DUMPSBOSS.

Nokia SR Linux EVPN and Data Center Interconnect

Nokia 4A0-D03

Version Demo

Total Demo Questions: 7

Total Premium Questions: 71

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

Which of the following statements about utilizing asymmetric routing in an L3 EVPN network is FALSE?

A. If a host interface's has two IP addresses, it must send out two separate EVPN route-type 2 updates for the MAC-VRF.

B. Each PE must have a full ARP table for all of the hosts in the L3 EVPN network.

C. All the MAC-VRFs connected to the L3 EVPN network must exist on each PE.

D. The ingress and egress PE routers will perform both MAC and IP forwarding.

ANSWER: D

Explanation:

The false statement is *The ingress and egress PE routers will perform both MAC and IP forwarding.* In asymmetric IRB routing, the ingress PE performs the Layer 3 lookup for traffic entering the EVPN service. Having selected the destination IP route, it resolves the destination host's MAC/IP binding and encapsulates the packet for EVPN transport in the destination MAC-VRF/VNI. Thus, the ingress PE performs both IP forwarding and the MAC-based forwarding operation needed to deliver the packet through the overlay. At the remote end, however, the egress PE decapsulates the EVPN packet and bridges it in the destination broadcast domain toward the destination host. Its role for that packet is Layer 2 MAC forwarding, rather than performing a second IP routing lookup. This separation of forwarding functions is the defining characteristic of asymmetric routing: routing is centralized at the ingress PE for the relevant traffic direction, while the egress PE provides local Ethernet delivery. The reverse flow is handled independently, with the PE receiving that return traffic acting as the ingress PE for that direction. EVPN MAC/IP advertisement behavior and integrated routing and bridging concepts are specified in RFC 7432 and further described for EVPN IP prefix and IRB procedures in RFC 9135.

QUESTION NO: 2

Consider the exhibit.

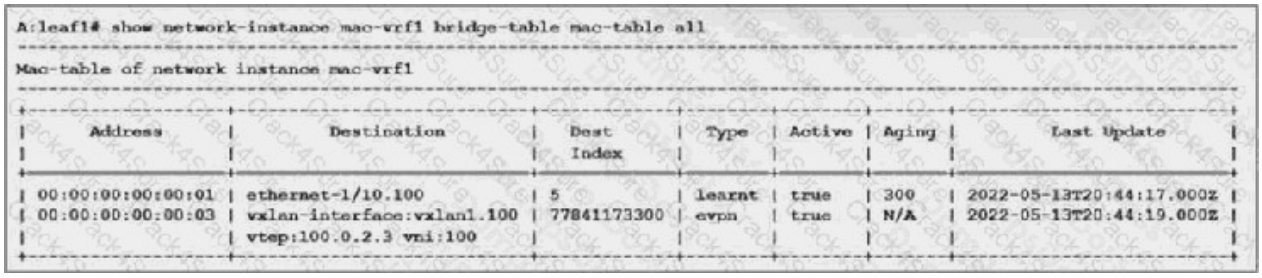

```
A:leaf1# show network-instance mac-vrf1 bridge-table mac-table all
-------------------------------------------------------------------------------------------------------------------
Mac-table of network instance mac-vrf1
-------------------------------------------------------------------------------------------------------------------
+-------------------+-------------------------------+-------------+---------+--------+--------+-------------------------------+
|      Address      |          Destination          |  Dest Index |  Type   | Active | Aging  |          Last Update          |
+===================+===============================+=============+=========+========+========+===============================+
| 00:00:00:00:00:01 | ethernet-1/10.100             | 5           | learnt  | true   | 300    | 2022-05-13T20:44:17.000Z      |
| 00:00:00:00:00:03 | vxlan-interface:vxlan1.100    | 77841173300 | evpn    | true   | N/A    | 2022-05-13T20:44:19.000Z      |
|                   | vtep:100.0.2.3 vni:100        |             |         |        |        |                               |
+-------------------+-------------------------------+-------------+---------+--------+--------+-------------------------------+
```

Which of the following statements is FALSE?

A. The remote host's IP address is 100.0.2.3.

B. The MAC address 00:00:00:00:00:01 was learned through the data plane.

C. The MAC address associated to the vxlan-interface will not age out.

D. The VXLAN network identifier configured for this MAC-VRF is 100.

ANSWER: A

Explanation:

The statement *"The remote host's IP address is 100.0.2.3."* is false because the IP address shown with a VXLAN-learned MAC entry identifies the remote VTEP used as the VXLAN tunnel endpoint, not the IP address of the endpoint host that owns the MAC address. In an EVPN-VXLAN MAC-VRF, the control plane distributes a remote MAC reachability advertisement. SR Linux installs that MAC in the forwarding database with a VXLAN destination and the remote VTEP address needed to encapsulate traffic toward the remote site. The VTEP address is therefore transport or overlay-tunnel information: it tells the local VTEP where to send VXLAN-encapsulated frames for that MAC. It does not reveal the attached host's IP address, which may be unknown to the MAC-VRF and can be unrelated to the VTEP loopback address. A MAC-

VRF is fundamentally a Layer 2 forwarding service, so it learns and distributes MAC reachability; host IP-to-MAC bindings are handled separately when IP route or ARP/ND EVPN route types are used. This distinction between a remote host and a remote VTEP is central to correctly interpreting SR Linux EVPN MAC-table output. See Nokia's SR Linux EVPN-VXLAN documentation and RFC 7348 for VXLAN tunnel endpoint behavior.

QUESTION NO: 3

An EVPN fabric uses iBGP route reflectors. Leaf1 originates a MAC/IP Advertisement route with its VTEP address as the BGP next hop, and the route reflector reflects the route to Leaf2.

Which statement is TRUE?

A. Leaf2 must have underlay reachability to Leaf1's VTEP address.

B. Leaf2 sends VXLAN traffic to the route reflector after receiving the reflected route.

C. The route reflector replaces Leaf1's VTEP address with its own address.

D. Leaf2 does not require underlay reachability to remote VTEP addresses.

ANSWER: A

Explanation:

The BGP next hop remains Leaf1's VTEP address when the route reflector reflects the iBGP EVPN route. Consequently, Leaf2 must have underlay IP reachability to Leaf1's VTEP address in order to encapsulate VXLAN traffic toward Leaf1.

A route reflector distributes EVPN control-plane routes but does not become the VXLAN data-plane endpoint for reflected routes. Route reflection also does not remove the need for VTEP loopback reachability in the underlay.

QUESTION NO: 4

A leaf router has learned a host MAC address remotely through an EVPN MAC/IP Advertisement route carrying a MAC mobility sequence number of 17. The same MAC address is then learned on a local access interface.

Which action should the leaf router take?

A. Advertise the local MAC address with a sequence number of 18.

B. Advertise the local MAC address with a sequence number of 17.

C. Advertise the local MAC address with a sequence number of 16.

D. Withdraw the remote MAC route without advertising a new route.

ANSWER: A

Explanation:

The leaf router must advertise the newly local MAC address with a MAC mobility sequence number greater than 17. By increasing the sequence number, it identifies its locally learned location as more recent than the previously received remote advertisement. Remote PEs can then update their forwarding state toward the new location.

Re-advertising sequence number 17 does not establish that the local learning is newer, and decrementing the sequence number would make the route less preferable. Withdrawing the previously received remote route does not distribute the new host location to the other PEs.

QUESTION NO: 5

Consider the exhibit.

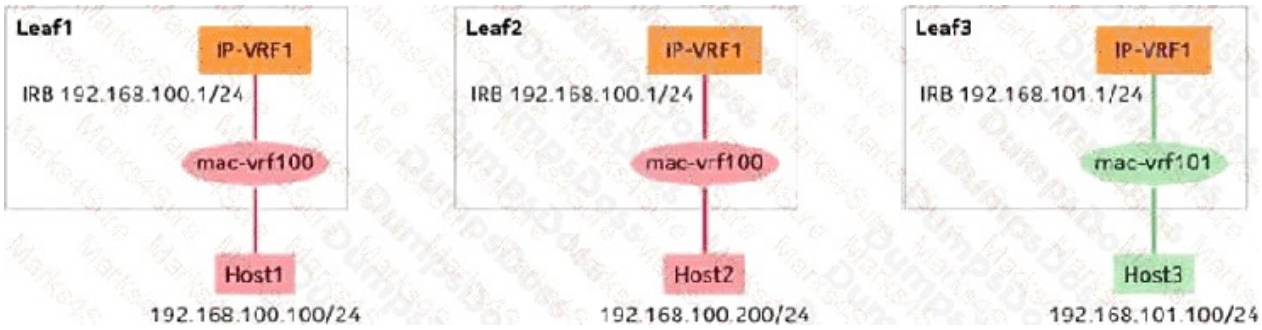

The network is configured for interface-less symmetric routing with an ECMP of 4 enabled on all leaf routers.

Which of the following statements is FALSE?

A. Anycast gateway is an optional configuration on the Leaf1 and Leaf2 IRB interfaces.

B. The IRB interfaces on Leaf1 and Leaf2 must be configured to advertise learned local host-routes.

C. MAC-VRF100 on Leaf1 and Leaf2 must have BGP EVPN configured, and IP-VRF1 must have BGP VPN configured.

D. IP-VRF1 on Leaf1, Leaf2 and Leaf3 must be configured with a VXLAN-routed interface.

ANSWER: A

Explanation:

“Anycast gateway is an optional configuration on the Leaf1 and Leaf2 IRB interfaces” is the false statement. In this interface-less symmetric-routing design, Leaf1 and Leaf2 both provide the distributed Layer 3 gateway for the same locally extended subnet. They therefore need a consistent default-gateway identity for attached hosts: the same gateway IP address and virtual gateway MAC address must be presented at each participating leaf. This is the anycast gateway function. It allows a host to use its closest leaf as the first-hop gateway while retaining identical gateway behavior if the host is connected through the other leaf or moves between leaves. The ingress leaf can then perform the required IRB lookup and symmetric VXLAN routing through the IP-VRF without requiring hosts to change their default-gateway configuration. ECMP provides multiple underlay paths between leaf and spine devices, but it does not replace the distributed gateway identity required at the IRB boundary. Nokia’s EVPN-VXLAN documentation describes EVPN’s distributed bridging and routing model, including gateway and IRB-related behavior; see the Nokia SR Linux documentation portal. The EVPN control-plane framework used for MAC/IP reachability is standardized in RFC 7432.

QUESTION NO: 6

Which of the following statements about utilizing VXLAN for the data plane in the data center is FALSE?

A. It allows the creation of a Layer 2 overlay network that can span the entire data center.

B. VXLAN was developed to support EVPN networks in the data center.

C. It may use ECMP to provide efficient utilization of the underlay interfaces within the data center.

D. It has the capability to isolate up to 16 million different overlays.

ANSWER: B

Explanation:

“VXLAN was developed to support EVPN networks in the data center.” is the false statement. VXLAN is a data-plane encapsulation mechanism that was defined to extend Layer 2 network segments across an IP network. It encapsulates an inner Ethernet frame in UDP/IP and identifies the logical overlay using a 24-bit VXLAN Network Identifier (VNI). This design enables large-scale multi-tenant segmentation and allows the IP underlay to use routed forwarding and equal-cost multipath forwarding.

EVPN is not the original purpose or prerequisite for VXLAN. EVPN is a standards-based BGP control plane that can distribute MAC, IP, multicast, and IP-prefix reachability information for a VXLAN overlay. Pairing EVPN with VXLAN is a widely deployed modern data-center architecture: VXLAN provides the forwarding encapsulation, while EVPN provides control-plane signaling and reduces reliance on data-plane MAC flooding. However, VXLAN was specified independently as an overlay encapsulation format, and EVPN can support multiple encapsulation technologies. RFC 7348 describes VXLAN and its VNI-based overlay model: RFC 7348. The EVPN BGP control-plane model is defined in RFC 7432.

QUESTION NO: 7

Consider the exhibit.

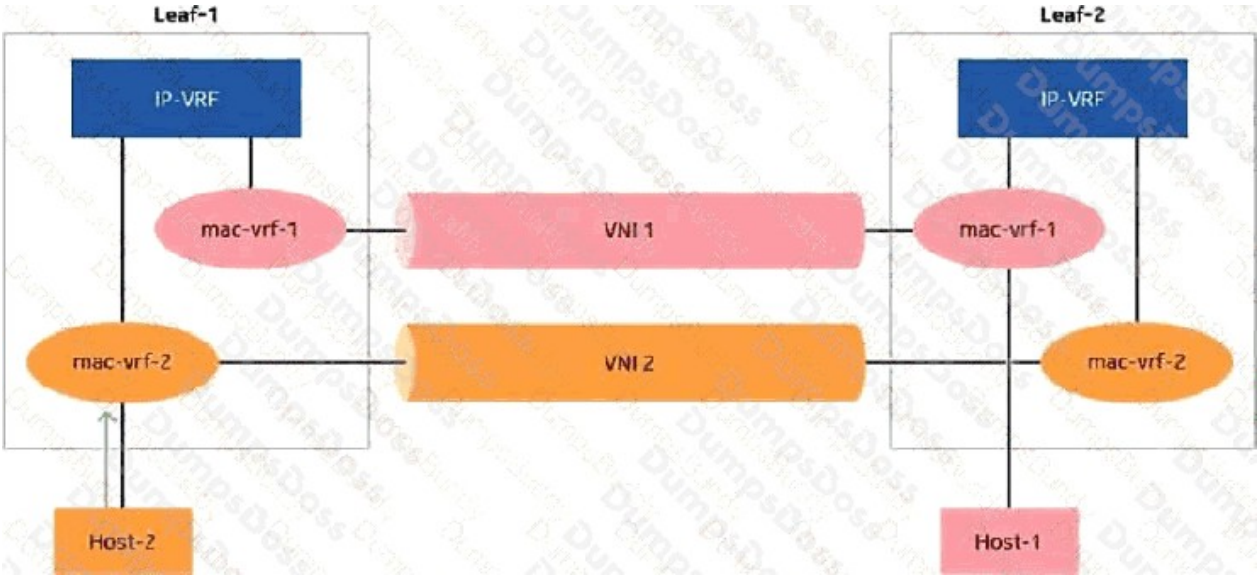

Host-2 is sending data to Host-1. The network is designed to use asymmetric routing.

Which of the following statements about the operation of the data plane is TRUE?

A. Leaf-1 forwards the data for mac-vrf-2 towards Leaf-2 with a VNI of 2.

B. Leaf-2 forwards the data for mac-vrf-2 towards the IP-VRF using its IRB interface.

C. Leaf-2 IP-VRF consults the ARP table to learn Host-1 MAC address.

D. Leaf-2 does a MAC look-up using mac-vrf-1 and forwards the data towards Host-1.

ANSWER: D

Explanation:

In asymmetric EVPN integrated routing and bridging, inter-subnet traffic is routed only at the ingress leaf. For traffic originating at Host-2 and destined for Host-1, the ingress leaf uses the IP-VRF routing information to select the destination subnet and resolves the destination endpoint through EVPN. It then encapsulates the packet in VXLAN using the VNI associated with the destination MAC-VRF. The packet is therefore delivered across the overlay in the bridging context of Host-1's subnet rather than being routed again at the egress leaf.

When Leaf-2 receives the decapsulated traffic, it acts as the egress VTEP and performs ordinary Layer 2 forwarding for the destination segment. Consequently, it looks up Host-1's destination MAC address in `mac-vrf-1`, identifies the local attachment/interface for Host-1, and forwards the frame to that host. This split of responsibilities—ingress IP routing followed by egress destination-MAC bridging—is the defining data-plane behavior of asymmetric IRB. EVPN's IP-prefix and MAC/IP route procedures supporting inter-subnet forwarding are specified in RFC 9135; the VXLAN encapsulation model is specified in RFC 7348.