipv4
            neighbor 10.3.4.4
                type external
                peer-as 64497
            exit
```

D.
```
ASBR1>config>router>bgp# info
----------------------------------------------
        enable-inter-as-vpn
        group "ASBR"
            family ipv4
            neighbor 10.3.4.4
                next-hop-self
                type external
                peer-as 64497
            exit
```

A. Option A

B. Option B

C. Option C

D. Option D

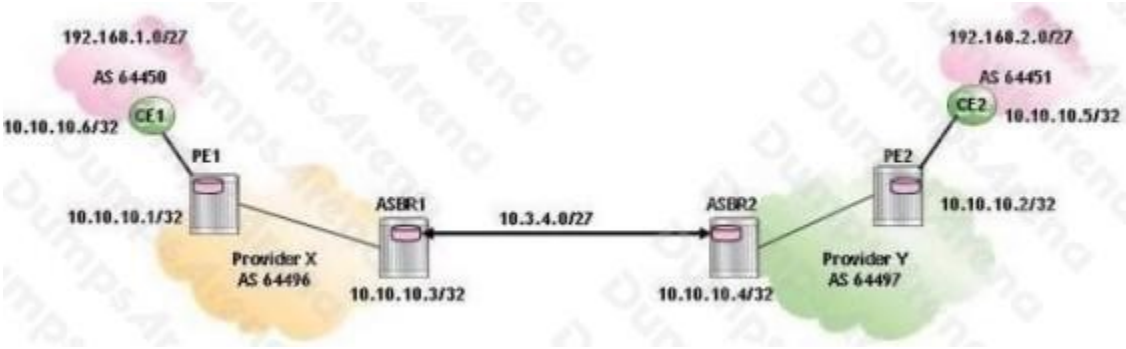
ANSWER: C

Explanation:

Option C is correct because an inter-AS VPRN using model B requires the AS border routers to exchange VPN-IPv4 routing information directly through a multiprotocol eBGP session. ASBR1 must therefore have BGP configured toward the remote ASBR using the VPN-IPv4 address family, rather than merely exchanging ordinary IPv4-unicast reachability. The VPN-IPv4 NLRI preserves each route's route distinguisher and carries the MPLS label information needed for the receiving ASBR to forward traffic toward the appropriate remote VPN. The ASBR also acts as the inter-AS MPLS forwarding boundary: it advertises VPN routes to its eBGP peer with itself as the usable next hop and provides the label binding required for packets arriving from the neighboring autonomous system. This is the defining control-plane behavior of inter-provider model B, where VPN routes are exchanged between ASBRs by MP-eBGP without requiring the provider edge routers in one AS to peer directly with those in the other AS. RFC 4364 describes this model's use of eBGP to distribute labeled VPN-IPv4 routes between ASBRs: RFC 4364, BGP/MPLS IP VPNs. Nokia's VPRN documentation is available through the Nokia SR OS documentation portal.

QUESTION NO: 17

Click the exhibit.



For the inter-AS model A VPRN shown, which of the following about the route advertisement of 192.168.1.0/27 is FALSE?

A. PE1 advertises 64496:1:192.168.1.0/27 route update to ASBR1.

B. PE2 advertises 192.168.1.0/27 route update to CE2.

C. ASBR1 advertises 64496:1:192.168.1.0/27 route update to ASBR2.

D. ASBR2 advertises 64497:1:192.168.1.0/27 route update to PE2.

ANSWER: C

Explanation:

ASBR1 advertises 64496:1:192.168.1.0/27 route update to ASBR2. is the false statement. In an inter-AS VPRN model A deployment, the ASBRs are connected through per-VPRN routing instances (VRFs). The inter-AS connection exchanges the customer IPv4 route in the relevant VPRN context, rather than carrying a VPN-IPv4 route with the originating AS's route distinguisher across the AS boundary. Consequently, after ASBR1 receives the VPN-IPv4 advertisement for the customer prefix from PE1 and imports it into its local VPRN, it advertises the IPv4 prefix 192.168.1.0/27 to ASBR2 through the model A inter-AS VRF-to-VRF peering. ASBR2 then imports that route into its own VPRN and can advertise it within its AS using its local VPN-IPv4 route distinguisher, 64497:1. This preserves separation of the VPRN routing domains while allowing the customer route to be exchanged between autonomous systems. RFC 4364 describes inter-provider VPN model A as back-to-back VRFs on the AS border routers; see RFC 4364, Section 10. Nokia's VPRN documentation is available at Nokia SR OS VPRN Guide.

QUESTION NO: 18

When a Service Provider offers VPRN services to its customers, which of the following functions are expected to be the responsibility of the Service Provider? (Choose three)

A. Distributing the customer generated labels between sites

B. Distributing the customer routing information between

C. Forwarding the customer originated data packets to the appropriate destination

D. Forwarding the provider originated data packets to the appropriate customer site

E. Providing secure layer 3 routing exchange between sites

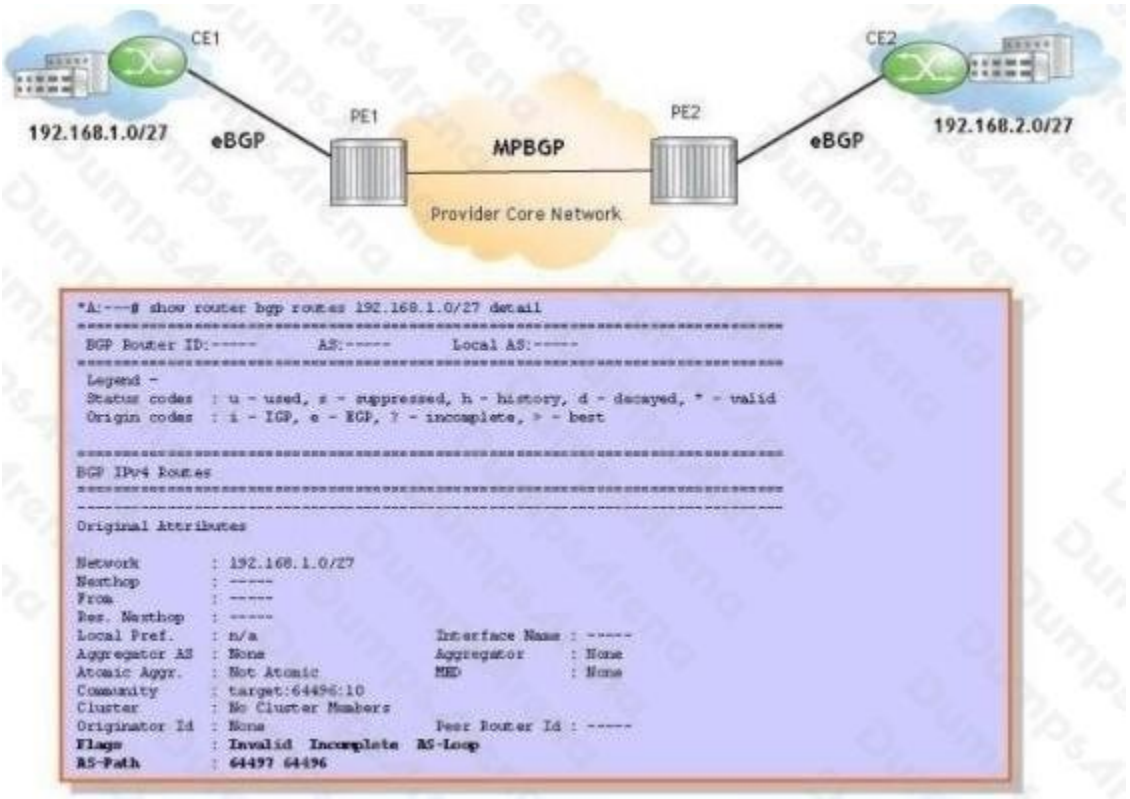
ANSWER: B C E

Explanation:

In a VPRN, the service provider operates the provider-edge functions that make multiple customer sites behave as though they are connected by a private routed network. Therefore, distributing the customer routing information between sites is a provider responsibility: PE routers import routes learned from customer-facing interfaces into the appropriate VPN routing and forwarding instance and distribute VPN routes across the provider MPLS/IP backbone, commonly using MP-BGP. The provider also forwards customer-originated data packets to the appropriate destination. PE routers apply the correct VPN label and transport label, and the provider core transports the packet to the egress PE, which forwards it into the destination customer site. Finally, providing secure layer 3 routing exchange between sites is fundamental to a VPRN service. VPN route targets and VRF separation ensure that routes and traffic are exchanged only among sites belonging to the intended customer VPN, preserving logical isolation while using shared provider infrastructure. These responsibilities align with the RFC definition of BGP/MPLS IP VPNs, in which the service provider supplies VPN route distribution, forwarding, and isolation across its backbone. See RFC 4364: BGP/MPLS IP Virtual Private Networks and Nokia's Layer 3 Services documentation.

QUESTION NO: 19

Click the exhibit.



According to the display, a BGP loop is detected for route 192.168.1.0/27. On which router is this command executed?

A. CE1

B. CE2

C. PE1

D. PE2

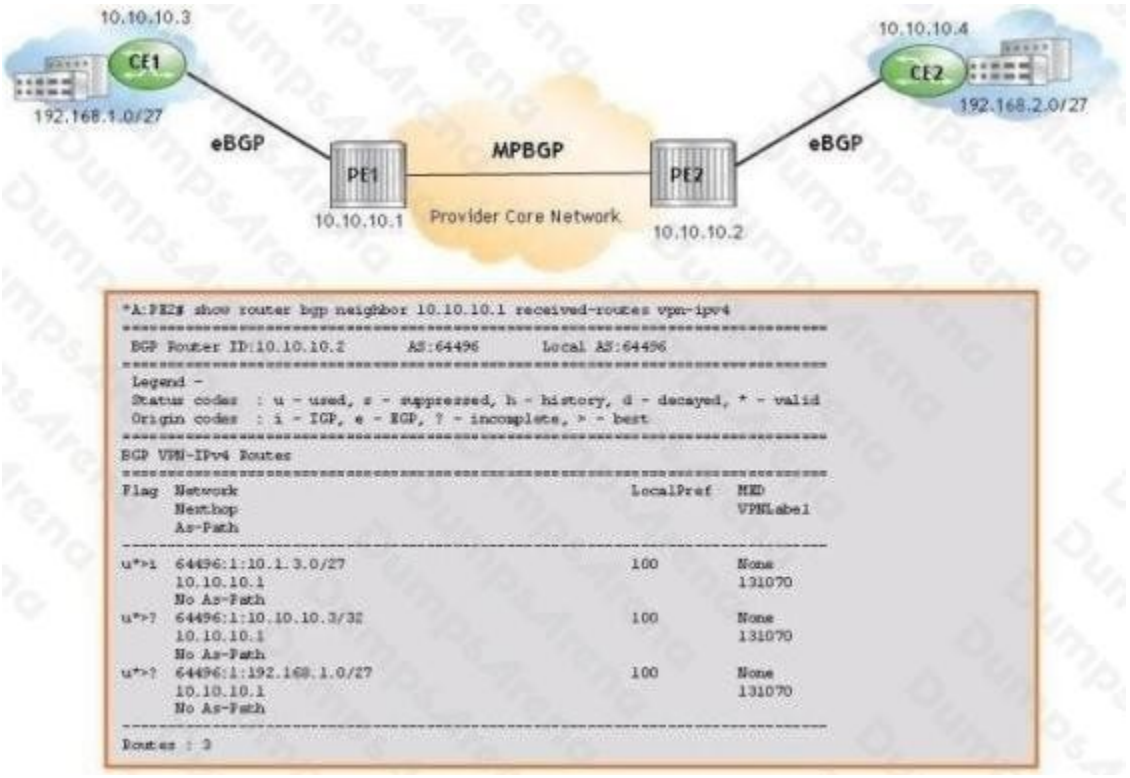
ANSWER: B

Explanation:

CE2 is correct. The displayed route status identifies an AS-path loop for 192.168.1.0/27, meaning that the router on which the command is run sees its own autonomous-system number in the received BGP AS path. BGP performs this check before accepting and installing a route, preventing the same advertisement from circulating back into an autonomous system and creating a persistent control-plane loop. In the illustrated path, the AS-path information shown for this prefix corresponds to the customer-edge router at the far end, CE2; therefore, CE2 is the device reporting the loop-detected condition. The route can be visible in BGP diagnostic output even though it is not eligible for normal route selection, because the output is specifically exposing the reason the received path is rejected. This behavior is a fundamental BGP loop-prevention mechanism: each eBGP speaker prepends its AS number when advertising a route, and a receiving speaker rejects a path containing its local AS. Nokia's BGP documentation describes AS-path handling and loop prevention in the SR OS BGP guide. General BGP path-vector and AS-loop behavior is also standardized in RFC 4271.

QUESTION NO: 21

Click the exhibit.




```
*A:PE2# show router bgp neighbor 10.10.10.1 received-routes vpn-ipv4
===============================================================================
 BGP Router ID:10.10.10.2       AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
Flag  Network                                    LocalPref   MED
      Nexthop                                                VPNLabel
      As-Path
-------------------------------------------------------------------------------
u*>i  64496:1:10.1.3.0/27                        100         None
      10.10.10.1                                           131070
      No As-Path
u*>?  64496:1:10.10.10.3/32                      100         None
      10.10.10.1                                           131070
      No As-Path
u*>?  64496:1:192.168.1.0/27                     100         None
      10.10.10.1                                           131070
      No As-Path
-------------------------------------------------------------------------------
Routes : 3
```

Which of the following best describes the output?

A. The AS-Path is not present because a policy is configured on PE2 to set it to null.

B. The AS-Path is not present because a policy is configured on CE1 to set it to null.

C. The AS-Path is not present because a policy is configured on PE1 to set it to null.

D. The AS-Path is not present because VPN-IPv4 routes do not propagate the IPv4 AS-Path BGP attribute.

ANSWER: C

Explanation:

The AS-Path is not present because a policy is configured on PE1 to set it to null. PE1 is the ingress provider-edge router for the customer route and is therefore the point at which routing policy can modify the BGP attributes before the route is advertised into the provider VPN route-distribution domain. A policy action that sets the AS path to null removes the AS-path information from the route advertisement. Consequently, the displayed VPN route has no AS-path value. The VPN-IPv4

address family itself does carry ordinary BGP path attributes, including AS_PATH when that attribute is present; the absence shown is the result of explicit policy handling at PE1 rather than an inherent limitation of VPN-IPv4. This behavior is consistent with SR OS BGP policy processing, where import/export policies may alter route attributes, and with the MPLS L3VPN control-plane model, which distributes VPN routes using Multiprotocol BGP. See Nokia's SR OS BGP documentation and RFC 4364: BGP/MPLS IP Virtual Private Networks.

QUESTION NO: 22

In a VPRN, how does the service provider ensure that customer addresses are unique when multiple customers use the same address space?

A. It is the customer responsibility to ensure their address space is unique

B. The Route Distinguisher ensures that addresses are unique between customers.

C. The Route Target ensures that addresses are unique between customers

D. The provider will advise the customers to change their addresses

E. The provider will use NAT to ensure that addresses are unique between customers

ANSWER: B

Explanation:

The Route Distinguisher ensures that addresses are unique between customers. In an MPLS Layer 3 VPN such as an Alcatel-Lucent/Nokia VPRN, customers may legitimately use identical IPv4 prefixes, including RFC 1918 private addressing. The provider edge router does not advertise the customer prefix by itself in the provider's VPN route distribution system. Instead, it combines the customer IPv4 prefix with an eight-octet Route Distinguisher to form a VPN-IPv4 address. For example, two separate VPRNs can both use 10.1.1.0/24, but assigning each VPRN a distinct Route Distinguisher creates distinct VPN-IPv4 routes. This preserves route uniqueness in MP-BGP while allowing the overlapping customer address plans to remain unchanged.

The Route Distinguisher identifies the VPN routing context for route uniqueness; it is not a customer-facing address translation mechanism. Route Targets are then used as extended BGP communities to control which VPN routes are imported and exported by VPRNs. This division enables scalable isolation and controlled route exchange across the provider MPLS network. The VPN-IPv4 construction and the role of the Route Distinguisher are specified in RFC 4364, BGP/MPLS IP Virtual Private Networks. Nokia's VPRN documentation also describes the use of route distinguishers and route targets for VPN route handling: Nokia SR OS VPRN overview.

QUESTION NO: 23

When an Nokia 7750 SR receives a BGP update containing an AS-Path loop, what is the default action?

A. The update is accepted and the route installed in the route table.

B. The update is accepted but the route will not be installed in the route table

C. The update is rejected

D. The peer session is dropped

ANSWER: C

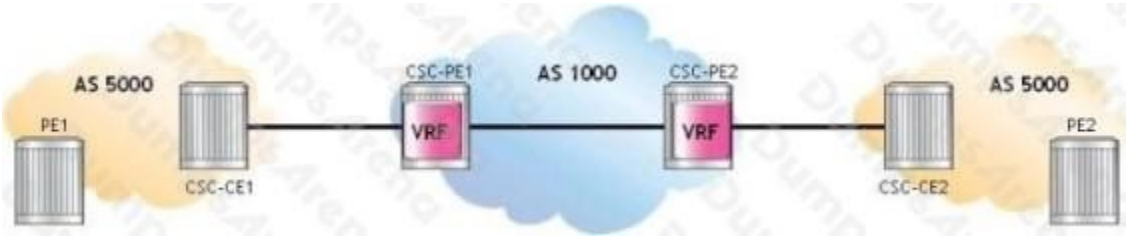
Explanation:

The update is rejected. BGP uses the AS_PATH attribute as a fundamental loop-prevention mechanism. When a Nokia 7750 SR receives an advertisement whose AS_PATH already contains its own autonomous-system number, accepting the route could cause traffic to be routed back toward an AS that has already propagated the route. Therefore, by default, the router treats the advertisement as an AS-path loop and does not accept it as a usable BGP route. This behavior applies during BGP UPDATE processing, before the route can be selected as a best path or installed into the routing table. The applicable

BGP standard specifies that a speaker that detects its own AS number in the AS_PATH attribute must not advertise or accept the route in a way that creates an inter-AS routing loop. Nokia SR OS BGP behavior implements this standard protection; configurations such as AS-path loop allowance may deliberately alter the behavior in specific designs, but they are not the default. See RFC 4271, BGP loop detection and the Nokia SR OS BGP documentation.

QUESTION NO: 25

Click the exhibit.



For the Carrier Supporting Carrier (CSC) VPRN, which routes are advertised between CSC-CE1 andCSC-PE1?

A. VPN-IPv4 routes for all PE system addresses in AS 1000

B. BGP labeled IPv4 routes for all PE system addresses in AS 1000

C. VPN-IPv4 routes for all PE system addresses in AS 5000

D. BGP labeled IPv4 routes for all PE system addresses in AS 5000

ANSWER: D

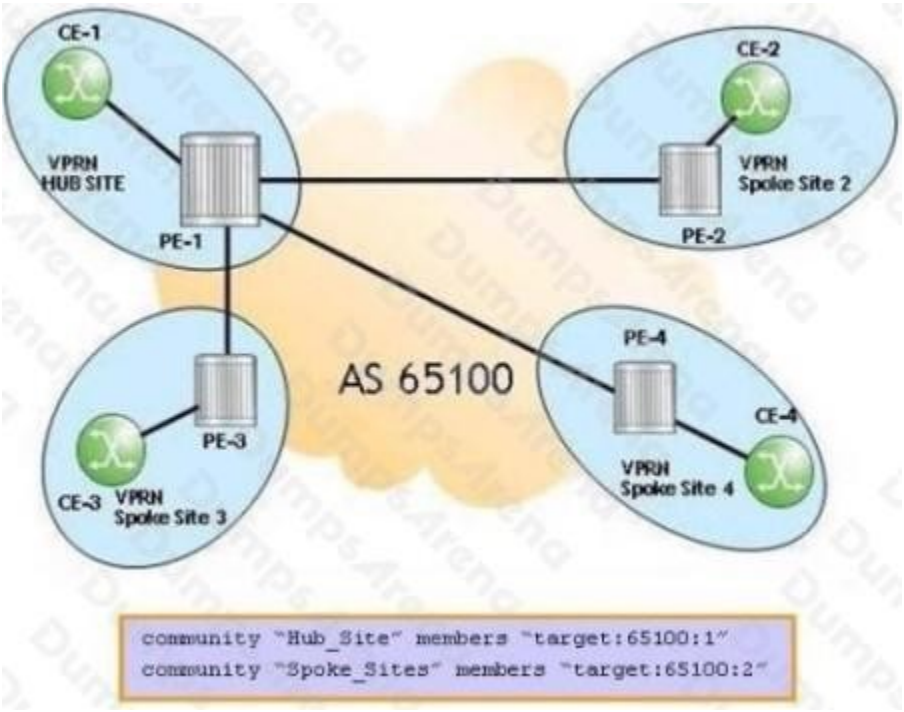
Explanation:

BGP labeled IPv4 routes for all PE system addresses in AS 5000 is correct. In a Carrier Supporting Carrier VPRN, the customer carrier (AS 5000 in the exhibit) requires labeled reachability to the loopback/system addresses of its own PE routers across the provider carrier network. CSC-CE1 therefore advertises its customer-carrier PE loopback routes to CSC-PE1 using BGP IPv4 labeled-unicast routes. Each route carries a BGP label allocated by the customer carrier, allowing the provider carrier to transport packets toward the appropriate customer PE while preserving the label operations required for the customer carrier's MPLS VPN services.

This exchange does not distribute the customer's VPN-IPv4 customer routes to the CSC provider edge. Instead, it establishes labeled IPv4 transport reachability for the customer carrier's PE system addresses. The CSC provider uses that reachability to carry the customer carrier's MPLS-encapsulated traffic through its VPRN. Nokia's CSC VPRN documentation describes the use of labeled routes between the CSC-CE and CSC-PE, while the BGP specification for labeled-unicast explains that labels are associated with IPv4 NLRI to provide MPLS forwarding information. See Nokia SR OS CSC VPRN documentation and RFC 8277: BGP MPLS Label Stack Encoding.

QUESTION NO: 26

Click the exhibit.



Community lists are configured for the PE hub and spoke VPRN. Which communities should be exported by the VPRN service on PE1?

A. "Spoke.Sites" and "Hub_Site"

B. "Spoke_Sites"

C. "Hub_Site"

D. No communities should be exported

ANSWER: C

Explanation:

"Hub_Site" is correct. In a hub-and-spoke VPRN design, BGP communities provide a route-role classification that the service import policies use to control connectivity. PE1 is the hub PE in the exhibit, so routes exported from its VPRN must be marked with the community identifying routes that originate from the hub site. The export community is attached when the PE advertises the route into the VPN routing domain; it allows spoke VPRNs to recognize and import routes learned from the hub. This classification is essential to the intended traffic model: spokes can reach the hub, while the configured policies can prevent direct spoke-to-spoke route distribution. The community attached on export describes the source role of the route, rather than all site roles that may later be allowed to import it. Nokia SR OS VPRN policy processing supports using BGP communities in service route import and export policy entries, including matching and adding communities for VPN route distribution. See Nokia's VPRN Guide and the BGP Guide for SR OS policy and community behavior.

QUESTION NO: 27

Which of the following statements are true regarding P devices in an MPLS VPRN? (Choose two.)

A. Participate in service provider core routing

B. P devices are not required to be MPLS enabled. MPLS is only required on the PE devices

C. Run a common routing protocol with the CE router

D. Must support MP-BGP

E. Do not have any connections to the CE

F. Must be aware of the VPRNs

ANSWER: A E

Explanation:

In an MPLS VPRN, P devices are provider-core routers. Therefore, they participate in service provider core routing: they use the provider IGP to establish reachability across the backbone and use MPLS label forwarding to transport traffic between PE routers. Their role is to provide scalable transit through the provider network rather than to hold customer VPN routing information. This separation is a fundamental MPLS L3VPN design principle, because it allows the core to forward labeled packets without maintaining per-VPRN route tables.

P devices also do not have any connections to the CE. Customer-edge routers attach to provider-edge routers, where VPRN service contexts, VRFs, customer-facing routing adjacencies, and VPN route import/export policies reside. A P router forwards the MPLS-encapsulated traffic received from the PE toward the egress PE based on the transport label, without participating in the CE-facing service edge.

RFC 4364 describes the PE/P/CE model and explains that provider backbone routers can forward VPN traffic without needing VPN-route knowledge: RFC 4364 — BGP/MPLS IP Virtual Private Networks. Nokia's service-router documentation provides MPLS and VPRN configuration and architecture material: Nokia Service Router documentation.

QUESTION NO: 28

Which of the following about BGP Site of Origin (SoO) is FALSE?

A. The SoO attribute can be used to uniquely identify the site from which a PE learns routes.

B. Before a PE redistributes a VPN-IPv4 route, it can assign an SoO attribute to the route.

C. The SoO attribute can be used to prevent loops from a single site with multiple CE-PE connections.

D. The SoO attribute can be used as matching criteria when redistributing routes from the CE to the PE.

ANSWER: D

Explanation:

The statement "The SoO attribute can be used as matching criteria when redistributing routes from the CE to the PE." is false because Site of Origin is applied by the provider edge router to routes it learns from a customer site. Its primary loop-prevention function is evaluated when a route is considered for advertisement *back toward* a customer edge router: if the route carries the same Site of Origin value as the receiving site, it must not be advertised to that site. This prevents a prefix learned through one attachment of a multihomed customer site from returning through another attachment and creating a routing loop. Thus, SoO is an origin marker placed on customer-learned routes and checked on delivery to the customer site, rather than a criterion for the initial redistribution direction from customer edge to provider edge. RFC 4364 describes the multihomed-site loop issue and the requirement to identify routes originating at the customer site; it also describes using a BGP extended community for this purpose. The encoding and transitive handling of BGP extended communities are defined in RFC 4360. See RFC 4364, Section 4.6.1 and RFC 4360.