PassQueen
a.k.a dumpsqueen.com

Implementing Cisco Service Provider VPN Services (300-515 SPVI)

Cisco 300-515

Version Demo

Total Demo Questions: 11

Total Premium Questions: 114

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Service Provider VPNs	14
Topic 2, Layer 2 VPNs	35
Topic 3, Layer 3 VPNs	42
Topic 4, Multicast VPNs	11
Topic 5, Segment Routing	2
Topic 6, Mix Questions	10
Total	114

QUESTION NO: 1

In a typical service provider environment, which two tools are used to help scale PE router connectivity requirements? (Choose two.)

A. route reflectors

B. VPNv4 address family

C. originator ID

D. cluster ID

E. confederations

ANSWER: A E

Explanation:

route reflectors and **confederations** are both BGP scaling mechanisms that reduce the number of internal BGP peerings required among PE routers. In an MPLS Layer 3 VPN service-provider network, PE routers must exchange VPN routing information through Multiprotocol BGP. A full iBGP mesh would require every PE to establish a session with every other PE, causing the number of BGP sessions to grow rapidly as PEs are added. Route reflectors centralize route advertisement: PE routers can peer with one or more route reflectors rather than with all other PEs. This substantially reduces the peer-session count while preserving distribution of VPNv4 or VPNv6 reachability. Route reflection behavior, including loop-prevention attributes used by reflectors, is standardized in RFC 4456.

Confederations provide another approach by partitioning a large public autonomous system into multiple internal subautonomous systems. Within each sub-AS, BGP peering and policy can be managed independently, while the overall network continues to appear externally as one public AS. This hierarchical design limits the scope of required iBGP connectivity and can make a very large provider BGP deployment more manageable. Confederation operation is specified in RFC 5065.

QUESTION NO: 2

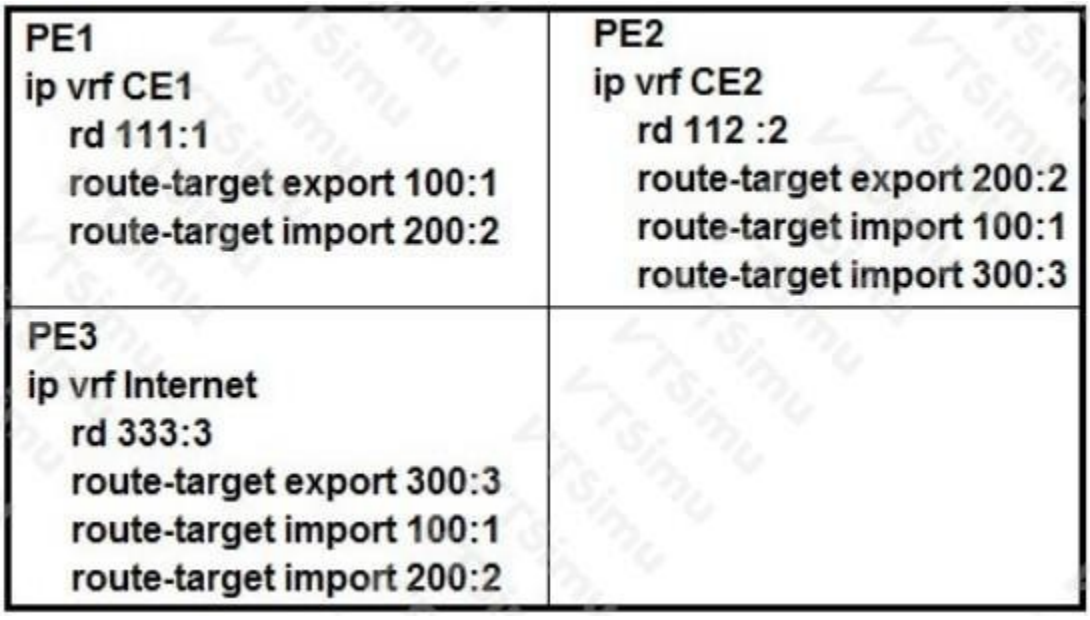

Refer to the exhibit. PE1 and PE2 are exchanging VPNv4 routes for CE1 and CE2, and PE3 contains the default route to the internet. If the three devices are operating normally, which two conclusions describe this configuration? (Choose two.)

A. The CE1 and CE2 VRFs can exchange routes only between their respective VRFs on PE1 and PE2.

B. All three routers must be running a distance-vector routing protocol.

C. All three routers must be running MP-BGP.

D. The CE1 and CE2 VRFs can access the default route provided by the Internet VRF.

E. Only the CE2 VRF can access the default route provided by the Internet VRF.

ANSWER: A C

Explanation:

The CE1 and CE2 VRFs can exchange routes only between their respective VRFs on PE1 and PE2 because VPN route targets determine which VRFs import VPNv4 prefixes. A PE exports a route with its configured export route target, and another PE installs that route into a VRF only when the VRF has a matching import route target. Thus, the shared CE route-target policy permits the CE1 and CE2 customer VRFs to exchange their VPN routes while retaining VRF-based separation from routes that do not carry the required import target.

All three routers must be running MP-BGP because VPNv4 route distribution requires Multiprotocol BGP support. MP-BGP carries the VPNv4 address family, consisting of an IPv4 prefix together with a route distinguisher, and advertises the associated VPN label and extended communities such as route targets. PE1 and PE2 use this capability to exchange the CE VPN routes, while PE3 must also participate in the VPNv4 control plane to advertise VPN routes, including any Internet-VRF route intended for import by another PE. Cisco documents VPNv4 address-family configuration and route-target import/export behavior in its MPLS Layer 3 VPN BGP documentation and VRF configuration documentation.

QUESTION NO: 3

What do routers on the network use to avoid routing loops when OSPF is running as the PE-CE routing protocol on a service provider network?

A. the AS-Override feature

B. the DN bit with type 3, 5, or 7 LSA

C. the domain tag for type 2 LSA

D. sham links to create a super backbone over the service provider network

ANSWER: B

Explanation:

The DN bit with type 3, 5, or 7 LSA is used to prevent routing loops in an MPLS Layer 3 VPN when OSPF operates between provider-edge and customer-edge routers. When a provider-edge router redistributes a VPN route into the customer's OSPF domain, it sets the Down (DN) bit in the applicable interarea-summary, external, or NSSA-external LSA. The DN bit identifies that the route originated from the MPLS VPN backbone and has already passed through the provider network.

If another provider-edge router receives such an LSA from the customer site, the DN-bit marking ensures that it does not redistribute that route back into the VPN routing-and-forwarding table and then into Multiprotocol BGP. This prevents a route from leaving the provider VPN, traversing a customer OSPF domain, and being injected back into the same VPN as though it were a newly learned customer route. The mechanism is especially important where a VPN has multiple provider-edge/customer-edge OSPF connections. Cisco documents the DN bit as the loop-prevention mechanism for OSPF PE-CE routing in MPLS VPNs. See Cisco's OSPF DN-bit overview and Cisco IOS OSPF MPLS VPN documentation.

QUESTION NO: 4 - (DRAG DROP)

DRAG DROP

Drag and drop the EVPN components from the left onto the correct planes on the right.

Select and Place:

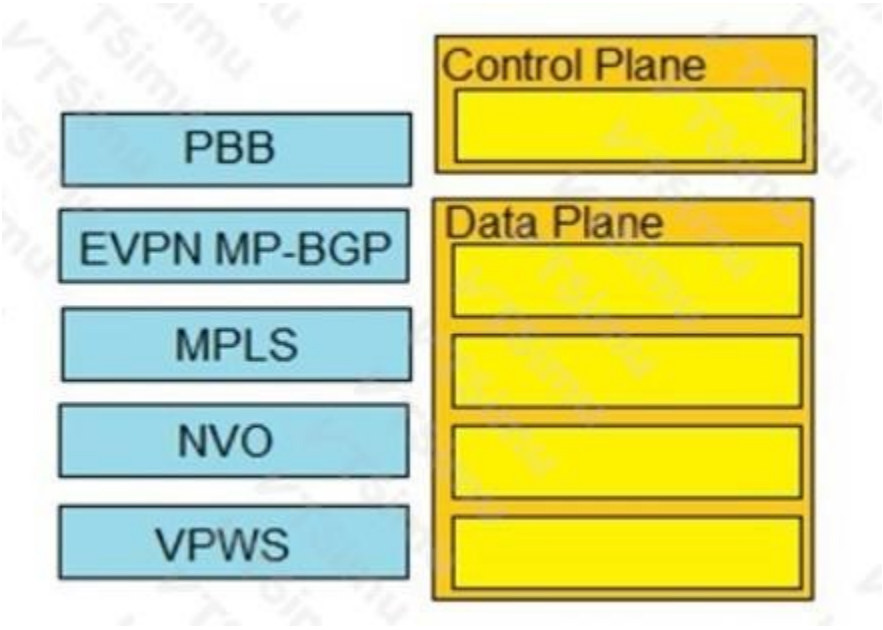

ANSWER:

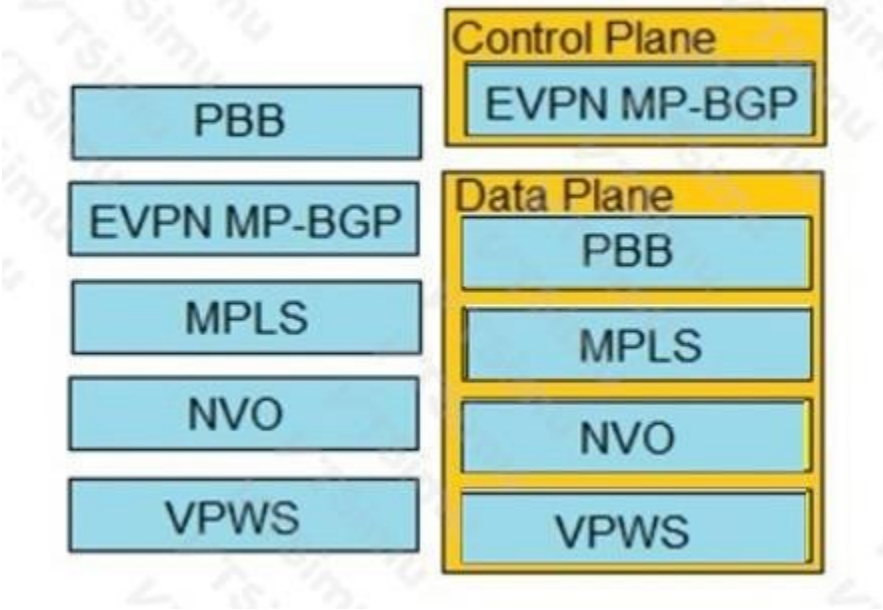

Explanation:

EVPN separates the distribution of reachability and service information from the forwarding technology used to transport customer traffic. EVPN MP-BGP belongs in the control plane because it is the BGP-based signaling mechanism that advertises EVPN Network Layer Reachability Information between provider edge devices. This information includes MAC/IP reachability, Ethernet segment information, and service membership details. The receiving devices use those BGP advertisements to build the EVPN forwarding state consistently and to support multihoming behavior.

PBB, MPLS, NVO, and VPWS belong in the data plane because they identify the forwarding or encapsulation environments in which EVPN services can operate. EVPN can use an MPLS transport data plane for labeled forwarding. It can also be used with network virtualization overlays, commonly represented by NVO-based overlay forwarding. Provider Backbone Bridging provides a data-plane framework for carrying Ethernet services, and VPWS represents point-to-point Layer 2 service forwarding. These technologies carry encapsulated user frames after the control plane has distributed the required reachability and service information.

The completed placement therefore correctly captures EVPN's key architectural idea: EVPN MP-BGP distributes control information, and the other listed technologies provide the traffic-forwarding context. Cisco's EVPN material describes EVPN as a BGP-based control plane that can operate over multiple data-plane technologies. See the Cisco reference material in Cisco's EVPN knowledge-network presentation and the overview in Cisco Ethernet VPN documentation.

QUESTION NO: 5

What must match in the EVPN and L2VPN configuration mode when configuring EVPN native in a router?

A. interface

B. address family

C. bridge domain

D. EVI

ANSWER: D

Explanation:

EVI must match between the EVPN configuration and the associated L2VPN bridge-domain configuration for a native EVPN service. The Ethernet VPN instance identifier is the service identifier that ties the EVPN control-plane definition to the Layer 2 forwarding construct. In native EVPN on Cisco IOS XR, the EVPN configuration defines the EVI and its BGP EVPN behavior, while the L2VPN bridge domain associates local attachment circuits and other Layer 2 service attributes with that same EVPN instance. Using the identical EVI ensures that MAC/IP route advertisements, inclusive multicast Ethernet tag routes, and other EVPN signaling are applied to the intended bridge domain.

The EVI is locally significant for correlating the EVPN instance with the L2VPN service configuration; it is not required to be globally unique across every provider-edge router. However, each provider-edge router participating in the same EVPN service must consistently map its local bridge domain to the appropriate service through its configured EVI and EVPN settings. Cisco's native EVPN configuration examples show the EVI under both the EVPN configuration hierarchy and the relevant L2VPN bridge-domain context. See Cisco's L2VPN and EVPN Configuration Guide and the Cisco IOS XR EVPN documentation.

QUESTION NO: 6

Which two BGP route types are used by an MVPN control plane to signal provider multicast tunnel information and customer multicast state? (Choose two.)

A. Intra-AS I-PMSI A-D route

B. C-multicast route

C. VPNv4 unicast route

D. EVPN MAC/IP Advertisement route

E. LDP Address message

ANSWER: A B

Explanation:

The **Intra-AS I-PMSI A-D route** advertises information about an inclusive provider multicast service interface tunnel. It allows PE routers in an MVPN to discover the provider multicast tunnel that can carry multicast traffic for the VPN.

The **C-multicast route** advertises customer multicast state, such as source and group information, so that the provider network can establish or select the required multicast distribution path. MVPN uses Multiprotocol BGP route types to distribute this information while retaining customer multicast separation between VRFs. MVPN BGP signaling is specified in RFC 6514.

QUESTION NO: 7

While configuring the VRF Selection feature, you get an error message after typing the below statement: Router(config)#no vrf selection source 172.16.0.0 255.255.0.0 vrf VRF1

Which action caused this message?

A. the entry of an inconsistent IP address and mask for VRF Selection

B. an attempt to configure a VRF instance on an interface that already has VRF Selection configured

C. an attempt to remove a VRF Selection entry that does not exist

D. an attempt to configure a VRF Selection table that does not exist

ANSWER: C

Explanation:

An attempt to remove a VRF Selection entry that does not exist is correct. In Cisco IOS, the `vrf selection source` command creates a source-based VRF-selection mapping, using the specified source prefix, mask, and VRF name as the identifying values for that mapping. Prefix-based matching is not sufficient for the negated form of the command: `no vrf selection source 172.16.0.0 255.255.0.0 vrf VRF1` is intended to delete an already configured entry with those exact parameters. If IOS cannot find that matching source-selection entry in the VRF Selection configuration, it reports an error because there is no configuration object to remove. Before issuing the removal command, verify the active configuration with `show running-config` and confirm both the configured source network and the associated VRF. Cisco's VRF Selection feature documentation describes source-address-based VRF selection and its configuration syntax, including the source prefix, mask, and VRF association. See the Cisco IOS VRF Selection documentation and the Cisco IOS IP Addressing and VRF command reference.

QUESTION NO: 8

Which pseudowire feature provides an optional sequence number field for detecting out-of-order packets?

A. control word

B. route distinguisher

C. router alert label

D. entropy label

ANSWER: A

Explanation:

The **control word** is correct. The pseudowire control word is an optional four-byte field inserted between the pseudowire label and the encapsulated payload. Depending on the pseudowire type and configuration, it can carry sequence-number information that enables the receiving PE to detect packet reordering.

The control word also helps avoid problems caused by payload fields being interpreted as MPLS-specific values by devices in the provider network. Both pseudowire endpoints must agree on control-word use because it changes the format of the encapsulated pseudowire payload. Pseudowire control-word procedures are described in RFC 4447.

QUESTION NO: 9

Which two statements about MPLS L3 VPN RDs are true? (Choose two.)

A. They enable EIGRP to use address families to separate traffic between IPv4 and VPNv4.

B. They are represented as 32-bit values

C. They are represented as 64-bit values.

D. They enable OSPF to import and export routes into the global routing table of a router.

E. They allow BGP to uniquely identify duplicate routes.

ANSWER: C E

Explanation:

An MPLS Layer 3 VPN route distinguisher is a 64-bit value that is prepended to a customer IPv4 prefix to create a 96-bit VPN-IPv4 address. The route distinguisher does not alter the customer's IP forwarding prefix itself; instead, it makes the VPN route globally unique in the service-provider BGP control plane. Its 64-bit structure can use formats based on an autonomous system number or an IPv4-address administrator field together with an assigned number, allowing the provider to allocate distinct values per VRF, site, or customer as needed.

The statement that route distinguishers are represented as 64-bit values is therefore correct. The statement that they allow BGP to uniquely identify duplicate routes is also correct: when different VPN customers use the same RFC 1918 or public IPv4 prefix, attaching different route distinguishers produces separate VPN-IPv4 NLRIs. Multiprotocol BGP can then advertise, retain, and select paths for those otherwise identical customer prefixes without treating them as the same global route. VPN route targets, rather than route distinguishers, control VPN route import and export policy between VRFs. See RFC 4364, section 4.2 and Cisco MPLS Layer 3 VPN overview.

QUESTION NO: 10

Which two statements describe EVPN all-active multihoming? (Choose two.)

A. Multiple PEs can forward traffic for the same Ethernet Segment.

B. Participating PEs advertise the same Ethernet Segment Identifier.

C. Only the designated forwarder can transmit known unicast traffic.

D. It requires a separate bridge domain for each multihomed link.

E. It prevents a CE from using link aggregation.

ANSWER: A B

Explanation:

In EVPN all-active multihoming, multiple PE devices can forward traffic simultaneously for the same Ethernet Segment. This allows the service to use multiple CE-to-PE links for forwarding while retaining redundancy if one PE or attachment circuit fails.

The participating PEs advertise the same Ethernet Segment Identifier (ESI) for the multihomed Ethernet Segment. The ESI identifies the shared CE attachment to the EVPN control plane and allows remote PEs to apply multihoming procedures such as aliasing and split-horizon filtering. All-active multihoming is commonly used with a CE device that supports link aggregation. EVPN multihoming procedures are defined in RFC 7432.

QUESTION NO: 11

Refer to the exhibit.

Which statement about this command is true?

A. It must be configured on each PE router to enable the PE routers to receive multicast traffic for this particular MVRF.

B. It is used to set the designated router on a link using PIM-SM.

C. It must be configured on the PE and CE router to enable MP-BGP to send labels for CSC.

D. It is used to set the router that will server as the root bridge for STP.

ANSWER: A

Explanation:

“It must be configured on each PE router to enable the PE routers to receive multicast traffic for this particular MVRF.” is correct. VRF-aware multicast requires multicast routing to be enabled in the relevant multicast VPN routing and forwarding instance on every provider-edge router that participates in forwarding multicast traffic for that customer VPN. This establishes the multicast-routing context for the VPN and allows the provider-edge router to create and maintain the multicast forwarding state associated with the VRF. In an mLDP-based multicast VPN deployment, that VRF multicast state is what enables customer multicast traffic to be received and carried across the provider multicast transport, with mLDP supplying the multipoint label-switched distribution tree used in the core. The configuration is therefore applied per participating provider-edge router and per applicable VRF, rather than being a single network-wide setting. Cisco’s multicast VPN and mLDP configuration material describes enabling multicast routing for the VRF as part of the required provider-edge setup before multicast flows can be transported for that VPN. See Cisco’s mLDP configuration guide and its VRF-aware multicast documentation.