Number: JN0-683
Passing Score: 800
Time Limit: 120
File Version: 3.0

Exam Code: JN0-683

Exam Name: Data Center Professional

Exam A

QUESTION 1
Exhibit.

Exhibit

```
user@device> show configuration routing-instances
Customer_B {
    instance-type vrf;
    routing-options {
        graceful-restart;
        multipath;
        auto-export;
    }
    protocols {
        evpn {
            irb-symmetric-routing {
                vni 10006;
            }
            ip-prefix-routes {
                advertise direct-nexthop;
                encapsulation vxlan;
                vni 10006;
                export export_policy;
            }
        }
    }
    interface irb.400;
    interface irb.800;
    interface lo0.3;
    route-distinguisher 172.16.0.2:20;
    vrf-target target:10006:1;
}
Customer_A {
    instance-type vrf;
    routing-options {
        graceful-restart;
        multipath;
        evpn {
            irb-symmetric-routing {
                vni 10000;
            }
            ip-prefix-routes {
                advertise direct-nexthop;
    instance-type vrf;
    routing-options {
        graceful-restart;
        multipath;
        auto-export;
    }
    protocols {
        evpn {
            irb-symmetric-routing {
                vni 10000;
            }
            ip-prefix-routes {
                advertise direct-nexthop;
                encapsulation vxlan;
                vni 10000;
                export export_policy;
```

Referring to the configuration shown in the exhibit, assume that there is no external router present, and that the configuration is fabric-only.
Which two statements are true about the example configuration? (Choose two.)

A. VNI 10006 is assigned to vlan 800 (irb.800).
B. Devices in irb.400 (vlan 400) are not able to communicate directly with devices in routing instance Customer A.
C. Devices in routing instance Customer A are able to communicate with devices in routing instance Customer B
D. Devices in irb.400 (vlan 400) and irb.800 (vlan 800) are able to communicate over the fabric.

Correct Answer: B, D
Section:
Explanation:
Understanding the Configuration:
The exhibit shows configurations for two VRFs (Customer_A and Customer_B) with specific VLANs and VNIs assigned. Each VRF has interfaces (IRBs) associated with particular VLANs.
Communication Between VLANs and Routing Instances:
Option B: VLAN 400 (irb.400) is part of Customer_B, and there is no direct connection or routing between Customer_A and Customer_B in the configuration provided. Therefore, devices in irb.400 cannot communicate directly with devices in the Customer_A routing instance.
Option D: Since irb.400 (VLAN 400) and irb.800 (VLAN 800) are part of the same routing instance (Customer_B), they can communicate over the fabric using VXLAN encapsulation.
Conclusion:
Option B: Correct---There is no direct communication between devices in irb.400 (Customer_B) and routing instance Customer_A.
Option D: Correct---Devices in VLAN 400 and VLAN 800 can communicate within the Customer_B routing instance over the fabric.

QUESTION 2
Exhibit.

Exhibit

```
user@leaf1> show evpn database
Instance: evpn-1
VLAN  DomainId  MAC address        Active source                  Timestamp        IP address
     10001      00:1c:73:00:00:01  irb.4000                       Apr 16 11:46:14  10.4.4.1
     10001      40:00:dc:01:00:01  00:02:00:00:00:00:04:00:00:04  Apr 16 11:46:14  10.4.4.2
     10001      40:00:dc:01:00:02  00:02:00:00:00:00:04:00:00:04  Apr 16 11:46:14  10.4.4.3
     10001      40:00:dc:01:00:03  00:02:00:00:00:00:04:00:00:04  Apr 16 11:46:14  10.4.4.4
     10001      40:00:dc:01:00:04  00:02:00:00:00:00:04:00:00:04  Apr 16 11:46:14  10.4.4.5
     10001      40:00:dc:01:00:05  00:02:00:00:00:00:04:00:00:04  Apr 16 11:46:14  10.4.4.6
     10001      44:11:01:00:00:01  00:02:00:00:00:00:04:00:00:04  Apr 16 11:46:14
     10001      44:11:01:00:00:02  00:02:00:00:00:00:04:00:00:04  Apr 16 11:46:14
     10001      44:11:01:00:00:03  00:02:00:00:00:00:04:00:00:04  Apr 16 11:46:14
     10001      44:11:01:00:00:04  00:02:00:00:00:00:04:00:00:04  Apr 16 11:46:14
     10001      44:11:01:00:00:05  00:02:00:00:00:00:04:00:00:04  Apr 16 11:46:14
     10001      44:12:01:00:00:01  00:02:00:00:00:00:03:00:00:03  Apr 16 11:46:14
     10001      44:12:01:00:00:02  00:02:00:00:00:00:03:00:00:03  Apr 16 11:46:14
     10001      44:12:01:00:00:03  00:02:00:00:00:00:03:00:00:03  Apr 16 11:46:14
     10001      44:12:01:00:00:04  00:02:00:00:00:00:03:00:00:03  Apr 16 11:46:14
     10001      44:12:01:00:00:05  00:02:00:00:00:00:03:00:00:03  Apr 16 11:46:14
     10002      00:1c:73:00:00:01  irb.300                        Apr 16 11:46:14  10.3.3.1
     10002      30:00:dc:01:00:01  00:02:00:00:00:00:01:00:00:01  Apr 16 11:46:14
     10002      30:00:dc:01:00:02  00:02:00:00:00:00:01:00:00:01  Apr 16 11:46:14
     10002      30:00:dc:01:00:03  00:02:00:00:00:00:01:00:00:01  Apr 16 11:46:14
     10002      30:00:dc:01:00:04  00:02:00:00:00:00:01:00:00:01  Apr 16 11:46:14
```

The exhibit shows the truncated output of the show evpn database command.
Given this output, which two statements are correct about the host with MAC address 40:00:dc:01:00:04? (Choose two.)

A. The host is assigned IP address 10.4.4.5.
B. The host is originating from irb.300.
C. The host is located on VN110002.
D. The host is originating from an ESI LAG.

Correct Answer: A, D
Section:
Explanation:
Understanding the Output:
The show evpn database command output shows the MAC address, VLAN, active source, timestamp, and IP address associated with various hosts in the EVPN instance.
Analysis of the MAC Address:
Option A: The MAC address 40:00:dc:01:00:04 is associated with the IP address 10.4.4.5, as indicated by the output in the IP address column. This confirms that this host has been assigned the IP 10.4.4.5.
Option D: The active source for the MAC address 40:00:dc:01:00:04 is listed as 00:02:00:00:00:04:00:04:00:00:04:00:04, which indicates that the host is connected via an ESI (Ethernet Segment Identifier) LAG (Link Aggregation Group). This setup is typically used in multi-homing scenarios to provide redundancy and load balancing across multiple physical links.
Conclusion:
Option A: Correct---The host with MAC 40:00:dc:01:00:04 is assigned IP 10.4.4.5.
Option D: Correct---The host is originating from an ESI LAG, as indicated by the active source value.

QUESTION 3
Exhibit.

Exhibit

```
user@leaf1> show ethernet-switching vxlan-tunnel-end-point remote
Logical System Name       Id  SVTEP-IP         IFL   L3-Idx    SVTEP-Mode    ELP-SVTEP-IP
                  0   192.168.100.11   lo0.0    0
 RVTEP-IP         L2-RTT                  IFL-Idx   Interface    NH-Id   RVTEP-Mode  ELP-IP
       Flags
 192.168.100.13   default-switch          571       vtep.32769   1758    RNVE
    VNID          MC-Group-IP
    5010          0.0.0.0
    5020          0.0.0.0
user@leaf1> show interfaces vtep.32769
  Logical interface vtep.32769 (Index 571) (SNMP ifIndex 534)
    Flags: Up SNMP-Traps Encapsulation: ENET2
    VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 192.168.100.13, L2 Routing Instance:
default-switch, L3 Routing Instance: default
    Input packets : 0
    Output packets: 19
 ...
user@leaf1> show evpn database
Instance: default-switch
VLAN  DomainId  MAC address        Active source                  Timestamp        IP address
     5010       00:00:5e:00:01:01  05:00:00:fd:e9:00:00:13:92:00  Apr 15 22:27:02  10.1.1.254
     5010       00:0c:29:e8:b7:39  xe-0/0/4.0                     Apr 15 19:41:27  10.1.1.1
     5010       02:05:86:a7:4c:00  irb.10                         Apr 15 18:50:45  10.1.1.101
     5020       00:00:5e:00:01:01  05:00:00:fd:e9:00:00:13:9c:00  Apr 15 22:26:51  10.1.2.254
     5020       00:0c:29:08:04:a0  192.168.100.13                 Apr 15 23:07:22  10.1.2.1
     5020       02:05:86:a7:4c:00  irb.20                         Apr 15 22:26:51  10.1.2.101
user@leaf1> show route table bgp.evpn.0 evpn-mac-address 00:0c:29:08:04:a0
bgp.evpn.0: 28 destinations, 42 routes (28 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
2:192.168.100.13:1::5020::00:0c:29:08:04:a0/304 MAC/IP
                   *[BGP/170] 00:49:55, localpref 100, from 192.168.100.1
                      AS path: I, validation-state: unverified
                    >  to 172.16.1.0 via xe-0/0/0.0
                       to 172.16.1.6 via xe-0/0/1.0
user@leaf1> show route forwarding-table matching 10.1.2.1
...
Destination        Type RtRef Next hop           Type Index    NhRef Netif
10.1.2.1/32        dest     0 0:c:29:8:4:a0      ucst     1775     1 vtep.32769
```

Referring to the exhibit, Host1 (10.1.1.1) is failing to communicate with Host2 (10.1.2.1) in a data center that uses an ERB architecture. What do you determine from the output?

A. The traffic is failing because load balancing is not configured correctly.
B. The traffic is entering the VXLAN tunnel.
C. Host1 and Host2 are directly connected to leaf1.
D. The irb.20 interface is not configured on leaf1.

Correct Answer: B
Section:
Explanation:

Understanding the Problem:
Host1 (10.1.1.1) is failing to communicate with Host2 (10.1.2.1) within an EVPN-VXLAN environment using ERB architecture.
Analysis of the Exhibit:
The provided output includes information from the show route forwarding-table matching command for IP 10.1.2.1. The next hop is shown as vtep.32769, which indicates that the traffic destined for 10.1.2.1 is being forwarded into the VXLAN tunnel with the correct VTEP (VXLAN Tunnel Endpoint).
Conclusion:
Option B: Correct---The traffic from Host1 is entering the VXLAN tunnel, as evidenced by the next hop pointing to a VTEP. However, the issue could lie elsewhere, possibly with the remote VTEP, routing configurations, or the receiving leaf/spine devices.

QUESTION 4
Exhibit.

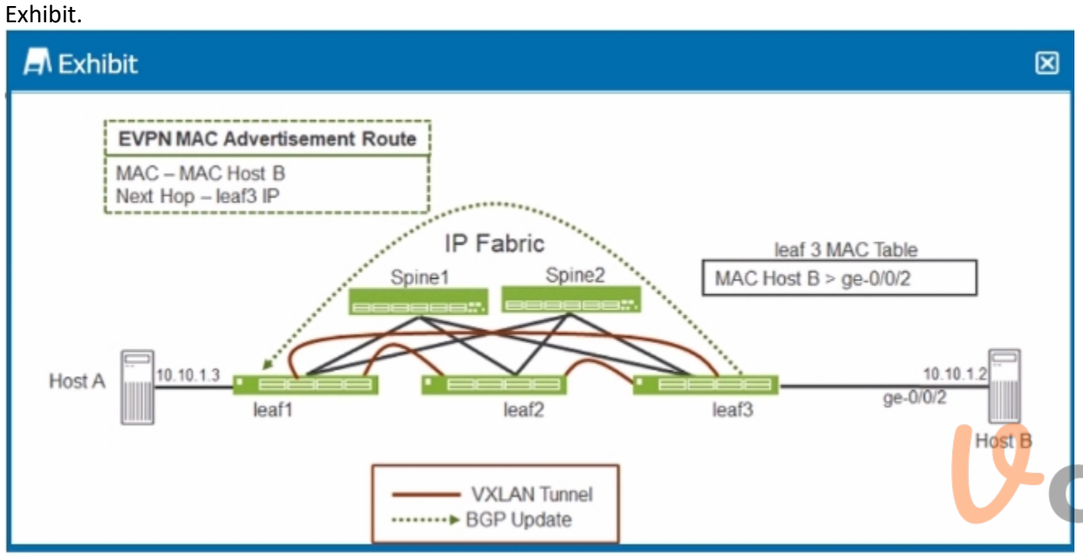

Referring to the exhibit, when Host A sends an ARP request for Host B's IP address, which Junos feature does leaf1 require to send an ARP response back to Host A without having to send a broadcast frame over the fabric?

A. proxy ARP
B. proxy NDP
C. GARP
D. DAD

Correct Answer: A
Section:
Explanation:
Scenario Overview:
In the exhibit, Host A is trying to resolve Host B's IP address (10.10.1.2) through ARP (Address Resolution Protocol). Normally, an ARP request would be broadcasted over the network, and the host owning the IP address (Host B) would respond.
Role of Proxy ARP:
Option A: Proxy ARP allows a router or switch (in this case, leaf1) to respond to ARP requests on behalf of another host. Leaf1, knowing the MAC address of Host B through the EVPN MAC advertisement, can reply to Host A's ARP request directly without broadcasting the request across the entire network fabric. This feature reduces unnecessary traffic and increases network efficiency.
Conclusion:
Option A: Correct---Proxy ARP enables leaf1 to respond to Host A's ARP request for Host B's IP without broadcasting over the IP fabric, thus providing the ARP response locally.

QUESTION 5
You are deploying a Clos IP fabric with an oversubscription ratio of 3:1.
In this scenario, which two statements are correct? (Choose two.)

A. The oversubscription ratio remains the same when you remove spine devices.
B. The oversubscription ratio decreases when you add spine devices.
C. The oversubscription ratio increases when you remove spine devices.
D. The oversubscription ratio remains the same when you add spine devices.

Correct Answer: B, C
Section:
Explanation:
Understanding Oversubscription in a Clos Fabric:
The oversubscription ratio in a Clos IP fabric measures the ratio of the amount of edge (leaf) bandwidth to the core (spine) bandwidth. An oversubscription ratio of 3:1 means that there is three times more edge bandwidth compared to core bandwidth.
Impact of Adding/Removing Spine Devices:
Option C: If you remove spine devices, the total available core bandwidth decreases, while the edge bandwidth remains the same. This results in an increase in the oversubscription ratio because there is now less core bandwidth to handle the same amount of edge traffic.
Option B: Conversely, if you add spine devices, the total core bandwidth increases. This decreases the oversubscription ratio because more core bandwidth is available to handle the edge traffic.
Conclusion:
Option C: Correct---Removing spine devices increases the oversubscription ratio.
Option B: Correct---Adding spine devices decreases the oversubscription ratio.

QUESTION 6
Exhibit.

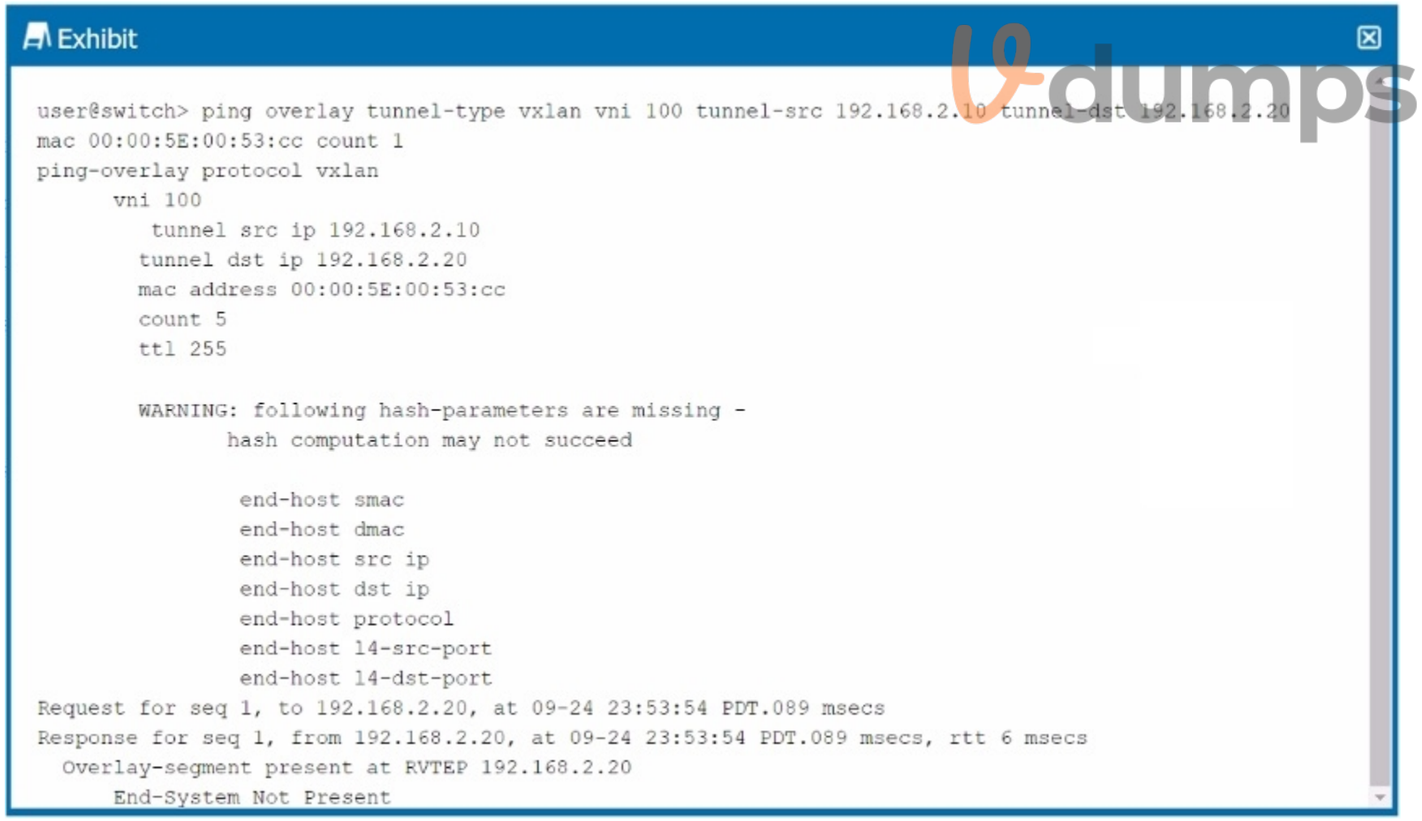

Referring to the exhibit, which statement is correct?

A. VNI 100 is not configured on the remote VTEP.
B. The MAC address is unknown and not in the forwarding table of the remote VTEP.
C. The remote VTEP is not responding.
D. The MAC address is known but not reachable by the remote VTEP

Correct Answer: B
Section:
Explanation:
Analyzing the Exhibit Output:
The command ping overlay tunnel-type vxlan is used to test the VXLAN tunnel between two VTEPs (VXLAN Tunnel Endpoints). The output shows a warning about missing hash parameters, but more importantly, it displays the result: End-System Not Present.
Understanding the Response:
The message End-System Not Present indicates that the remote VTEP (192.168.2.20) did not find the MAC address 00:00:5E:00:53:CC in its forwarding table. This typically means that the MAC address is unknown to the remote VTEP, and as a result, it could not forward the packet to the intended destination.
Conclusion:
Option B: Correct---The MAC address is unknown and is not in the forwarding table of the remote VTEP, which is why the system reports that the 'End-System' is not present.

QUESTION 7
You are asked to interconnect Iwo data centers using a method that provides EVPN Type 2 connectivity, is highly scalable, and limits VXLAN tunnels between border leaf devices. What will satisfy these requirements?

A. over the top full-mesh interconnect
B. EVPN Type 2 stretch
C. IP VPN
D. Type 2 seamless stitching

Correct Answer: D
Section:
Explanation:
Requirement Analysis:
The scenario requires a solution to interconnect two data centers that supports EVPN Type 2 connectivity. The solution must be highly scalable and must minimize the number of VXLAN tunnels between border leaf devices.
Understanding Type 2 Seamless Stitching:
Option D: Type 2 seamless stitching is a method used in EVPN to provide Layer 2 connectivity (such as MAC address mobility) across different VXLAN segments. It is scalable because it allows only necessary tunnels to be established between border leaf devices, reducing the overhead of maintaining a full mesh of VXLAN tunnels.
Conclusion:
Option D: Correct---Type 2 seamless stitching satisfies the requirement by enabling scalable, efficient interconnection of two data centers with minimal VXLAN tunnels.

QUESTION 8
Exhibit.

Exhibit

```
QFX10k-1
routing-instances {
    EVPN-VXLAN {
        instance-type vrf;
        interface irb.100;
        interface lo0.1;
        route-distinguisher 10.10.10.70:5000;
        vrf-target target:300:5000;
        protocols {
            evpn {
                ip-prefix-routes {
                    advertise direct-nexthop;
                    encapsulation vxlan;
                    vni 5000;
                }
            }
        }
    }
}
QFX10k-2
routing-instances {
    EVPN-VXLAN {
        instance-type vrf;
        interface irb.400;
        interface lo0.1;
        route-distinguisher 10.10.10.26:5000;
        vrf-target target:300:5000;
        protocols {
            evpn {
                ip-prefix-routes {
                    advertise direct-nexthop;
                    encapsulation vxlan;
                    vni 5000;
                }
            }
        }
    }
}
```


You have a sample configuration for connecting two sites through EVPN-VXLAN by exchanging IP prefix routes.
Referring to the exhibit, which two statements regarding the configuration are true? {Choose two.)

A. The advertise direct-nexthop option enables the receiver to resolve the next-hop route using only information carried in the Type 5 route.
B. The advertise direct-nexthop option enables the receiver to resolve the next-hop route using only information carried in the Type 2 route.
C. The VNI must match on all devices for the same customer.
D. The VNI should be unique on all devices for each customer site.

Correct Answer: A, C
Section:

Explanation:
EVPN-VXLAN Configuration:
The configuration provided in the exhibit shows an EVPN-VXLAN setup where IP prefix routes are exchanged between two sites. The advertise direct-nexthop option and the VNI (Virtual Network Identifier) settings are crucial in this context.
Advertise Direct-Nexthop:
Option A: The advertise direct-nexthop option ensures that the next-hop route is resolved using only the information carried in the EVPN Type 5 route. Type 5 routes are used for IP prefix advertisement in EVPN, which is key to enabling Layer 3 interconnectivity between different VXLAN segments.
VNI Consistency:
Option C: For the same customer across different devices, the VNI must be consistent. This consistency ensures that all devices can correctly map traffic to the appropriate VXLAN segment, maintaining seamless Layer 2 and Layer 3 connectivity.

QUESTION 9
Exhibit.

Exhibit

```
{master:0}[edit]
user@leaf1# show policy-options
...
policy-statement load-balance {
    term 1 {
        then {
            load-balance per-packet;
        }
    }
}
{master:0}[edit]
user@leaf1# show routing-options
router-id 192.168.100.11;
autonomous-system 65100;
{master:0}[edit]
user@leaf1# show protocols
bgp {
    group spine {
        type external;
        export direct;
        local-as 65003;
        multipath {
            multiple-as;
        }
        neighbor 172.16.1.5 {
            peer-as 65001;
        }
        neighbor 172.16.1.17 {
            peer-as 65002;
        }
    }
}
```


You are troubleshooting an IP fabric (or your data center. You notice that your traffic is not being load balanced to your spine devices from your leaf devices. Referring to the configuration shown in the exhibit, what must be configured to solve this issue?

A. The load-balance policy must be applied to the forwarding table under the routing-options hierarchy.
B. The multipast multiple -as configuration must be configured for each peer in the BGP spine group.
C. The load-balance policy must be applied as an export policy to your BGP
D. The load-balance policy must have a from statement that matches on protocol bgp.

Correct Answer: B
Section:
Explanation:
IP Fabric Load Balancing:
In the provided configuration, traffic is not being load-balanced to the spine devices. The issue likely relates to how BGP routes are being selected and whether Equal-Cost Multi-Path (ECMP) is functioning correctly.
Multipath Multiple-AS:
Option B: The multipath multiple-as configuration is essential when using BGP in an IP fabric where devices belong to different Autonomous Systems (AS). This setting allows BGP to consider multiple paths (even across different AS numbers) as equal cost, enabling ECMP and proper load balancing across spine devices.
Conclusion:
Option B: Correct---The multipath multiple-as configuration is necessary for achieving ECMP and effective load balancing in a multi-AS BGP environment.

QUESTION 10
You are implementing seamless stitching between two data centers and have a proposed configuration for a border leaf device.
In this scenario, which two statements are correct? {Choose two.)

A. The translation-vni must match in both data centers.
B. The translation-vni must be different in each data center.
C. The ESI must be different in each data center.
D. The ESI must match in both data centers.

Correct Answer: B, D
Section:
Explanation:
Understanding Seamless Stitching:
Seamless stitching is used in EVPN to interconnect two data centers, allowing for consistent Layer 2 and Layer 3 connectivity across them. This is often achieved by translating VNIs (Virtual Network Identifiers) between the data centers.
Translation-VNI:
Option B: The translation VNI must be different in each data center to ensure that traffic can be correctly routed and distinguished as it crosses between the data centers. This differentiation helps to maintain the integrity of the traffic flows and prevents any potential overlap or conflict in VNIs.
Ethernet Segment Identifier (ESI):
Option D: The ESI must match in both data centers to ensure that the same Ethernet segment (which could be multihomed) is recognized consistently across the data centers. Matching ESIs are crucial for maintaining a unified view of the Ethernet segment across the interconnected fabric.
Conclusion:
Option B: Correct---Translation VNIs must be unique to each data center for proper traffic distinction.
Option D: Correct---Matching ESIs are necessary to maintain consistent Ethernet segment identification across both data centers.

QUESTION 11
Exhibit.

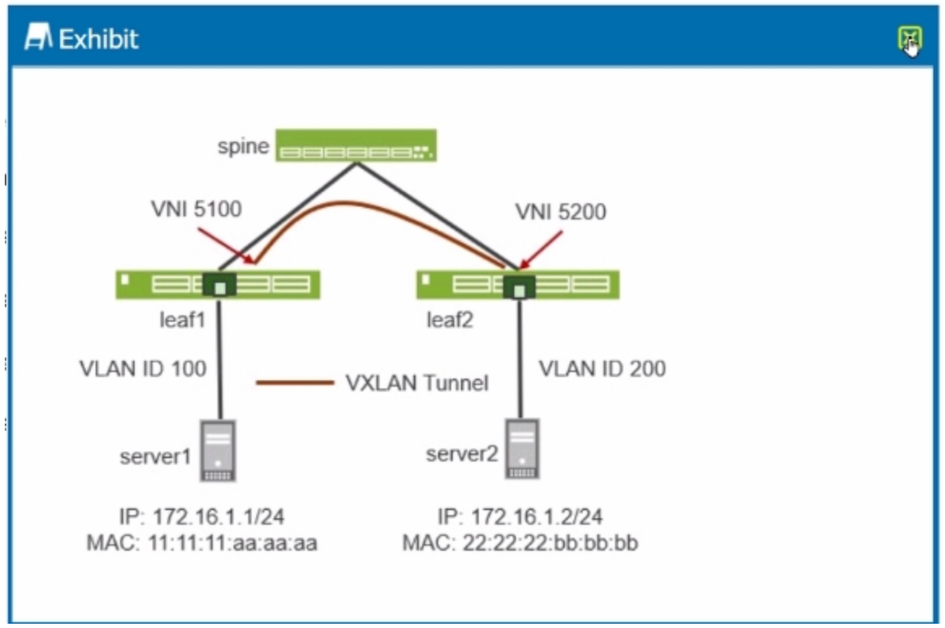

A VXLAN tunnel has been created between leaf1 and leaf2 in your data center. Referring to the exhibit, which statement is correct?

A. Traffic sent from server1 to server2 will be dropped on leaf2.
B. Traffic sent from server1 to server2 will be tagged with VLAN ID 100 on leaf2 and forwarded to server2.
C. Traffic sent from server1 to server2 will be tagged with VLAN ID 200 on leaf2 and forwarded to server2.
D. Traffic sent from server1 to server2 will be dropped on leaf1.

Correct Answer: C
Section:
Explanation:
Understanding VXLAN Tunneling:
VXLAN (Virtual Extensible LAN) is a network virtualization technology that addresses the scalability issues associated with traditional VLANs. VXLAN encapsulates Ethernet frames in UDP, allowing Layer 2 connectivity to extend across Layer 3 networks.
Each VXLAN network is identified by a unique VXLAN Network Identifier (VNI). In this exhibit, we have two VNIs, 5100 and 5200, assigned to the VXLAN tunnels between leaf1 and leaf2.
Network Setup Details:
Leaf1: Connected to Server1 with VLAN ID 100 and associated with VNI 5100.
Leaf2: Connected to Server2 with VLAN ID 200 and associated with VNI 5200.
Spine: Acts as the interconnect between leaf switches.
Traffic Flow Analysis:
When traffic is sent from Server1 to Server2, it is initially tagged with VLAN ID 100 on leaf1.
The traffic is encapsulated into a VXLAN packet with VNI 5100 on leaf1.
The packet is then sent across the network (via the spine) to leaf2.
On leaf2, the VXLAN header is removed, and the original Ethernet frame is decapsulated.
Leaf2 will then associate this traffic with VLAN ID 200 before forwarding it to Server2.
Correct Interpretation of the Exhibit:
The traffic originating from Server1, which is tagged with VLAN ID 100, will be encapsulated into VXLAN and transmitted to leaf2.
Upon arrival at leaf2, it will be decapsulated, and since it is associated with VNI 5200 on leaf2, the traffic will be retagged with VLAN ID 200.
Therefore, the traffic will reach Server2 tagged with VLAN ID 200, which matches the network configuration shown in the exhibit.
Data Center
Reference:

This configuration is typical in data centers using VXLAN for network virtualization. It allows isolated Layer 2 segments (VLANs) to be stretched across Layer 3 boundaries while maintaining distinct VLAN IDs at each site.
This approach is efficient for scaling large data center networks while avoiding VLAN ID exhaustion and enabling easier segmentation.
In summary, the correct behavior, as per the exhibit and the detailed explanation, is that traffic sent from Server1 will be tagged with VLAN ID 200 when it reaches Server2 via leaf2. This ensures proper traffic segmentation and handling across the VXLAN-enabled data center network.

QUESTION 12
Exhibit.

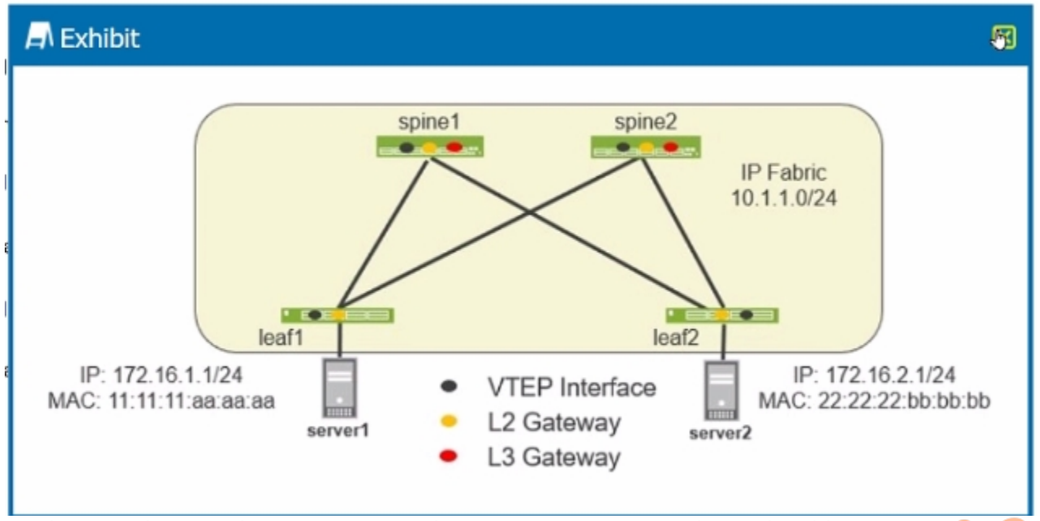

You have implemented an EVPN-VXLAN data center. Device served must be able to communicate with device server2.
Referring to the exhibit, which two statements are correct? (Choose two.)

A. An IRB interface must be configured on spinel and spine2.
B. Traffic from server1 to server2 will transit a VXLAN tunnel to spinel or spine2. then a VXLAN tunnel from spinel or spine2 to leaf2.
C. An IRB Interface must be configured on leaf1 and leaf2.
D. Traffic from server! to server2 will transit the VXLAN tunnel between leaf1 and leaf2.

Correct Answer: C, D
Section:
Explanation:
Understanding the Exhibit Setup:
The network diagram shows an EVPN-VXLAN setup, a common design for modern data centers enabling Layer 2 and Layer 3 services over an IP fabric.
Leaf1 and Leaf2 are the leaf switches connected to Server1 and Server2, respectively, with each server in a different subnet (172.16.1.0/24 and 172.16.2.0/24).
Spine1 and Spine2 are part of the IP fabric, interconnecting the leaf switches.
EVPN-VXLAN Basics:
EVPN (Ethernet VPN) provides Layer 2 and Layer 3 VPN services using MP-BGP.
VXLAN (Virtual Extensible LAN) encapsulates Layer 2 frames into Layer 3 packets for transmission across an IP network.
VTEP (VXLAN Tunnel Endpoint) interfaces on leaf devices handle VXLAN encapsulation and decapsulation.
Integrated Routing and Bridging (IRB):
IRB interfaces are required on leaf1 and leaf2 (where the endpoints are directly connected) to route between different subnets (in this case, between 172.16.1.0/24 and 172.16.2.0/24).
The IRB interfaces provide the necessary L3 gateway functions for inter-subnet communication.
Traffic Flow Analysis:
Traffic from Server1 (172.16.1.1) destined for Server2 (172.16.2.1) must traverse from leaf1 to leaf2.
The traffic will be VXLAN encapsulated on leaf1, sent over the IP fabric, and decapsulated on leaf2.
Since the communication is between different subnets, the IRB interfaces on leaf1 and leaf2 are crucial for routing the traffic correctly.

Correct Statements:
C . An IRB Interface must be configured on leaf1 and leaf2: This is necessary to perform the inter-subnet routing for traffic between Server1 and Server2.
D . Traffic from server1 to server2 will transit the VXLAN tunnel between leaf1 and leaf2: This describes the correct VXLAN operation where the traffic is encapsulated by leaf1 and decapsulated by leaf2.
Data Center
Reference:
In EVPN-VXLAN architectures, the leaf switches often handle both Layer 2 switching and Layer 3 routing via IRB interfaces. This allows for efficient routing within the data center fabric without the need to involve the spine switches for every routing decision.
The described traffic flow aligns with standard EVPN-VXLAN designs, where direct VXLAN tunnels between leaf switches enable seamless and scalable communication across a data center network.

QUESTION 13
Which statement is correct about a collapsed fabric EVPN-VXLAN architecture?

A. Fully meshed back-to-back links are needed between the spine devices.
B. It supports multiple vendors in the fabric as long as all the spine devices are Juniper devices deployed with L2 VTEPs
C. Using Virtual Chassis at the leaf layer increases resiliency.
D. Border gateway functions occur on border leaf devices.

Correct Answer: D
Section:
Explanation:
Collapsed Fabric Architecture:
A collapsed fabric refers to a simplified architecture where the spine and leaf roles are combined, often reducing the number of devices and links required.
In this architecture, the spine typically handles core switching, while leaf switches handle both access and distribution roles.
Understanding Border Gateway Functionality:
Border gateway functions include connecting the data center to external networks or other data centers.
In a collapsed fabric, these functions are usually handled at the leaf level, particularly on border leaf devices that manage the ingress and egress of traffic to and from the data center fabric.
Correct Statement:
D . Border gateway functions occur on border leaf devices: This is accurate in collapsed fabric architectures, where the border leaf devices take on the role of managing external connections and handling routes to other data centers or the internet.
Data Center
Reference:
The collapsed fabric model is advantageous in smaller deployments or scenarios where simplicity and cost-effectiveness are prioritized. It reduces complexity by consolidating functions into fewer devices, and the border leaf handles the critical task of interfacing with external networks.
In conclusion, border gateway functions are effectively managed at the leaf layer in collapsed fabric architectures, ensuring that the data center can communicate with external networks seamlessly.

QUESTION 14
You are deploying an EVPN-VXLAN overlay. You must ensure that Layer 3 routing happens on the spine devices. In this scenario, which deployment architecture should you use?

A. ERB
B. CRB
C. bridged overlay
D. distributed symmetric routing

Correct Answer: B
Section:
Explanation:
Understanding EVPN-VXLAN Architectures:
EVPN-VXLAN overlays allow for scalable Layer 2 and Layer 3 services in modern data centers.
CRB (Centralized Routing and Bridging): In this architecture, the Layer 3 routing is centralized on spine devices, while the leaf devices focus on Layer 2 switching and VXLAN tunneling. This setup is optimal when the goal is to

centralize routing for ease of management and to avoid complex routing at the leaf level.
ERB (Edge Routing and Bridging): This architecture places routing functions on the leaf devices, making it a distributed model where each leaf handles routing for its connected hosts.
Architecture Choice for Spine Routing:
Given the requirement to ensure Layer 3 routing happens on the spine devices, the CRB (Centralized Routing and Bridging) architecture is the correct choice. This configuration offloads routing tasks to the spine, centralizing control and potentially simplifying the overall design.
With CRB, the spine devices perform all routing between VXLAN segments. Leaf switches handle local switching and VXLAN encapsulation, but routing decisions are centralized at the spine level.
This model is particularly advantageous in scenarios where centralized management and routing control are desired, reducing the complexity and configuration burden on the leaf switches.
Data Center
Reference:
The CRB architecture is commonly used in data centers where centralized control and simplified management are key design considerations. It allows the spines to act as the primary routing engines, ensuring that routing is handled in a consistent and scalable manner across the fabric.

QUESTION 15
You want to ensure that VXLAN traffic from the xe-0/0/12 interlace is being encapsulated by logical vlep.32770 and sent to a remote leaf device in this scenario, which command would you use to verify that traffic is flowing?

A. monitor traffic interface xe-0/0/12
B. show interface terse vtep.32770
C. show interfaces terse vtep.32770 statistics
D. show interfaces vtep.32770 detail

Correct Answer: C
Section:
Explanation:
VXLAN Traffic Verification:
To ensure VXLAN traffic from the xe-0/0/12 interface is correctly encapsulated by the logical vtep.32770 and sent to a remote leaf device, it is essential to monitor the relevant interface statistics.
The command show interfaces terse vtep.32770 statistics provides a concise overview of the traffic statistics for the specific VTEP interface, which can help verify whether traffic is being correctly encapsulated and transmitted.
This command is particularly useful for quickly checking the traffic counters and identifying any potential issues with VXLAN encapsulation or transmission.
It allows you to confirm that traffic is flowing as expected, by checking the transmitted and received packet counters.
Data Center
Reference:
Monitoring interface statistics is a crucial step in troubleshooting and validating network traffic, particularly in complex overlay environments like EVPN-VXLAN.

QUESTION 16
Exhibit.

Exhibit

```
user@Leaf-1> show configuration switch-options
service-id 1;
vtep-source-interface lo0.0;
route-distinguisher 192.168.100.51:1;
vrf-target target:65000:1;
user@Leaf-2> show configuration switch-options
vtep-source-interface lo0.0;
route-distinguisher 192.168.100.51:1;
vrf-target target:65000:2;
```

Connections between hosts connected to Leaf-1 and Leaf-2 are not working correctly.

A. Referring to the exhibit, which two configuration changes are required to solve the problem? (Choose two.)
B. Configure the set switch-options vtep-source-interface irb.0 parameter on Leaf-1.
C. Configure the set switch-options vrf-target target:65000:l parameteron Leaf-2.
D. Configure the set switch-options route-distinguisher i92.168.100.50:i parameter on Leaf-1.
E. Configure the set switch-options service-id 1 parameter on Leaf-2.

Correct Answer: C, E
Section:
Explanation:
Issue Analysis:
The problem in the exhibit suggests a mismatch in configuration parameters between Leaf-1 and Leaf-2, leading to communication issues between hosts connected to these leaf devices.
Configuration Mismatches:
Service-ID: Leaf-1 has service-id 1 configured, while Leaf-2 does not have this parameter. For consistency and proper operation, the service-id should be the same across both leaf devices.
VRF Target: Leaf-1 is configured with vrf-target target:65000:1, while Leaf-2 is configured with vrf-target target:65000:2. To allow proper VRF import/export between the two leafs, these should match.
Corrective Actions:
C . Configure the set switch-options vrf-target target:65000:1 parameter on Leaf-2: This aligns the VRF targets between the two leaf devices, ensuring they can correctly import and export routes.
E . Configure the set switch-options service-id 1 parameter on Leaf-2: This ensures that both Leaf-1 and Leaf-2 use the same service ID, which is necessary for consistency in the EVPN-VXLAN setup.
Data Center
Reference:
Correct configuration of VRF targets and service IDs is critical in EVPN-VXLAN setups to ensure that routes and services are correctly shared and recognized between different devices in the network fabric.

QUESTION 17

What are two ways in which an EVPN-signaled VXLAN is different from a multicast-signaled VXLAN? (Choose two.)

A. An EVPN-signaled VXLAN can perform autodiscovery of VTEPs using IS-IS.
B. An EVPN-signaled VXLAN can perform autodiscovery of VTEPs using BGP.
C. An EVPN-signaled VXLAN is less resource intensive.
D. An EVPN-signaled VXLAN features slower and more complete convergence.

Correct Answer: B, C
Section:
Explanation:
Multicast-Signaled VXLAN:
In traditional multicast-signaled VXLAN, VTEPs (VXLAN Tunnel Endpoints) use multicast to flood and learn about remote VTEPs. This method relies on multicast in the underlay network to distribute BUM (Broadcast, Unknown unicast, and Multicast) traffic.
This approach can be resource-intensive due to the need for multicast group management and increased network traffic, especially in large deployments.
EVPN-Signaled VXLAN:
EVPN-signaled VXLAN uses BGP (Border Gateway Protocol) to signal the presence of VTEPs and distribute MAC address information. BGP is used for VTEP autodiscovery and the distribution of endpoint information.
This method is more efficient because it reduces the reliance on multicast, instead using BGP control-plane signaling to handle VTEP discovery and MAC learning, which reduces the overhead on the network and improves scalability.
Correct Statements:
B . An EVPN-signaled VXLAN can perform autodiscovery of VTEPs using BGP: This is correct because EVPN uses BGP for VTEP autodiscovery, making it more efficient and scalable compared to multicast-based methods.
C . An EVPN-signaled VXLAN is less resource-intensive: This is correct because it eliminates the need for multicast flooding in the underlay, instead using BGP for signaling, which is less demanding on network resources.
Incorrect Statements:
A . An EVPN-signaled VXLAN can perform autodiscovery of VTEPs using IS-IS: This is incorrect because EVPN relies on BGP, not IS-IS, for VTEP discovery and signaling.
D . An EVPN-signaled VXLAN features slower and more complete convergence: This is incorrect; EVPN with BGP typically provides faster convergence due to its use of a control plane rather than relying on data plane learning.
Data Center
Reference:
EVPN-VXLAN is widely adopted in modern data center designs due to its scalability, efficiency, and reduced resource consumption compared to multicast-based VXLAN solutions. It leverages the strengths of BGP for control-plane-driven operations, resulting in more efficient and scalable networks.

QUESTION 18
You are implementing VXLAN broadcast domains in your data center environment. Which two statements are correct in this scenario? (Choose two.)

A. A VXLAN packet does not contain a VLAN ID.
B. The VNI must match the VLAN tag to ensure that the remote VTEP can decapsulate VXLAN packets.
C. Layer 2 frames are encapsulated by the source VTEP.
D. The VNI is a 16-bit value and can range from 0 through 16.777.215.

Correct Answer: A, C
Section:
Explanation:
VXLAN Overview:
VXLAN (Virtual Extensible LAN) is a network virtualization technology that encapsulates Layer 2 Ethernet frames into Layer 3 UDP packets for transmission over an IP network. It allows the creation of Layer 2 overlay networks across a Layer 3 infrastructure.
Understanding VXLAN Components:
VTEP (VXLAN Tunnel Endpoint): A VTEP is responsible for encapsulating and decapsulating Ethernet frames into and from VXLAN packets.
VNI (VXLAN Network Identifier): A 24-bit identifier used to distinguish different VXLAN segments, allowing for up to 16 million unique segments.
Correct Statements:
C . Layer 2 frames are encapsulated by the source VTEP: This is correct. In a VXLAN deployment, the source VTEP encapsulates the original Layer 2 Ethernet frame into a VXLAN packet before transmitting it over the IP network to the destination VTEP, which then decapsulates it.

A . A VXLAN packet does not contain a VLAN ID: This is correct. The VXLAN header does not carry the original VLAN ID; instead, it uses the VNI to identify the network segment. The VLAN ID is local to the switch and does not traverse the VXLAN tunnel.
Incorrect Statements:
B . The VNI must match the VLAN tag to ensure that the remote VTEP can decapsulate VXLAN packets: This is incorrect. The VNI is independent of the VLAN tag, and the VLAN ID does not need to match the VNI. The VNI is what the remote VTEP uses to identify the correct VXLAN segment.
D . The VNI is a 16-bit value and can range from 0 through 16,777,215: This is incorrect because the VNI is a 24-bit value, allowing for a range of 0 to 16,777,215.
Data Center
Reference:
VXLAN technology is critical for modern data centers as it enables scalability and efficient segmentation without the constraints of traditional VLAN limits.

QUESTION 19
You are deploying an IP fabric using EBGP and notice that your leaf devices are advertising and receiving all the routes. However, the routes are not installed in the routing table and are marked as hidden.
Which two statements describe how to solve the issue? (Choose two.)

A. You need to configure as-override.
B. You need to configure a next-hop self policy.
C. You need to configure loops 2.
D. You need to configure multipath multiple-as.

Correct Answer: B, D
Section:
Explanation:
Issue Overview:
The leaf devices in an IP fabric using eBGP are advertising and receiving all routes, but the routes are not being installed in the routing table and are marked as hidden. This typically indicates an issue with the BGP configuration, particularly with next-hop handling or AS path concerns.
Corrective Actions:
B . You need to configure a next-hop self policy: This action ensures that the leaf devices modify the next-hop attribute to their own IP address before advertising routes to their peers. This is particularly important in eBGP setups where the next-hop may not be directly reachable by other peers.
D . You need to configure multipath multiple-as: This setting allows the router to accept multiple paths from different autonomous systems (ASes) and use them for load balancing. Without this, the BGP process might consider only one path and mark others as hidden.
Incorrect Statements:
A . You need to configure as-override: AS-override is used to replace the AS number in the AS-path attribute to prevent loop detection issues in MPLS VPNs, not in a typical eBGP IP fabric setup.
C . You need to configure loops 2: There is no specific BGP command loops 2 relevant to resolving hidden routes in this context. It might be confused with allowas-in, which is used to allow AS path loops under certain conditions.
Data Center
Reference:
Proper BGP configuration is crucial in IP fabrics to ensure route propagation and to prevent routes from being marked as hidden. Configuration parameters like next-hop self and multipath multiple-as are common solutions to ensure optimal route installation and load balancing in a multi-vendor environment.

QUESTION 20
In your EVPN-VXAN environment, you want to prevent a multihomed server from receiving multiple copies of BUM traffic in active/active scenarios. Which EVPN route type would satisfy this requirement?

A. Type 8
B. Type 7
C. Type 4
D. Type 5

Correct Answer: C
Section:

Explanation:
Understanding the Scenario:
In an EVPN-VXLAN environment, when using multi-homing in active/active scenarios, there's a risk that a multihomed server might receive duplicate copies of Broadcast, Unknown unicast, and Multicast (BUM) traffic. This is because multiple VTEPs might forward the same BUM traffic to the server.
EVPN Route Types:
Type 4 Route (Ethernet Segment Route): This route type is used to advertise the Ethernet Segment (ES) to which the device is connected. It is specifically used in multi-homing scenarios to signal the ES and its associated Ethernet Tag to all the remote VTEPs. The Type 4 route includes information that helps prevent BUM traffic duplication in active/active multi-homing by using a split-horizon mechanism, which ensures that traffic sent to a multihomed device does not get looped back.
The Type 4 route is crucial for ensuring that in a multi-homed setup, particularly in an active/active configuration, BUM traffic does not result in duplication at the server. The route helps coordinate which VTEP is responsible for forwarding the BUM traffic to the server, thereby preventing duplicate traffic.
Data Center
Reference:
Type 4 routes are essential for managing multi-homing in EVPN to avoid the issues of BUM traffic duplication, which could otherwise lead to inefficiencies and potential network issues.

QUESTION 21
You want to convert an MX Series router from a VXLAN Layer 2 gateway to a VXLAN Layer 3 gateway for VNI 100. You have already configured an IRB interface. In this scenario, which command would you use to accomplish this task?

A. set protocols isis interface irb.100 passive
B. set vlans VLAN-100 13-interface irb.100
C. set bridge-domains VLAN-100 routing-interface irb.100
D. set protocols ospf area 0.0.0.0 interface irb.100 passive

Correct Answer: C
Section:
Explanation:
Scenario Overview:
Converting an MX Series router from a VXLAN Layer 2 gateway to a VXLAN Layer 3 gateway involves transitioning the router's functionality from simply bridging traffic within a VXLAN segment to routing traffic between different segments.
Key Configuration Requirement:
IRB (Integrated Routing and Bridging) Interface: An IRB interface allows for both Layer 2 switching and Layer 3 routing. To enable routing for a specific VNI (VXLAN Network Identifier), the IRB interface must be associated with the routing function in the corresponding bridge domain.
Correct Command:
C . set bridge-domains VLAN-100 routing-interface irb.100: This command correctly binds the IRB interface to the bridge domain, enabling Layer 3 routing functionality within the VXLAN for VNI 100. This effectively transitions the device from operating solely as a Layer 2 gateway to a Layer 3 gateway.
Data Center
Reference:
This configuration step is essential when converting a Layer 2 VXLAN gateway to a Layer 3 gateway, enabling the MX Series router to route between VXLAN segments.

QUESTION 22
You manage an IP fabric with an EVPN-VXLAN overlay. You have multiple tenants separated using multiple unique VRF instances. You want to determine the routing information that belongs in each routing instance's routing table.
In this scenario, which property is used for this purpose?

A. the VRF target community
B. the routing instance type
C. the VRF table label
D. the route distinguisher value

Correct Answer: D
Section:
Explanation:
Understanding VRF and Routing Instances:
In an EVPN-VXLAN overlay network, multiple tenants are separated using unique VRF (Virtual Routing and Forwarding) instances. Each VRF instance maintains its own routing table, allowing for isolated routing domains within the same network infrastructure.
Role of Route Distinguisher:
Route Distinguisher (RD): The RD is a unique identifier used in MPLS and EVPN environments to distinguish routes belonging to different VRFs. The RD is prepended to the IP address in the route advertisement, ensuring that routes from different tenants remain unique even if they use the same IP address range.
Correct Property:
the route distinguisher value: This is the correct answer because the RD is crucial in determining which routing information belongs to which VRF instance. It ensures that each VRF's routing table only contains relevant routes, maintaining isolation between tenants. Data Center
Reference: The RD is a key element in MPLS and EVPN-based multi-tenant environments, ensuring proper routing segregation and isolation for different VRFs within the data center fabric.

QUESTION 23
Exhibit.

Exhibit

```
user@Border-Leaf-1> show configuration protocols bgp
group UNDERLAY {
    type external;
    export LOOPBACKS;
    local-as 65205;
    multipath {
        multiple-as;
    }
    neighbor 172.16.1.5 {
        peer-as 65102;
    }
}
group OVERLAY {
    type external;
    local-address 192.168.100.4;
    family evpn {
        signaling;
    }
    local-as 65101;
    neighbor 192.168.100.1 {
        peer-as 65102;
    }
    neighbor 192.168.100.22 {
        description Border-Leaf-2;
        peer-as 65222;
    }
    accept-remote-nexthop;
}
group PROVIDER {
    type external;
    peer-as 65001;
    local-as 65002;
    neighbor 172.16.1.224;
}
```


You are troubleshooting a DCI connection to another data center The BGP session to the provider is established, but the session to Border-Leaf-2 is not established. Referring to the exhibit, which configuration change should be made to solve the problem?

A. set protocols bgp group overlay export loopbacks
B. delete protocols bgp group UNDERLAY advertise-external
C. set protocols bgp group PROVIDER export LOOPBACKS
D. delete protocols bgp group OVERLAY accept-remote-nexthop

Correct Answer: D
Section:
Explanation:
Understanding the Configuration:
The exhibit shows a BGP configuration on a Border-Leaf device. The BGP group UNDERLAY is used for the underlay network, OVERLAY for EVPN signaling, and PROVIDER for connecting to the provider network.
The OVERLAY group has the accept-remote-nexthop statement, which is designed to accept the next-hop address learned from the remote peer as is, without modifying it.

Problem Identification:
The BGP session to Border-Leaf-2 is not established. A common issue in EVPN-VXLAN environments is related to next-hop reachability, especially when accept-remote-nexthop is configured.
In typical EVPN-VXLAN setups, the next-hop address should be reachable within the overlay network. However, the accept-remote-nexthop can cause issues if the next-hop IP address is not directly reachable or conflicts with the expected behavior in the overlay.
Corrective Action:
delete protocols bgp group OVERLAY accept-remote-nexthop: Removing this command will ensure that the device uses its own IP address as the next-hop in BGP advertisements, which is standard practice in many EVPN-VXLAN setups. This change should help establish the BGP session with Border-Leaf-2. Data Center
Reference: Proper handling of BGP next-hop attributes is critical in establishing and maintaining stable BGP sessions, especially in complex multi-fabric environments like EVPN-VXLAN. Removing accept-remote-nexthop aligns with best practices in many scenarios.

QUESTION 24
Exhibit.

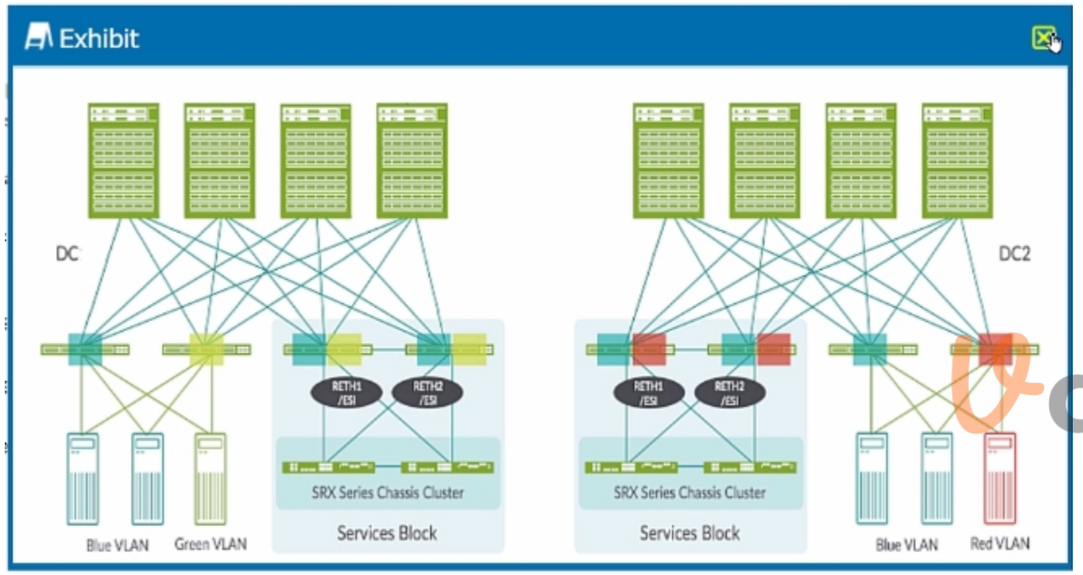

Both DC and DC2 ate using EVPN-VXLAN technology deployed using an ERB architecture. A server on the Red VLAN must communicate with a server on the Green VLAN. The Blue VLAN in DC and DC2 needs to be the same VLAN.
Which statement is correct in this scenario?

A. The eight spine devices must be configured as border spine devices; a full mush interconnect must exist between all eight spine devices and the Blue VLAN must be stitched together
B. An interconnect is required between the four SRX Series devices; the Blue VLAN must be stretched and a transit VNI must be added for the Red and Green VLANs.
C. An interconnect is required between four leaf devices in the services blocks; the Red VLAN and the Green VLAN must be stitched and the Blue VLAN must be stretched.
D. A lean super spine device must be added to DC and DC2; all VLANs must be stretched to the lean super spine device and the lean super spine devices must stitch all the VLANs together.

Correct Answer: B
Section:
Explanation:
ERB Architecture in EVPN-VXLAN:
ERB (Edge Routed Bridging) architecture is commonly used in data center networks where routing decisions are made at the network edge (leaf or border devices), while bridging (Layer 2 forwarding) is extended across the fabric. This architecture allows for efficient L3 routing while still enabling L2 services like VLANs to span across multiple locations.
VLAN and VNI Configuration:
The scenario specifies that a server on the Red VLAN needs to communicate with a server on the Green VLAN. Since these VLANs are in different data centers (DC and DC2), and given the use of EVPN-VXLAN, the communication between these VLANs will require a transit VNI (Virtual Network Identifier). This transit VNI will allow traffic to traverse the VXLAN tunnel across the DCI (Data Center Interconnect).

Interconnect between SRX Series Devices:
The exhibit shows SRX Series Chassis Clusters used as service devices (likely for firewalling or other security services). These devices need to be interconnected between the two data centers to ensure that VLANs can communicate effectively. The Blue VLAN needs to be stretched between DC and DC2 to maintain the same Layer 2 domain across both data centers.
Conclusion:
Option B: Correct---Interconnecting the SRX Series devices will ensure the necessary service chaining, while stretching the Blue VLAN and adding a transit VNI for the Red and Green VLANs will enable the required communication across the data centers.

QUESTION 25
You are adding a server lo a tenant's network within your data center and must limit access to a specific traffic type within the tenant network without pushing all tenant traffic through a firewall.
What will satisfy this requirement?

A. Use route leaking with EVPN and a routing policy.
B. Use filter-based forwarding.
C. Put the new server on a unique subnet within the tenant's network.
D. Use a static route in the tenant VRF with a firewall as the next hop for traffic to the new server.

Correct Answer: B
Section:
Explanation:
Controlling Traffic Within a Tenant's Network:
The requirement is to limit access to specific traffic types within a tenant's network without routing all tenant traffic through a firewall. This requires a selective method that can direct specific types of traffic to different paths based on the nature of the traffic.
Filter-Based Forwarding (FBF):
FBF is a technique that allows for routing decisions based on filters applied to the traffic, such as matching on source IP addresses, destination IP addresses, or even specific application types (like HTTP or FTP). This allows specific types of traffic to be forwarded to a specific next hop (e.g., a firewall) without affecting the entire traffic flow within the tenant's network.
Conclusion:
Option B: Correct---Filter-based forwarding allows for granular control of traffic, ensuring that only specific types of traffic within the tenant's network are redirected through a firewall, satisfying the requirement.

QUESTION 26
Why is a designated forwarder required in a multihomed CE-to-PE VXLAN environment using EVPN signalling?

A. The designated forwarder is required to prevent packets from looping between the PEs.
B. The designated forwarder is required to prevent flooding of MAC addresses to multihomed hosts.
C. The designated forwarder is required to prevent a traffic storm from being received on multihomed hosts.
D. The designated forwarder is required to prevent duplicate packets from being received on multihomed hosts.

Correct Answer: D
Section:
Explanation:
Understanding Multihomed CE-to-PE VXLAN Environment:
In a VXLAN environment using EVPN signaling, multiple PEs (Provider Edge devices) can be connected to the same CE (Customer Edge device). This setup is referred to as multihoming, where a CE device has multiple connections to the network to ensure redundancy and load balancing.
Role of the Designated Forwarder:
The designated forwarder (DF) is a mechanism used in EVPN to manage the forwarding of broadcast, unknown unicast, and multicast (BUM) traffic in a multihomed environment. The DF is selected to ensure that only one of the PEs forwards this type of traffic to the CE, preventing loops and unnecessary duplicate packets.
Avoiding Duplicate Packets:
Without a designated forwarder, all PEs connected to a multihomed CE could potentially forward the same packet to the CE, resulting in duplicate packets. This duplication can cause issues with packet processing on the CE, leading to inefficiencies and potential network problems.
Conclusion:

Option D: Correct---The designated forwarder is essential to prevent duplicate packets from being received on multihomed hosts, ensuring that only one PE forwards BUM traffic to the CE.

QUESTION 27
You are asked to deploy 100 QFX Series devices using ZTP Each OFX5120 requires a different configuration. In this scenario, what are two components that you would configure on the DHCP server? (Choose two.)

A. the IP address of the FTP server
B. the MAC address for each OFX5120
C. the MAC address of the FTP server
D. the management IP address for each OFX5120

Correct Answer: B, D
Section:
Explanation:
Zero Touch Provisioning (ZTP):
ZTP allows for the automated configuration of network devices, like QFX Series switches, without manual intervention. During ZTP, a switch will obtain its configuration from a DHCP server and then download the required software and configuration files from a specified server (e.g., FTP, HTTP).
DHCP Server Configuration:
Option B: The DHCP server needs to know the MAC address for each QFX5120 to provide a specific configuration based on the device identity. By mapping the MAC address to a particular configuration, the DHCP server can ensure that each switch gets the correct configuration.
Option D: The management IP address for each QFX5120 must also be assigned by the DHCP server. This IP address allows the device to communicate on the network and access the configuration files and other required resources during the ZTP process.
Conclusion:
Option B: Correct---MAC addresses allow the DHCP server to identify each QFX5120 and assign the appropriate configuration.
Option D: Correct---Management IP addresses are essential for network communication during ZTP.

QUESTION 28
Which two statements are correct about an IP fabric? (Choose two.)

A. All leaf devices can use the same AS number in an IP fabric without making any adjustments to the EBGP configuration
B. The multipath multiple-as statement is required to enable ECMP if every device has a different AS number.
C. Only a single point to point EBGP session is required between peers in an IP fabric.
D. FBGP is only required to route most routing information to external devices outside the fabric.

Correct Answer: A, B
Section:
Explanation:
BGP in IP Fabric:
In an IP fabric, Border Gateway Protocol (BGP) is used to manage the routing between leaf and spine devices. Each device can have the same or different Autonomous System (AS) numbers depending on the network design.
Multipath Multiple-AS:
Option B: If every device in the fabric has a different AS number, then enabling Equal-Cost Multi-Path (ECMP) routing requires the multipath multiple-as statement. This configuration allows BGP to consider multiple paths across different AS numbers as equal cost, enabling efficient load balancing across the network.
Same AS Number Configuration:
Option A: It's possible for all leaf devices to use the same AS number in an IP fabric, which simplifies the configuration. EBGP (External BGP) will still function correctly in this setup because BGP considers the peering relationship rather than strictly enforcing different AS numbers in this specific use case.
Conclusion:
Option B: Correct---This statement is essential for enabling ECMP in a multi-AS environment.
Option A: Correct---Leaf devices can share the same AS number without needing special EBGP configuration.

QUESTION 29

As part of the onboarding process for new switches being added to your data centers, your company uses Juniper Networks' ZTP process. As part of the ZTP process, a script is executed by the devices being onboarded. Which statement is correct in this scenario?

A. The Junos ZTP process supports Shell. JScript. and Ansible.
B. The Junos ZTP process supports Python, SLAX, and Perl.
C. The Junos ZTP process supports JScript. Ansible. and Perl.
D. The Junos ZTP process supports Shell. Python, and SLAX.

Correct Answer: D
Section:
Explanation:
Zero Touch Provisioning (ZTP):
Juniper Networks' ZTP (Zero Touch Provisioning) process automates the deployment of new devices by allowing them to fetch and execute scripts for configuration and setup as they are powered on and connected to the network.
Supported Scripting Languages:
The Junos OS supports several scripting languages that can be used during the ZTP process:
Shell scripts are often used for general automation tasks.
Python is a widely supported language in Junos, offering powerful scripting capabilities for automating network tasks.
SLAX (Service Logic Execution Environment) is a scripting language specific to Junos, designed to automate configuration tasks and simplify network operations.
Conclusion:
Option D: Correct---Junos ZTP supports Shell, Python, and SLAX, making it the correct choice among the provided options.

QUESTION 30
Exhibit.

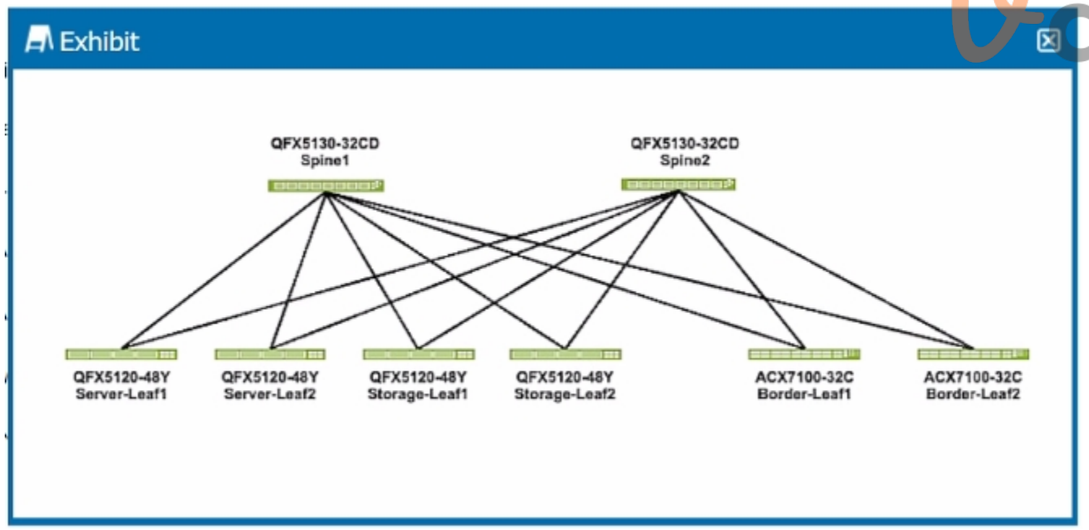

You are deploying a VXLAN overlay with EVPN as the control plane in an ERB architecture.
Referring to the exhibit, which three statements are correct about where the VXLAN gateways will be placed? (Choose three.)

A. Only the spine devices will have L2 VXLAN gateways.
B. All leaf devices will have L2 VXLAN gateways.
C. All leaf devices will have L3 VXLAN gateways.
D. Only the border and leaf devices will have L3 VXLAN gateways.
E. Spine devices will have no VXLAN gateways.

Correct Answer: B, C, E
Section:
Explanation:
Understanding ERB Architecture:
ERB (Edge Routed Bridging) architecture is a network design where the routing occurs at the edge (leaf devices) rather than in the spine devices. In a VXLAN overlay network with EVPN as the control plane, leaf devices typically act as both Layer 2 (L2) and Layer 3 (L3) VXLAN gateways.
Placement of VXLAN Gateways:
Option B: All leaf devices will have L2 VXLAN gateways to handle the bridging of VLAN traffic into VXLAN tunnels.
Option C: All leaf devices will also have L3 VXLAN gateways to route traffic between different VXLAN segments (VNIs) and external networks.
Option E: Spine devices in an ERB architecture generally do not function as VXLAN gateways. They primarily focus on forwarding traffic between leaf nodes and do not handle VXLAN encapsulation/decapsulation.
Conclusion:
Option B: Correct---All leaf devices will have L2 VXLAN gateways.
Option C: Correct---All leaf devices will have L3 VXLAN gateways.
Option E: Correct---Spine devices will not act as VXLAN gateways

QUESTION 31
Which parameter is used to associate a received route with a local VPN route table?

A. route-target community
B. VLAN ID
C. VNI
D. route-distinguisher

Correct Answer: A
Section:
Explanation:
Understanding VPN Route Table Association:
In MPLS/VPN and EVPN networks, the route-target community is a BGP extended community attribute used to control the import and export of VPN routes. It associates received routes with the appropriate VPN route tables on the PE (Provider Edge) routers.
Function of Route-Target Community:
The route-target community tag ensures that routes are imported into the correct VRF (Virtual Routing and Forwarding) instance, allowing them to be correctly routed within the VPN.
Conclusion:
Option A: Correct---The route-target community is used to associate received routes with a local VPN route table.

QUESTION 32
You are designing an IP fabric tor a large data center, and you are concerned about growth and scalability. Which two actions would you take to address these concerns? (Choose two.)

A. Design a five-stage Clos IP fabric.
B. Design a three-stage Clos IP fabric.
C. Use EX4300 Series devices as the spine devices.
D. Use OFX5700 Series devices as the super spines.

Correct Answer: B, D
Section:
Explanation:
Clos IP Fabric Design:
A Clos fabric is a network topology designed for scalable, high-performance data centers. It is typically arranged in multiple stages, providing redundancy, high bandwidth, and low latency.
Three-Stage Clos Fabric:
Option B: A three-stage Clos fabric, consisting of leaf, spine, and super spine layers, is widely used in data centers. This design scales well and allows for easy expansion by adding more leaf and spine devices as needed.

Super Spines for Scalability:
Option D: Using high-capacity devices like the QFX5700 Series as super spines can handle the increased traffic demands in large data centers and support future growth. These devices provide the necessary bandwidth and scalability for large-scale deployments.
Conclusion:
Option B: Correct---A three-stage Clos fabric is a proven design that addresses growth and scalability concerns in large data centers.
Option D: Correct---QFX5700 Series devices are suitable for use as super spines in large-scale environments due to their high performance.

QUESTION 33
You are asked to configure telemetry on the OFX Series devices in your data center fabric. You want to use sensors that have a vendor-neutral data model Which type of sensor should you use in this scenario?

A. JTI OpenConfig sensors
B. JTI native sensors
C. Python sensors
D. analog sensors

Correct Answer: A
Section:
Explanation:
Telemetry in Data Centers:
Telemetry allows for real-time monitoring of network devices by collecting and exporting data such as interface statistics, routing table updates, and other key metrics.
Option A: JTI (Junos Telemetry Interface) OpenConfig sensors use a vendor-neutral data model, which is important for ensuring compatibility across different network devices and systems. OpenConfig is an industry-standard model, which facilitates integration with various telemetry collection systems.
Conclusion:
Option A: Correct---OpenConfig sensors provide a vendor-neutral solution for telemetry, ensuring broad compatibility and flexibility in data center environments.

QUESTION 34
What are two supported methods (or exporting data when using the Junos telemetry interface? (Choose two.)

A. using REST
B. using UDP
C. using SNMP
D. using gRPC

Correct Answer: B, D
Section:
Explanation:
Junos Telemetry Interface (JTI):
The Junos Telemetry Interface is a framework that allows network operators to collect real-time telemetry data from Juniper devices. This data can be used for monitoring, analytics, and network automation.
Data Export Methods:
Option B: UDP (User Datagram Protocol) is a lightweight, connectionless protocol used for exporting telemetry data quickly with minimal overhead. While it doesn't guarantee delivery, it is suitable for high-speed data transfer where occasional packet loss is acceptable.
Option D: gRPC (gRPC Remote Procedure Call) is a modern, high-performance method for data export that supports streaming and remote procedure calls, making it ideal for more complex telemetry data use cases.
Conclusion:
Option B: Correct---UDP is supported for exporting telemetry data.
Option D: Correct---gRPC is also supported, offering advanced streaming capabilities

QUESTION 35
You are deploying a new network lo support your Al workloads on devices that support at least 400 Gbps Ethernet. There is no requirement for any Layer 2 VLANs in this network. Which network architecture would satisfy this requirement?

A. an IP fabric using PIM-SM to signal VXLAN overlay
B. an IP fabric using the EVPN-MPLS architecture
C. an IP fabric with an EVPN-VXLAN architecture
D. an IP fabric using EBGP

Correct Answer: D
Section:
Explanation:
Requirements for AI Workloads:
The scenario requires a network that supports at least 400 Gbps Ethernet and does not require Layer 2 VLANs. This setup is well-suited for a pure Layer 3 network, which can efficiently route traffic between devices without the overhead or complexity of maintaining Layer 2 domains.
Choosing the Right Network Architecture:
Option D: An IP fabric using EBGP (External BGP) is ideal for this scenario. In a typical IP fabric, EBGP is used to handle routing between spine and leaf switches, creating a scalable and efficient network. Since there is no need for Layer 2 VLANs, the pure IP fabric design with EBGP provides a straightforward and effective solution.
Options A, B, and C involve more complex architectures (like VXLAN or EVPN), which are unnecessary when there's no requirement for Layer 2 overlays or VLANs.
Conclusion:
Option D: Correct---An IP fabric with EBGP is the most suitable and straightforward architecture for a network that needs to support high-speed AI workloads without Layer 2 VLANs.

QUESTION 36
You are preparing an sFlow monitoring system configuration.
In this scenario, what Information will be included in the datagram sent to the sFlow collector? (Choose two.)

A. the interlace through which the packets entered the agent
B. the sending device's serial number
C. the CRC from the sampled packet
D. the source and destination VLAN for sampled packets

Correct Answer: A, D
Section:
Explanation:
Understanding sFlow Monitoring:
sFlow is a packet sampling technology used to monitor traffic in a network. It sends sampled packet data and interface counters to an sFlow collector, which analyzes the traffic patterns.
Information Included in sFlow Datagram:
Option A: The datagram sent to the sFlow collector includes information about the interface through which the packets entered the agent (the switch or router). This is crucial for understanding where in the network the traffic was captured.
Option D: sFlow datagrams also include the source and destination VLAN for the sampled packets. This allows for detailed analysis of the traffic flow within different VLANs.
Conclusion:
Option A: Correct---The ingress interface is included in the sFlow datagram.
Option D: Correct---The source and destination VLANs are also included, providing context for the sampled traffic.

QUESTION 37
You are deploying multiple Juniper switches al the same location. Your switches are currently using the factory-default configuration.
In this scenario, which two statements are correct? (Choose two.)

A. The DHCP server configuration cannot provide Junos version requirements to DHCP clients.
B. The switch will try to request an IP address from a DHCP server using all interfaces that are connected and are operational.
C. The switch will try to request an IP address from a DHCP server using only the management interface.

D. The DHCP server configuration can provide Junos version requirements to DHCP clients.

Correct Answer: B, D
Section:
Explanation:
DHCP Behavior in Factory-Default Configuration:
Option B: In the factory-default configuration, Juniper switches are designed to send DHCP requests on all operational interfaces. This behavior ensures that the switch can obtain an IP address for management and further configuration from any available DHCP server.
Option D: The DHCP server can provide additional configuration parameters, including the required Junos version. This allows for automated provisioning and ensures that the switch is running the correct software version.
Conclusion:
Option B: Correct---The switch will use any operational interface to request an IP address via DHCP.
Option D: Correct---The DHCP server can specify Junos version requirements, enabling automated software management.