ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

Exhibit.

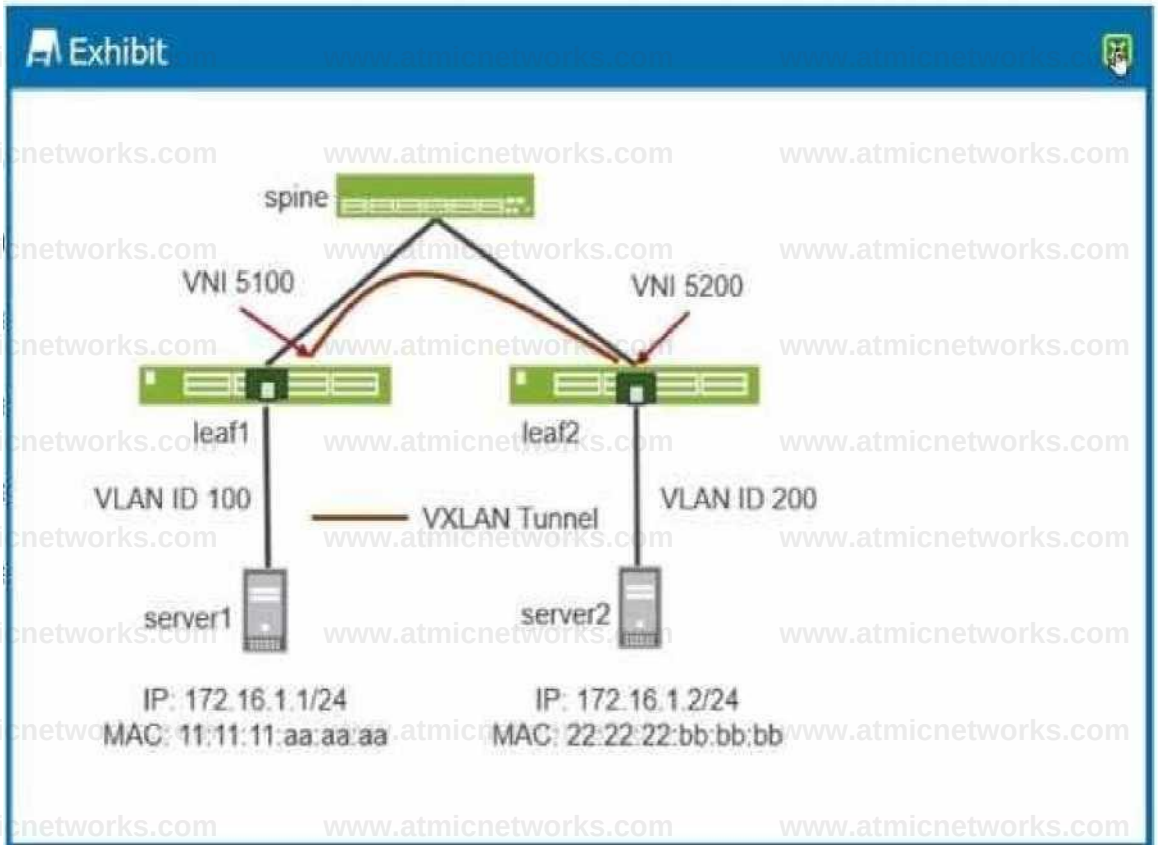

A VXLAN tunnel has been created between leaf1 and leaf2 in your data center. Referring to the exhibit, which statement is correct?

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.

Answer: C

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: 2

Exhibit.

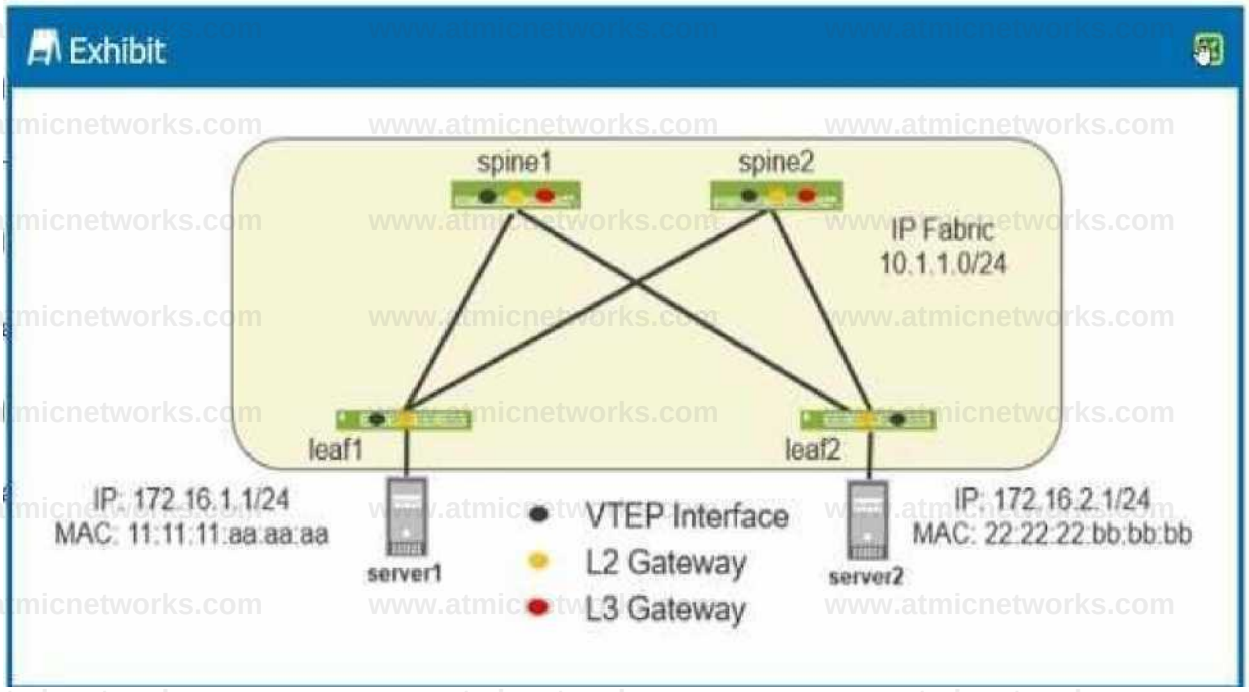

You have implemented an EVPN-VXLAN data center. Device served must be able to communicate with device server2.
Referring to the exhibit, which two statements are correct? (Choose two.)

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.

Answer: CD

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 intersubnet 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: 3

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.

Answer: D

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: 4

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

Answer: B

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.

Explanation:

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: 5

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

Answer: C

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.

Explanation:

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: 6

Exhibit.

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:I 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.

Answer: CE

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: 7

What are three actions available tor MAC move limiting? (Choose three.)

A. drop
B. filter
C. enable
D. log
E. shutdown

Answer: ADE

Explanation:

MAC Move Limiting:
MAC move limiting is a security feature used in network switches to detect and mitigate rapid changes in MAC address locations, which could indicate a network issue or an attack such as MAC flapping or spoofing. When a MAC address is learned on a different interface than it was previously learned, the switch can take various actions to prevent potential issues.
Available Actions:
A . drop: This action drops packets from the MAC address if it violates the move limit, effectively blocking communication from the offending MAC address.
D . log: This action logs the MAC move event without disrupting traffic, allowing network administrators to monitor and investigate the event.
E . shutdown: This action shuts down the interface on which the MAC address violation occurred, effectively stopping all traffic on that interface to prevent further issues.
Other Actions (Not Correct):
B . filter: Filtering is not typically associated with MAC move limiting; it generally refers to applying ACLs or other mechanisms to filter traffic.
C . enable: This is not an action related to MAC move limiting, as it does not represent a specific reaction to a MAC move event.
Data Center Reference:
MAC move limiting is crucial for maintaining network stability and security, particularly in environments with dynamic or large-scale Layer 2 networks where MAC addresses might frequently change locations.

Question: 8

Exhibit.

Referring to the exhibit, the spine1 device has an underlay BGP group that is configured to peer with its neighbors' directly connected interfaces. Which two statements are true in this scenario? (Choose two.)

A. The multihop statement is not required to establish the underlay BGP sessions.
B. Load balancing for the underlay is not configured correctly.
C. The multihop statement is required to establish the underlay BGP sessions.
D. Load balancing for the underlay is configured correctly.

Answer: AD

Explanation:

Understanding BGP Configuration in the Exhibit:
The exhibit shows a BGP configuration on spine1 with a group named underlay, configured to peer with directly connected interfaces of other devices in the network.
Multipath multiple-as: This statement allows the router to install multiple paths in the routing table for routes learned from different ASes, facilitating load balancing.
Key Statements:
A . The multihop statement is not required to establish the underlay BGP sessions: In this case, the BGP peers are directly connected (as indicated by their neighbor IP addresses), so the multihop statement is unnecessary. Multihop is typically used when BGP peers are not directly connected and packets need to traverse multiple hops.
D . Load balancing for the underlay is configured correctly: The multipath { multiple-as; } statement in the configuration enables load balancing across multiple paths from different autonomous systems, which is appropriate for underlay networks in data center fabrics.
Incorrect Statements:
C . The multihop statement is required to establish the underlay BGP sessions: This is incorrect because the peers are directly connected, making the multihop statement unnecessary.

B . Load balancing for the underlay is not configured correctly: This is incorrect because the configuration includes the necessary multipath settings for load balancing.
Data Center Reference:
BGP configurations in EVPN-VXLAN underlay networks are crucial for ensuring redundancy, load balancing, and efficient route propagation across the data center fabric.

Question: 9

You want to provide a OCI that keeps each data center routing domain isolated, while also supporting translation of VNIs. Which DCI scheme allows these features?

A. MPLS DCI label exchange
B. over the top (OTT) with VNI translation enabled
C. VXLAN stitching
D. over the top (OTT) with proxy gateways

Answer: C

Explanation:

Understanding DCI (Data Center Interconnect) Schemes:
DCI schemes are used to connect multiple data centers, enabling seamless communication and resource sharing between them. The choice of DCI depends on the specific requirements, such as isolation, VNI translation, or routing domain separation.
VXLAN Stitching:
VXLAN stitching involves connecting multiple VXLAN segments, allowing VNIs (VXLAN Network Identifiers) from different segments to communicate with each other while maintaining separate routing domains.
This approach is particularly effective for keeping routing domains isolated while supporting VNI translation, making it ideal for scenarios where you need to connect different data centers or networks without merging their control planes.
Other Options:
A . MPLS DCI label exchange: This option typically focuses on MPLS-based interconnections and does not inherently support VNI translation or isolation in the context of VXLAN.
B . Over the top (OTT) with VNI translation enabled: This could support VNI translation but does not inherently ensure routing domain isolation.
D . Over the top (OTT) with proxy gateways: This typically involves using external gateways for traffic routing and may not directly support VNI translation or isolation in the same way as VXLAN stitching. Data Center Reference:
VXLAN stitching is a powerful method in multi-data center environments, allowing for flexibility in connecting various VXLAN segments while preserving network isolation and supporting complex interconnect requirements.

Question: 10

Exhibit.

Given the configuration shown in the exhibit, why has the next hop remained the same for the EVPN routes advertised to the peer 203.0.113.2?

A. EVPN routes cannot have the next hop changed.
B. The export policy is incorrectly configured.
C. The vrf-export parameter must be applied.
D. The vpn-apply-export parameter must be applied to this peer.

Answer: D

Explanation:

Understanding the Configuration:
The configuration shown in the exhibit involves an EVPN (Ethernet VPN) setup using BGP as the routing protocol. The export policy named CHANGE_NH is applied to the BGP group evpn-peer, which includes a rule to change the next hop for routes that match the policy.
Issue with Next Hop Not Changing:
The policy CHANGE_NH is correctly configured to change the next hop to 203.0.113.10 for the matching routes. However, the next hop remains unchanged when advertising EVPN routes to the peer 203.0.113.2.
Reason for the Issue:
In Junos OS, when exporting routes for VPNs (including EVPN), the next-hop change defined in a policy will not take effect unless the vpn-apply-export parameter is used in the BGP configuration. This parameter ensures that the export policy is applied specifically to VPN routes.
The vpn-apply-export parameter must be included to apply the next-hop change to EVPN routes. Correct Answer Explanation:
D . The vpn-apply-export parameter must be applied to this peer: This is the correct solution because the next hop in EVPN routes won't be altered without this parameter in the BGP configuration. It instructs the BGP process to apply the export policy to the EVPN routes.
Data Center Reference:

This behavior is standard in EVPN deployments with Juniper Networks devices, where the export policies applied to VPN routes require explicit invocation using vpn-apply-export to take effect.

Question: 11

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.

Answer: BC

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 controlplane 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 multicastbased 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: 12

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.

Answer: AC

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: 13

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.

Answer: BD

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: 14

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

Answer: C

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. Explanation:
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: 15

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

Answer: C

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: 16

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

Answer: D

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:
E. 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: 17

Exhibit.

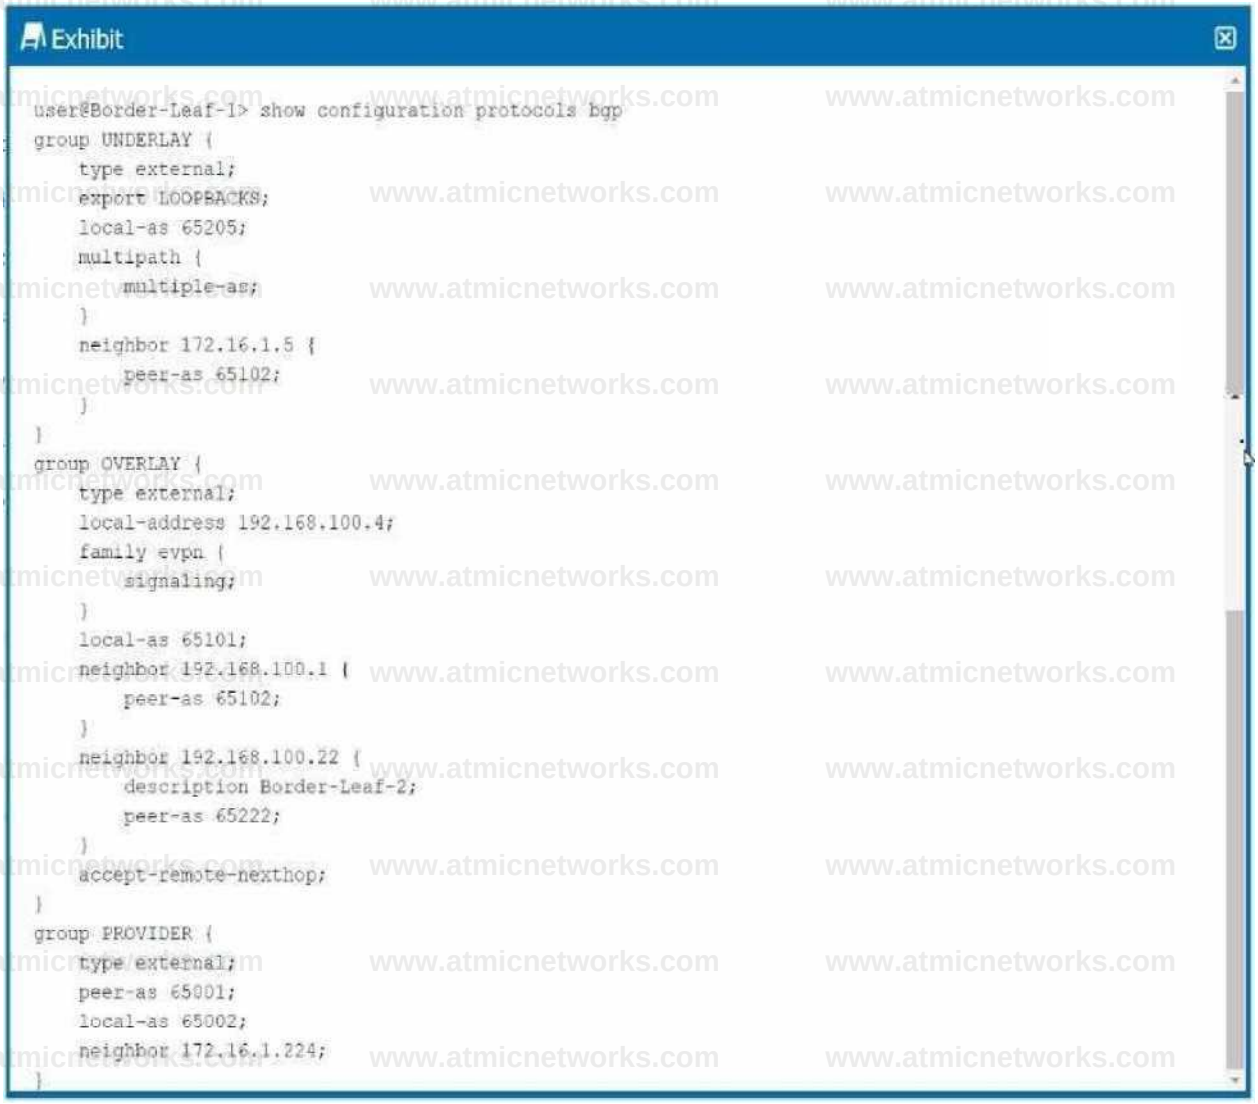

You are troubleshooting a DCI connection to another data center The BGP session to the provider is established, but the session to Border-Leaf-2 is not established. Referring to the exhibit, which configuration change should be made to solve the problem?

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

Answer: D

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 nexthop 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:

D. 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: 18

You are asked to automatically provision new Juniper Networks devices in your network with minimal manual intervention Before you begin, which two statements are correct? (Choose two.)

A. You must have a DHCP server that provides the location of the software image and configuration files.
B. You must have a system log (syslog) server to manage system log messages and alerts.
C. You must have an NTP server to perform time synchronization.
D. You must have a file server that stores software image and configuration files.

Answer: AD

Explanation:

Zero-Touch Provisioning (ZTP):

ZTP is a feature that allows for the automatic provisioning of devices with minimal manual intervention. It is widely used in large-scale deployments to quickly bring new devices online. Key Requirements for ZTP:

A . DHCP Server: A DHCP server is crucial for ZTP as it provides the necessary information to new devices, such as the IP address, the location of the software image, and configuration files.

D . File Server: The file server is where the software image and configuration files are stored. The device downloads these files during the provisioning process.

Incorrect Options:

B . Syslog Server: While a syslog server is important for logging and monitoring, it is not a requirement for the initial provisioning process.

C . NTP Server: An NTP server is used for time synchronization, which is essential for accurate logging and operation but not specifically required for ZTP.

Data Center Reference:

ZTP simplifies the deployment process by automating the initial configuration steps, relying heavily on DHCP for communication and a file server for delivering the necessary configuration and software.

Question: 19

You are selling up an EVPN-VXLAN architecture (or your new data center. this initial deployment will be less than 50 switches: however, it could scale up to 250 switches over time supporting 1024 VLANs. You are still deciding whether to use symmetric or asymmetric routing.

In this scenario, which two statements are correct? (Choose two.)

A. Symmetric routing needs an extra VLAN with an IRB interface for each L3 VRF instance.

B. Asymmetric routing is easier lo monitor because of the transit VNI.
C. Symmetric routing supports higher scaling numbers.
D. Asymmetric routing routes traffic on the egress switch.

Answer: CD

Explanation:

Symmetric vs. Asymmetric Routing in EVPN-VXLAN:
Symmetric Routing: Traffic enters and exits the VXLAN network through the same VTEP, regardless of the source or destination. This approach simplifies routing decisions, especially in large networks, and is generally more scalable.
Asymmetric Routing: The routing occurs on the egress VTEP. This method can be simpler to deploy in smaller environments but becomes complex as the network scales, particularly with larger numbers of VNIs and VLANs.
Correct Statements:
C . Symmetric routing supports higher scaling numbers: Symmetric routing is preferred in larger EVPN-VXLAN deployments because it centralizes routing decisions, which can be more easily managed and scaled.
D . Asymmetric routing routes traffic on the egress switch: This is accurate, as asymmetric routing means the routing decision is made at the final hop, i.e., the egress VTEP before the traffic reaches its destination.
Incorrect Statements:
A . Symmetric routing needs an extra VLAN with an IRB interface for each L3 VRF instance: This is not accurate. Symmetric routing does not require an extra VLAN per VRF; rather, it uses the same VLAN/VNI across the network, simplifying routing and VLAN management.
B . Asymmetric routing is easier to monitor because of the transit VNI: Asymmetric routing is not necessarily easier to monitor; in fact, it can add complexity due to the split routing logic between ingress and egress points.
Data Center Reference:
The choice between symmetric and asymmetric routing in an EVPN-VXLAN environment depends on network size, complexity, and specific operational requirements. Symmetric routing is generally more scalable and easier to manage in large-scale deployments.

Question: 20

Your organization is implementing EVPN-VXLAN and requires multiple overlapping VLAN-IDs. You decide to use a routing-instance type mac-vrf to satisfy this request.
Which two statements are correct in this scenario? (Choose two.)
A. Host-facing interfaces must be configured using a service-provider style configuration.
B. Host-facing interfaces must be configured using enterprise-style configuration.
C. Spine-facing interfaces must be configured using an enterprise-style configuration.
D. The routing-instance service type can be VLAN-based.

Answer: AD

Explanation:

Understanding the Scenario:
EVPN-VXLAN deployments often involve scenarios where multiple tenants or applications require overlapping VLAN IDs, which can be managed using the mac-vrf routing instance type. This allows you to segregate traffic within the same VLAN ID across different tenants.
Host-facing Interface Configuration:

A . Host-facing interfaces must be configured using a service-provider style configuration: This is correct. In mac-vrf configurations, host-facing interfaces (those connecting end devices) typically follow a service-provider style configuration, where each customer or tenant's traffic is isolated even if overlapping VLAN IDs are used.
B . Host-facing interfaces must be configured using enterprise-style configuration: This is incorrect for mac-vrf instances because enterprise-style configurations are more common in simpler, less segmented networks.
Routing Instance Service Type:
D . The routing-instance service type can be VLAN-based: This is correct. The service type in mac-vrf can indeed be VLAN-based, which is particularly useful in scenarios where VLAN ID overlap is needed between different tenants or services.
Data Center Reference:
The mac-vrf instance type is powerful for handling complex multi-tenant environments in EVPN- VXLAN, especially when dealing with overlapping VLAN IDs across different segments of the network.

Question: 21

You are using a single tenant data center with a bridged overlay architecture. In this scenario, how do hosts of the different virtual networks communicate with each other?

A. off-fabric using an external device
B. using anycast gateway addresses configured on the leaf devices
C. using EVPN Type 5 routes
D. using virtual gateway addresses configured on the spine

Answer: A

Explanation:

Understanding Bridged Overlay Architecture:
In a single-tenant data center using a bridged overlay architecture, virtual networks (VLANs) are typically isolated within the fabric, with traffic between these VLANs handled outside the fabric.
Communication Between Different Virtual Networks:
A . off-fabric using an external device: This is correct. In many bridged overlay architectures, communication between different virtual networks is handled off-fabric, often using an external router or firewall that connects the different VLANs. The fabric itself primarily provides Layer 2 connectivity within each VLAN, leaving inter-VLAN routing to be handled externally.
Data Center Reference:
This design is common in smaller or simpler data center environments where a single tenant does not require complex on-fabric routing and prefers to handle inter-VLAN routing through dedicated devices.

Question: 22

A local VTEP has two ECMP paths to a remote VTEP
Which two statements are correct when load balancing is enabled in this scenario? (Choose two.)

A. The inner packet fields are not used in the hash for load balancing.
B. The destination port in the UDP header is used to load balance VXLAN traffic.
C. The source port in the UDP header is used to load balance VXLAN traffic.

D. The inner packet fields are used in the hash for load balancing.

Answer: CD

Explanation:

Load Balancing in VXLAN:

VXLAN uses UDP encapsulation to transport Layer 2 frames over an IP network. For load balancing across Equal-Cost Multi-Path (ECMP) links, various fields in the packet can be used to ensure even distribution of traffic.

Key Load Balancing Fields:

C . The source port in the UDP header is used to load balance VXLAN traffic: This is correct. The source UDP port in the VXLAN packet is typically calculated based on a hash of the inner packet's fields. This makes the source port vary between packets, enabling effective load balancing across multiple paths.

D . The inner packet fields are used in the hash for load balancing: This is also correct. Fields such as the source and destination IP addresses, source and destination MAC addresses, and possibly even higher-layer protocol information from the inner packet can be used to generate the hash that determines the ECMP path.

Incorrect Statements:

A . The inner packet fields are not used in the hash for load balancing: This is incorrect as the inner packet fields are indeed critical for generating the hash used in load balancing.

B . The destination port in the UDP header is used to load balance VXLAN traffic: This is incorrect because the destination UDP port in VXLAN packets is typically fixed (e.g., port 4789 for VXLAN), and therefore cannot be used for effective load balancing.

Data Center Reference:

Effective load balancing in VXLAN is crucial for ensuring high throughput and avoiding congestion on specific links. By using a combination of the source UDP port and inner packet fields, the network can distribute traffic evenly across available paths.

Question: 23

Exhibit.

Exhibit

```
user@leaf1> show configuration
...
interfaces {
    ge-0/0/0 {
        description "facing_spine1:ge-0/0/1";
        speed 10g;
        mtu 9192;
        unit 0 {
            family inet {
                mtu 9170;
                address 172.16.0.9/31;
            }
        }
    }
    ge-0/0/1 {
        description "facing_spine2:ge-0/0/1";
        speed 10g;
        mtu 9192;
        unit 0 {
            family inet {
                mtu 9170;
                address 172.16.0.11/31;
            }
        }
    }
    irb {
        unit 200 {
            family inet {
                address 192.168.200.1/24;
            }
        }
    }
}
vlans {
    vn100 {
        vlan-id 100;
        description "BLUE";
    }
    vn200 {
        description RED;
        vlan-id 200;
        l3-interface irb.200;
    }
}
```

Host A is connected to vlan 100 on lead. Host B is connected to vlan 200 on leaf1. Host A and Host B ate unable to communicate. You have reviewed the touting and your hosts have the correct default route (.1) Referring to the exhibit, which two commands will solve the problem? (Choose two.)

A. delete vlans vn200 13-interface irb.200
B. set interfaces irb unit 100 family inet address 192-168.100.1
C. set routing-options static route 0.0.0.0/0 next-hop 192.168.200.10

D. set vlans vn100 13-interface irb.100

Answer: CD

Explanation:

In the provided network configuration, Host A is in VLAN 100 and Host B is in VLAN 200. The issue arises because these two hosts are unable to communicate, which indicates that either the interfaces are not properly linked to their respective VLANs, or there is a missing static route required for interVLAN routing.

Step-by-Step Analysis:

VLAN Assignment:

The exhibit shows that irb.200 is correctly associated with VLAN 200 in the configuration. However, there is no corresponding irb.100 for VLAN 100. Without irb.100, the network lacks the logical interface to handle routing for VLAN 100. Thus, adding irb.100 to VLAN 100 is necessary.

Command to solve this:

set vlans vn100 13-interface irb.100

Static Route Configuration:

For inter-VLAN routing to occur, a static route needs to be configured that allows traffic to pass between different subnets (in this case, between VLAN 100 and VLAN 200). The command set routing-options static route 0.0.0.0/0 next-hop 192.168.200.10 would add a static route that directs all traffic from VLAN 100 to the correct gateway (192.168.200.10), which is necessary to route traffic between the two VLANs.

Command to solve this:

set routing-options static route 0.0.0.0/0 next-hop 192.168.200.10

Explanation of Incorrect Options:

Option A (delete vlans vn200 13-interface irb.200): This would remove the logical interface associated with VLAN 200, which is not desired because we need VLAN 200 to remain active and properly routed.

Option B (set interfaces irb unit 100 family inet address 192-168.100.1): This command would incorrectly assign an IP address that does not correspond with the subnet of VLAN 100 (192.168.200.1/24). This could create a misconfiguration, leading to routing issues.

Data Center Reference:

For a Data Center, proper VLAN management and static routing are crucial for ensuring that different network segments can communicate effectively, especially when dealing with separated subnets or zones like in different VLANs. This aligns with best practices in DCIM (Data Center Infrastructure Management) which stress the importance of proper network configuration to avoid downtime and ensure seamless communication between all critical IT infrastructure components.

Ensuring that the correct interfaces are associated with the correct VLANs and having the proper static routes in place are both essential steps in maintaining a robust and reliable data center network.

This detailed analysis reflects best practices as noted in standard data center design and network configuration guides.

Question: 24
Exhibit.

Exhibit

```
routing-instances {
    tenant1 {
        instance-type vrf;
        routing-options {
            auto-export {
                family inet {
                    unicast;
                }
            }
        }
        protocols {
            evpn {
                ip-prefix-routes {
                    advertise direct-nexthop;
                    vni 10010;
                }
            }
        }
        interface lo0.10;
        route-distinguisher 192.168.100.14:5001;
        vrf-target target:65000:1;
    }
}
```

You want to enable the border leaf device to send Type 5 routes of local networks to the border leaf device in another data center. What must be changed to the configuration shown in the exhibit to satisfy this requirement?

A. Move vrf-target target: 65000:1 to the evpn hierarchy.
B. Add a VLAN configuration with an 13-interface to the tenant1 routing instance.
C. Add encapsulation vxlan to the evpn hierarchy.
D. Change: 5001 in the route-distinguisher to : 10010.

Answer: A

Explanation:

In this scenario, you want the border leaf device to advertise Type 5 EVPN routes to another border leaf in a different data center. Type 5 routes in EVPN are used to advertise IP prefixes, which means that for proper route advertisement, you need to configure the correct settings within the evpn hierarchy.

Step-by-Step Analysis:

Understanding EVPN Type 5 Routes:

EVPN Type 5 routes are used to advertise IP prefixes across EVPN instances, which allow different data centers or networks to exchange routing information effectively.

VRF Target Setting:

The vrf-target configuration is crucial because it defines the export and import policies for the VRF within the EVPN instance. For EVPN Type 5 routes to be advertised to other border leaf devices, the vrf-target needs to be correctly configured under the evpn hierarchy, not just within the routing instance.

Command to solve this:

move vrf-target target:65000:1 to evpn

Other Options:

Option B: Adding a VLAN configuration would not address the requirement to advertise Type 5 routes.

Option C: Adding VXLAN encapsulation may be necessary for other scenarios but does not directly address the Type 5 route advertisement.

Option D: Changing the route-distinguisher will differentiate routes but does not impact the advertisement of Type 5 routes to other data centers.

By moving the vrf-target to the evpn hierarchy, you enable the proper route advertisement, ensuring that the Type 5 routes for local networks are shared with other data center border leaf devices. This is aligned with best practices for multi-data center EVPN implementations, which emphasize the correct placement of routing policies within the EVPN configuration.

Question: 25

Exhibit.

Exhibit

```
user@leaf1> show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)
Ethernet switching table : 6 entries, 6 learned
Routing instance : default-switch
   Vlan             MAC                MAC      Logical                SVLBNH/       Active
   name             address            flags    interface              VENH Index    source
   v10              00:00:5e:00:01:01  DRP      esi.1777
                        05:00:00:fd:e9:00:00:13:92:00
   v10              00:0c:29:e8:b7:39  D        xe-0/0/4.0
   v10              02:05:86:d9:1b:00  DR       vtep.32769                          192.168.100.13
   v20              00:00:5e:00:01:01  DRP      esi.1759
                        05:00:00:fd:e9:00:00:13:9c:00
   v20              00:0c:29:08:04:a0  DR       vtep.32769                          192.168.100.13
   v20              02:05:86:d9:1b:00  DR       vtep.32769                          192.168.100.13
```

Referring to the exhibit, why is the active source field blank for the entry that uses the 00:0c:29:e8:b7:39 MAC address?

A. The EVPN route for this host does not have a valid next hop.
B. The ARP lookup for this host has failed.
C. The host for this entry is locally connected to leaf1.
D. This entry is associated with a multicast EVPN route.

Answer: A

Explanation:

In this scenario, the active source field is blank for the MAC address 00:0c:29:e8:b7:39, indicating an issue with how this MAC entry is being processed within the EVPN/VXLAN environment.

Step-by-Step Analysis:

Understanding the MAC Entry:

The active source field should normally indicate the source of the route advertisement for a specific MAC address within the EVPN. If it is blank, it suggests that there is a problem with how this entry is being learned or propagated.

Possible Issues:

Option A: If the EVPN route for this MAC address does not have a valid next hop, the entry might exist in the MAC table, but it will not have a valid path for forwarding, leading to a blank active source.

Option B: If the ARP lookup had failed, the entry might not even appear in the MAC table. However, the entry does exist, suggesting that ARP is not the primary issue here.

Option C: If the host were locally connected, the active source should reflect a local interface, but the field is blank, ruling out local connection as the cause.

Option D: Multicast EVPN routes typically do not appear in this manner in the MAC table, and this would not cause the active source to be blank.

Conclusion: The most logical explanation is that the EVPN route for this host exists but does not have a valid next hop, leading to the absence of an active source. This is consistent with how EVPN routing tables work in a VXLAN environment, where the lack of a valid next hop would prevent proper route advertisement and forwarding for the specific MAC address.

Question: 26

You are deploying an IP fabric with an oversubscription ratio of 3:1.

In this scenario, which two statements are correct? (Choose two.)

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

Answer: CD

Explanation:

Understanding Oversubscription Ratio in IP Fabrics:

The oversubscription ratio in an IP fabric typically refers to the ratio of the available bandwidth at the edge of the network (leaves) to the available bandwidth at the core or spine. A 3:1 oversubscription ratio means that for every 3 units of bandwidth at the leaves, there is 1 unit of bandwidth at the spine.

Impact of Adding or Removing Leaf Devices:

Removing Leaf Devices: When you remove leaf devices, the amount of total edge bandwidth decreases while the bandwidth in the spine remains constant. This causes the oversubscription ratio to increase because there is now less total bandwidth to distribute across the same amount of spine bandwidth.

Adding Leaf Devices: Conversely, when you add leaf devices, the total edge bandwidth increases.

Since the spine bandwidth remains the same, the oversubscription ratio would remain the same if the additional

leaves consume their share of the available bandwidth proportionally.

Conclusion:

Option C: Correct—Removing leaf devices increases the oversubscription ratio.

Option D: Correct—Adding leaf devices typically maintains the oversubscription ratio assuming uniform bandwidth distribution.

Question: 27

You are asked to interconnect two of your company's data centers across the IP backbone. Both data centers have their own unique IP space and do not require any bridging. In this scenario, which two actions would accomplish this task? (Choose two.)

A. Configure a Type 2 EVPN route for each unique prefix.
B. Configure peering for EVPN between border leaf nodes in each data center.
C. Configure a Type 5 EVPN route for each unique prefix.
D. Configure peering for EVPN between all leaf nodes within each data center.

Answer: BC

Explanation:

Interconnecting Data Centers:

The scenario requires interconnecting two data centers with unique IP spaces across an IP backbone.

The key point is that bridging is not required, so Layer 3 routing methods must be used.

EVPN Configuration:

Option B: Establishing EVPN peering between the border leaf nodes in each data center is the most appropriate solution as it allows for exchanging routing information between the two data centers. This ensures that the routes are properly distributed without the need for L2 bridging.

Option C: Configuring Type 5 EVPN routes is necessary for advertising IP prefixes (Layer 3 routes) across the EVPN. Type 5 routes allow for the exchange of IP prefixes between the two data centers, enabling the necessary routing functionality without the need for bridging.

Conclusion:

Option B: Correct—Peering between border leaf nodes sets up the necessary route exchange between data centers.

Option C: Correct—Type 5 EVPN routes are essential for exchanging Layer 3 prefixes between data centers.

Question: 28

Which three statements are correct about symmetric IRB routing with EVPN Type 2 routes? (Choose three.)

A. An L3 interface (IRB) is required for each local VLAN.
B. Symmetric routing requires MAC-VRF.
C. Symmetric routing supports the EVPN service VLAN bundle.
D. Symmetric routing requires an extra transit VNI for each VRF.
E. Symmetric routing is less efficient than asymmetric routing.

Answer: ABD

Explanation:

Symmetric IRB Routing with EVPN Type 2 Routes:

Symmetric Routing: In symmetric IRB (Integrated Routing and Bridging), routing occurs in both directions at the ingress and egress leaf nodes using the same routing logic. This is contrasted with asymmetric routing, where different routing logic is used depending on the direction of the traffic. Required Components:

Option A: An L3 IRB interface is necessary for each VLAN that participates in routing, as it handles the Layer 3 processing for the VLAN.

Option B: MAC-VRF is required for symmetric routing to maintain a mapping of MAC addresses to the appropriate VRF, ensuring correct forwarding within the EVPN.

Option D: A transit VNI (Virtual Network Identifier) is required for each VRF to encapsulate the Layer 3 traffic as it traverses the network, allowing the IP traffic to be appropriately forwarded.

Conclusion:

Option A: Correct—Each local VLAN needs an IRB interface for L3 processing.

Option B: Correct—MAC-VRF is necessary for handling MAC address resolution in symmetric routing.

Option D: Correct—Transit VNIs are required for routing VRF-specific traffic across the network.

Options C and E are incorrect because:

C: Symmetric routing can work with various VLAN models, including single or multiple VLANs within an EVPN instance.

F. Symmetric routing is generally more efficient than asymmetric routing as it uses consistent routing logic in both directions.

Question: 29

You are asked to interconnect two of your company's data centers across an IP backbone. Both data centers require Layer 2 and Layer 3 connectivity. In this scenario, which three actions would accomplish this task? (Choose three.)

A. Advertise Type 2 EVPN routes across the DCI.
B. Ensure border leaf nodes in each data center can exchange EVPN routes.
C. Ensure there is a full mesh of VTEPs between all spine nodes within both data centers.
D. Advertise Type 5 EVPN routes across the DCI.
E. Ensure there is a full mesh of VTEPs between all leaf nodes within data centers.

Answer: ABD

Explanation:

Layer 2 and Layer 3 Connectivity Requirements:

To interconnect two data centers across an IP backbone with both Layer 2 (L2) and Layer 3 (L3) connectivity, EVPN-VXLAN (Ethernet VPN with Virtual Extensible LAN) is the ideal solution. EVPN supports L2 VPNs while also enabling L3 connectivity across multiple locations.

Necessary EVPN Route Types:

Type 2 EVPN Routes: These routes are used to advertise MAC addresses for Layer 2 connectivity.

They are essential for enabling seamless L2 communication across data centers.

Type 5 EVPN Routes: These routes are necessary for advertising IP prefixes for Layer 3 connectivity between data

centers. They enable the exchange of L3 information across the IP backbone, ensuring routed traffic can reach its destination.

Border Leaf Nodes:

Border Leaf Nodes: Ensuring that the border leaf nodes (the entry and exit points for traffic between data centers) can exchange EVPN routes is critical for the correct dissemination of both L2 and L3 information across the data centers.

Conclusion:

Option A: Correct—Type 2 EVPN routes are required for Layer 2 MAC address learning and communication across the DCI (Data Center Interconnect).

Option B: Correct—Border leaf nodes need to exchange EVPN routes to maintain connectivity between data centers.

Option D: Correct—Type 5 EVPN routes are essential for Layer 3 connectivity across the DCI.

Options C and E are incorrect because they refer to establishing full mesh VTEPs (VXLAN Tunnel Endpoints) across all spine or leaf nodes, which is unnecessary for the scenario provided. The focus should be on border leaf nodes and appropriate route advertisements for L2 and L3 connectivity.

Question: 30

You are asked to build redundant gateways in your EVPN-VXLAN environment, but you must conserve address space because these gateways must span across seven PES. What should you implement on the PEs lo satisfy these requirements?

A. Use IRB interfaces with the same IP address and different MAC addresses.
B. Use IRB interfaces with the same IP and VGA.
C. Use IRB interfaces with the same IP and MAC address.
D. Use IRB interfaces with different IP addresses and the same VGA.

Answer: C

Explanation:

Redundant Gateways in EVPN-VXLAN:

In an EVPN-VXLAN environment, providing redundant gateway functionality typically involves the use of Anycast Gateway. This allows multiple PEs (Provider Edge devices) to use the same IP address and MAC address for the gateway, enabling seamless failover and redundancy without IP conflicts. Conserving Address Space:

Using the same IP address across multiple PEs conserves address space because only one IP address is needed for the gateway function, regardless of the number of PEs. The shared MAC address ensures that ARP resolution and forwarding behavior are consistent across all the PEs.

Conclusion:

Option C: Correct—Using IRB interfaces with the same IP and MAC address across all PEs satisfies the need for redundancy while conserving address space.

Options A, B, and D introduce unnecessary complexity or do not fully utilize the efficient Anycast Gateway approach, which is best practice for conserving IP space and providing redundancy.

Question: 31

You are asked to identify microburst traffic occurring in the network leading lo packet drops in your data center

switches Which two tools would be used in this scenario? (Choose two.)

A. port mirroring
B. Traceoptions
C. port buffer monitoring
D. syslog

Answer: AC

Explanation:

Identifying Microburst Traffic:

Microbursts are short spikes in network traffic that can overwhelm buffers and cause packet drops. Detecting and analyzing microbursts is crucial for understanding where packet loss might be occurring in a data center network.

Port Buffer Monitoring:

Port Buffer Monitoring: This tool specifically tracks the usage of switch buffers, helping to identify when microbursts are causing buffers to overflow, leading to packet drops.

Port Mirroring:

Port Mirroring: This tool allows you to monitor real-time traffic on a specific port by copying the traffic to another port where it can be analyzed, often with a packet analyzer. While port mirroring doesn't directly detect microbursts, it helps capture traffic patterns that can indicate microbursts. Conclusion:

Option C: Correct—Port buffer monitoring directly identifies buffer overflows caused by microbursts. Option A: Correct—Port mirroring allows for the detailed capture and analysis of traffic patterns, which can reveal microburst behavior.

Options B (Traceoptions) and D (Syslog) are less effective in identifying microburst traffic. Traceoptions focus on control plane traffic debugging, and Syslog is more about logging system events than detecting high-frequency traffic spikes.

Question: 32

Exhibit.

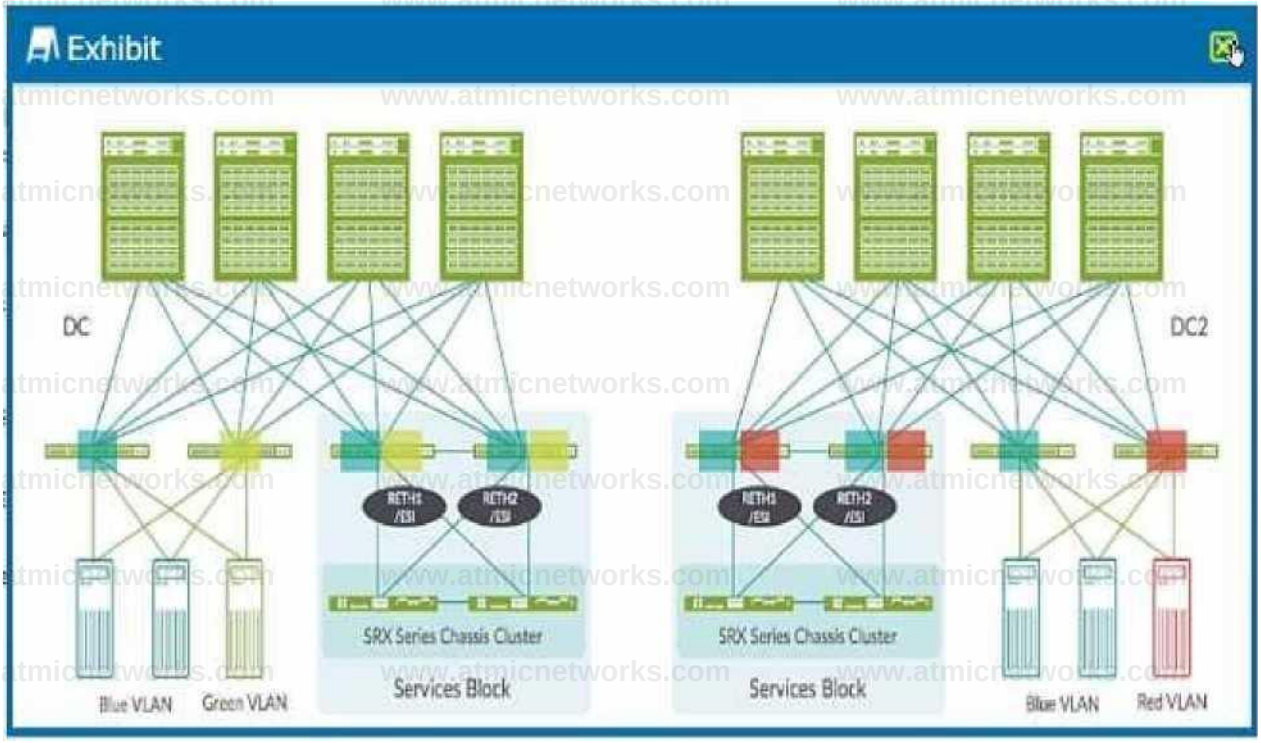

Both DC and DC2 ate using EVPN-VXLAN technology deployed using an ERB architecture. A server on the Red VLAN must communicate with a server on the Green VLAN. The Blue VLAN in DC and DC2 needs to be the same VLAN.

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.

Answer: B

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: 33

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.

Answer: B

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: 34

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.

Answer: D

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: 35

Exhibit.

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

Referring to the exhibit, which two configuration changes are required to solve the problem? (Choose two.)

A. Configure the set switch-options route-distinguisher 192.168.100.51:2 parameter on Leaf-1.
B. Configure the set switch-options service-id 1 parameter on Leaf-2.
C. Configure the set switch-options vtep-source-interface 100. 0 parameter on Leaf-1.
D. Configure the set switch-options vrf-target target: 65000:55 parameter on Leaf-2.

Answer: BD

Explanation:

Review of the Exhibit:
The exhibit shows the switch configuration for Leaf-1 and Leaf-2. The configurations include route distinguishers, VRF targets, and service IDs, all of which are crucial for ensuring proper operation in an EVPN-VXLAN environment.
Service-ID Consistency:
The service ID must be consistent across all participating leaf devices in the same EVPN instance to ensure that they are part of the same VXLAN overlay network.
VRF Target Consistency:
The vrf-target parameter must also be consistent across devices to ensure that VRFs (Virtual Routing and Forwarding instances) are correctly imported and exported between leaf nodes.
Conclusion:
Option B: Correct—Setting the same service-id on Leaf-2 ensures that it is part of the same VXLAN overlay as Leaf-1.
Option D: Correct—The vrf-target on Leaf-2 should match Leaf-1 to ensure consistent routing policies and proper route exchange.

Question: 36

Which two statements are true about a pure IP fabric? (Choose two.)

A. Devices in an IP fabric function as Layer 3 routers.
B. An IP fabric supports Layer 2 VLANs.
C. Devices in an IP fabric must be connected to a fabric controller.
D. An IP fabric does not support Layer 2 protocols.

Answer: AD

Explanation:

Understanding Pure IP Fabric:
A pure IP fabric is a network design where all devices operate at Layer 3, meaning that each device in the fabric is a router that makes forwarding decisions based on IP addresses.
Layer 2 Support:
In a pure IP fabric, traditional Layer 2 protocols such as Spanning Tree Protocol (STP) or VLANs are not supported. Instead, the network relies entirely on Layer 3 routing protocols to manage traffic between devices.
Routing Functionality:
Since devices in an IP fabric operate as Layer 3 routers, they handle IP routing and provide network services based on IP addresses, not on MAC addresses or Layer 2 switching.
Conclusion:
Option A: Correct—Devices in an IP fabric function as Layer 3 routers.

Option D: Correct—A pure IP fabric does not support traditional Layer 2 protocols, making it a purely routed environment.

Question: 37

Which two statements are true about IP fabrics using unnumbered BGP? (Choose two.)

A. Unnumbered BGP requires that family inet6 is configured on each interface.
B. Unnumbered BGP peering automatically provisions IPv6 peering.
C. Unnumbered BGP requires that family inet is configured on each interface.
D. Unnumbered BGP peering automatically provisions IPv4 peering.

Answer: CD

Explanation:

Understanding Unnumbered BGP:
Unnumbered BGP (Border Gateway Protocol) allows BGP peering between routers without assigning specific IP addresses to the interfaces. Instead, it uses the loopback address or another router identifier for the BGP session, making IP address management more straightforward in large-scale networks.
Family inet Configuration:
Option C: The family inet configuration is required on each interface involved in unnumbered BGP peering to support IPv4 address families. This ensures that IPv4 peering sessions can be established between devices.
Automatic IPv4 Peering:
Option D: Unnumbered BGP peering automatically provisions IPv4 peering sessions. This simplifies the configuration by eliminating the need to manually assign and manage IP addresses for BGP peering.
Conclusion:
Option C: Correct—Unnumbered BGP requires the family inet configuration for IPv4.
Option D: Correct—Unnumbered BGP automatically provisions IPv4 peering, simplifying setup.

Question: 38

You are asked to implement VXLAN group-based policies (GBPs) in your data center. Which two statements are correct in (his scenario? (Choose two.)

A. VXLAN GBP uses scalable group tags that must be configured statically on each switch and activated through 802.1X.
B. VXLAN GBP uses scalable group tags that may be configured on a RADIUS server and pushed to the switch through 802.1X.
C. VXLAN GBP ensures consistent application of security group policies throughout the network.
D. VXLAN GBP ensures consistent application of BGP groups throughout the network.

Answer: BC

Explanation:

VXLAN Group-Based Policies (GBP):

VXLAN Group-Based Policies are used to apply security policies consistently across the network.

These policies are often tied to user or device identities rather than static IP addresses, which allows for more dynamic and scalable security management.

Scalable Group Tags via RADIUS and 802.1X:

Option B: VXLAN GBP can use scalable group tags configured on a RADIUS server, which are then pushed to network devices through 802.1X. This allows for centralized and automated policy application based on user or device identity.

Consistent Security Policy Application:

Option C: GBP ensures that security policies are consistently applied across the network, regardless of where a user or device connects. This consistency is crucial in environments where security policies must follow the user or device.

Conclusion:

Option B: Correct—Group tags can be configured on a RADIUS server and pushed via 802.1X, enabling centralized policy management.

Option C: Correct—GBP ensures consistent application of security policies, which is essential for maintaining security across a dynamic network environment.

Question: 39

You are using E8GP peering in an underlay IP fabric. Which two statements are correct in this scenario? (Choose two.)

A. EBGP peering requires an IGP protocol tor adjacency establishment.
B. EBGP peering does not require an IGP protocol tor adjacency establishment.
C. Every leaf node has one peering session to every spine node.
D. Every leaf node has a peering session to every other leaf node.

Answer: BC

Explanation:

Understanding EBGP in an IP Fabric:

EBGP (External Border Gateway Protocol) is commonly used in IP fabrics to establish peering between routers, such as leaf and spine nodes, without relying on an Interior Gateway Protocol (IGP) like OSPF or IS-IS.

IGP Requirement for EBGP:

Option B: EBGP peering does not require an IGP for adjacency establishment. This is because EBGP peers are typically directly connected, and BGP establishes its own sessions without needing an underlying IGP.

Leaf-to-Spine Peering:

Option C: In a typical IP fabric, each leaf node establishes an EBGP session with every spine node.

This ensures full connectivity between leaves and spines, facilitating efficient routing and forwarding within the fabric.

Conclusion:

Option B: Correct—EBGP does not require an IGP for establishing peering sessions.

Option C: Correct—Each leaf node peers with every spine node, which is a standard practice in IP fabrics to ensure connectivity and redundancy.

Question: 40

Which two statements are true about EVPN routes for Data Center Interconnect? (Choose two.)

A. Type 5 EVPN routes require a VXLAN tunnel to the protocol next hop.
B. Type 2 EVPN routes do not require a VXLAN tunnel to the protocol next hop.
C. Type 2 EVPN routes require a VXLAN tunnel to the protocol next hop.
D. Type 5 EVPN routes do not require a VXLAN tunnel to the protocol next hop.

Answer: BD

Explanation:

Type 2 EVPN Routes:
Type 2 routes advertise MAC addresses within an EVPN instance and are used primarily for Layer 2 bridging. These routes do not require a VXLAN tunnel to the protocol next hop because they operate within the same Layer 2 domain.
Type 5 EVPN Routes:
Type 5 routes are used to advertise IP prefixes (Layer 3 routes) within EVPN. Similar to Type 2 routes, they do not require a VXLAN tunnel to the protocol next hop because they represent L3 routes, which are managed at the routing layer without the need for VXLAN encapsulation.
Conclusion:
Option B: Correct—Type 2 routes do not need a VXLAN tunnel to the next hop, as they are used for Layer 2.
Option D: Correct—Type 5 routes also do not need a VXLAN tunnel because they operate at Layer 3, handling IP prefixes.

Question: 41

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

Answer: BD

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: 42

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.

Answer: AB

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: 43

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

Answer: BD

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: 44

You are deploying a new network lo support your AI 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

Answer: D

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: 45

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

Answer: AD

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: 46

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.

Answer: BD

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.

Question: 47

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.

Answer: D

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: 48

Exhibit.

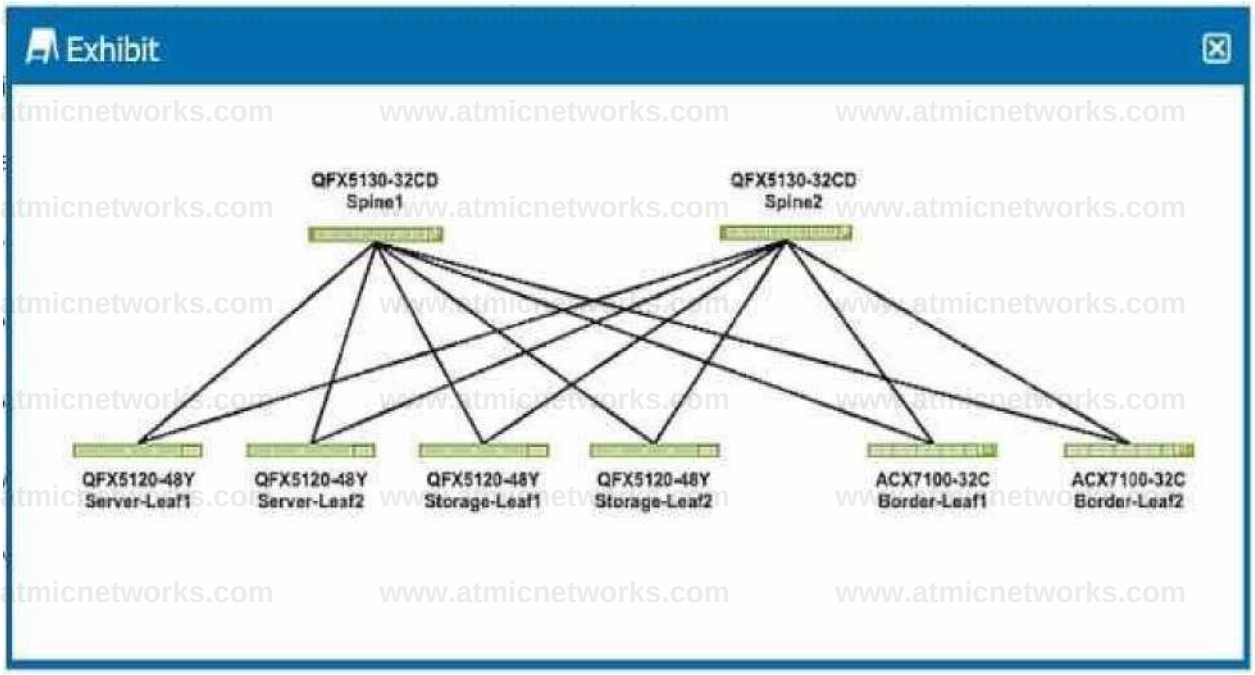

You are deploying a VXLAN overlay with EVPN as the control plane in an ERB architecture.

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.

Answer: BCE

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: 49

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

Answer: A

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: 50

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.

Answer: BD

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: 51

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

Answer: A

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.

Vendor-Neutral Data Models:

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: 52

Exhibit.

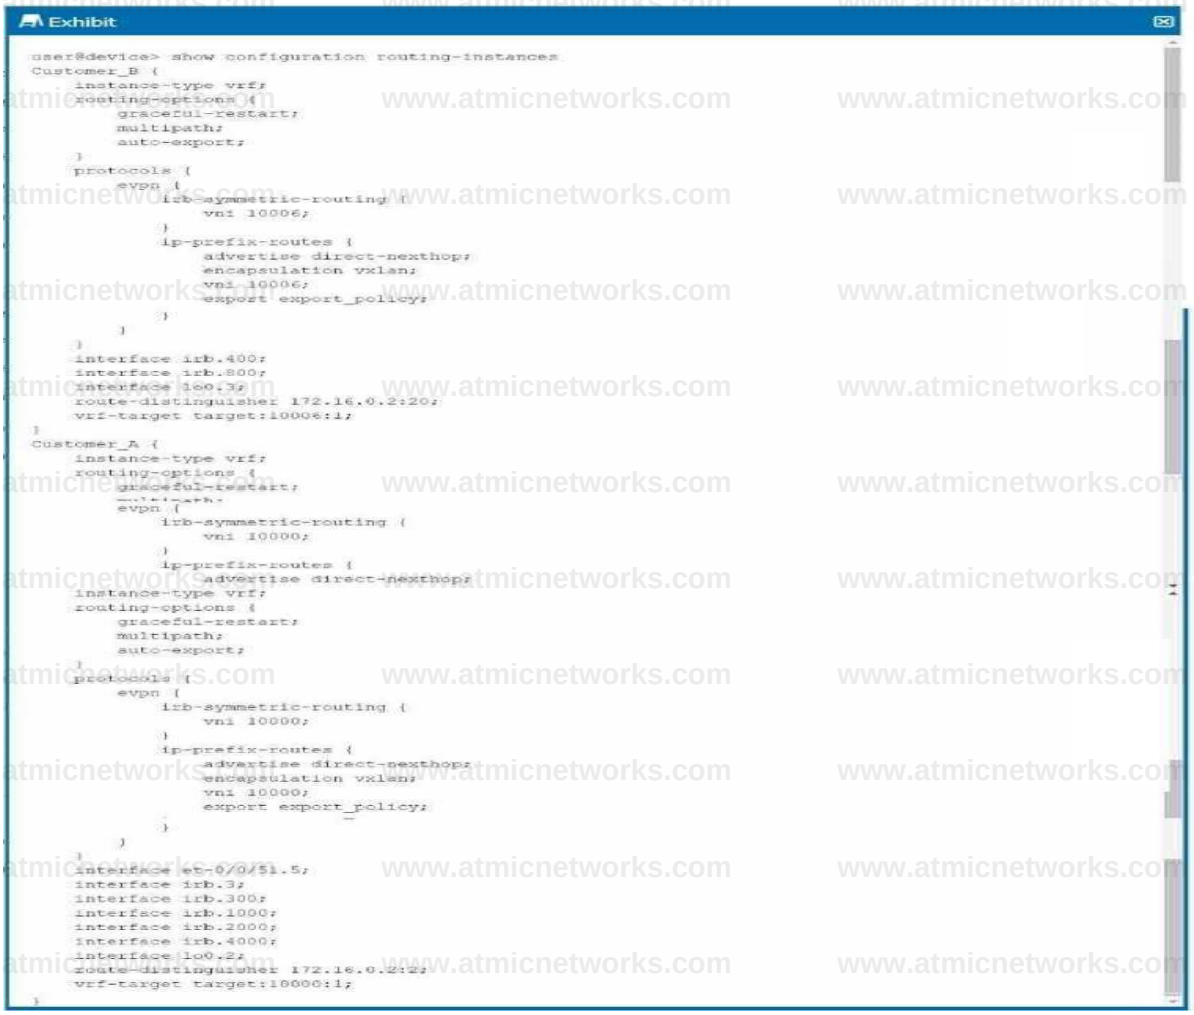

Referring to the configuration shown in the exhibit, assume that there is no external router present, and that the configuration is fabric-only.

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.

Answer: BD

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: 53

Exhibit.

W Exhibit

```
user!leafl> show evpn database Instance: evpn-1
VLM Domain I d  MAC address        Active source                    TimeECamp          IE addre
    10001       00:lc:73:00:00:01  irb.4000                         Apr 16 11:46:14    10.4.4.1
    10001       40:00:dc:01:00t01  00:02:00:00:00:00:04:00:00:04    Apr 16 11:46:14    10.4.4.2
    10001       40:DO:dc:01:00:Cl  00:02:00:00:00:00:04:00:00:04    "I- 16 11:46:14    10.4.4.3
    10001       40:D0:dc:01:0D:O3  GG:02:CO:GO:n0:00:04:00:00:04    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  GC:02:f):K':}3:0i:H:89:W:u4      Apr 16 11:46:14
    10001       44:11:01: ■'■;00:02 00:02:00:00:00:00:04:00:00:04   Api16 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  Oi: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    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 _  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:lc:73:00:00:01  lrb.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:D0:dc:01:00:02  00:02:00:00:00:00:01:00:00:01    Apr 16 11:46:14
    10002       3D: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.

Answer: AD

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 multihoming 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: 54

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.

Answer: B

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 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: 55

Exhibit.

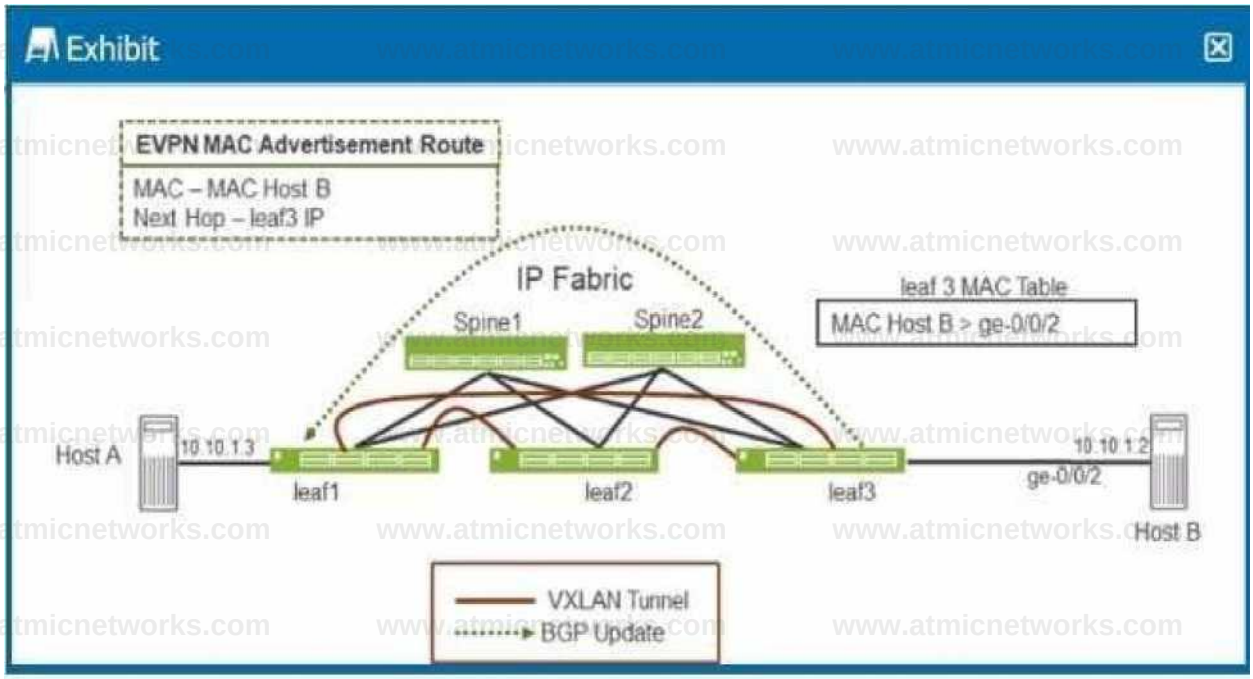

Referring to the exhibit, when Host A sends an ARP request for Host B's IP address, which Junos feature does leaf1

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

Answer: A

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: 56

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.

Answer: BC

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: 57

Exhibit.

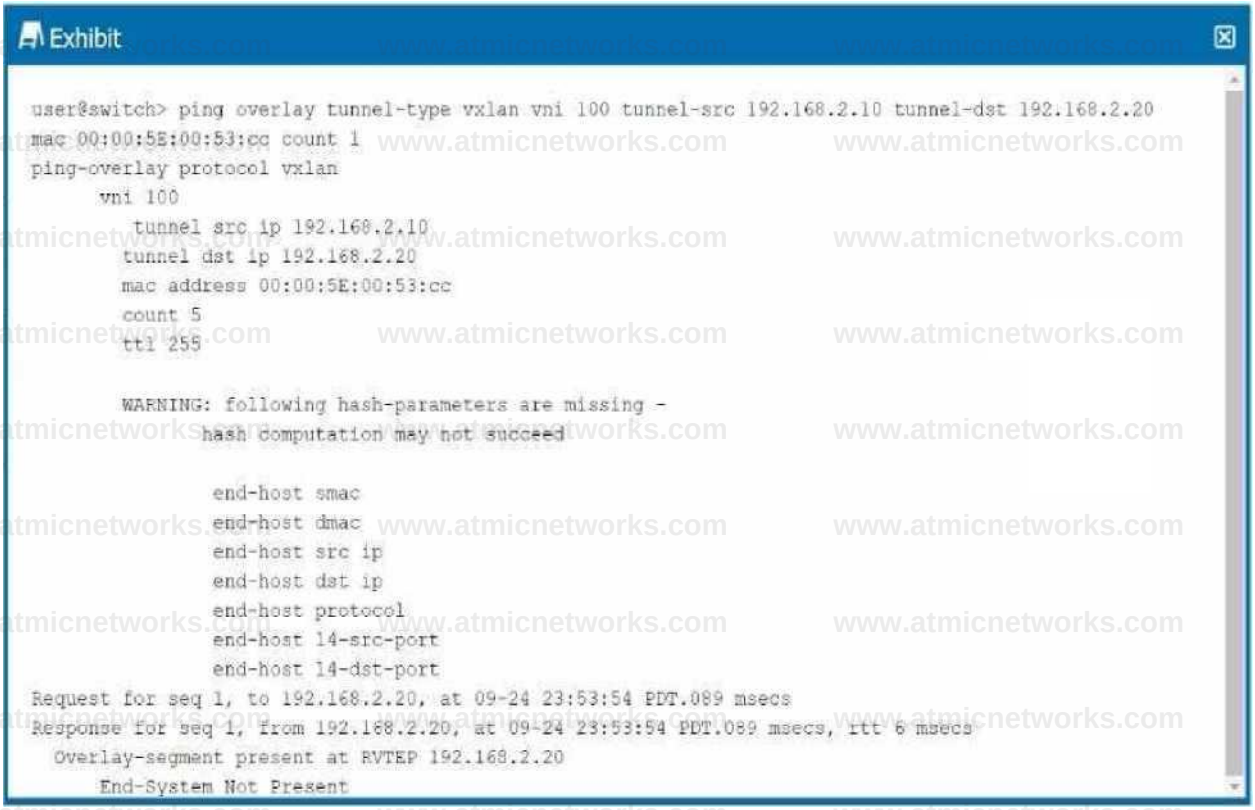

Referring to the exhibit, which statement is correct?

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

Answer: B

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: 58

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

Answer: D

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: 59

Exhibit.

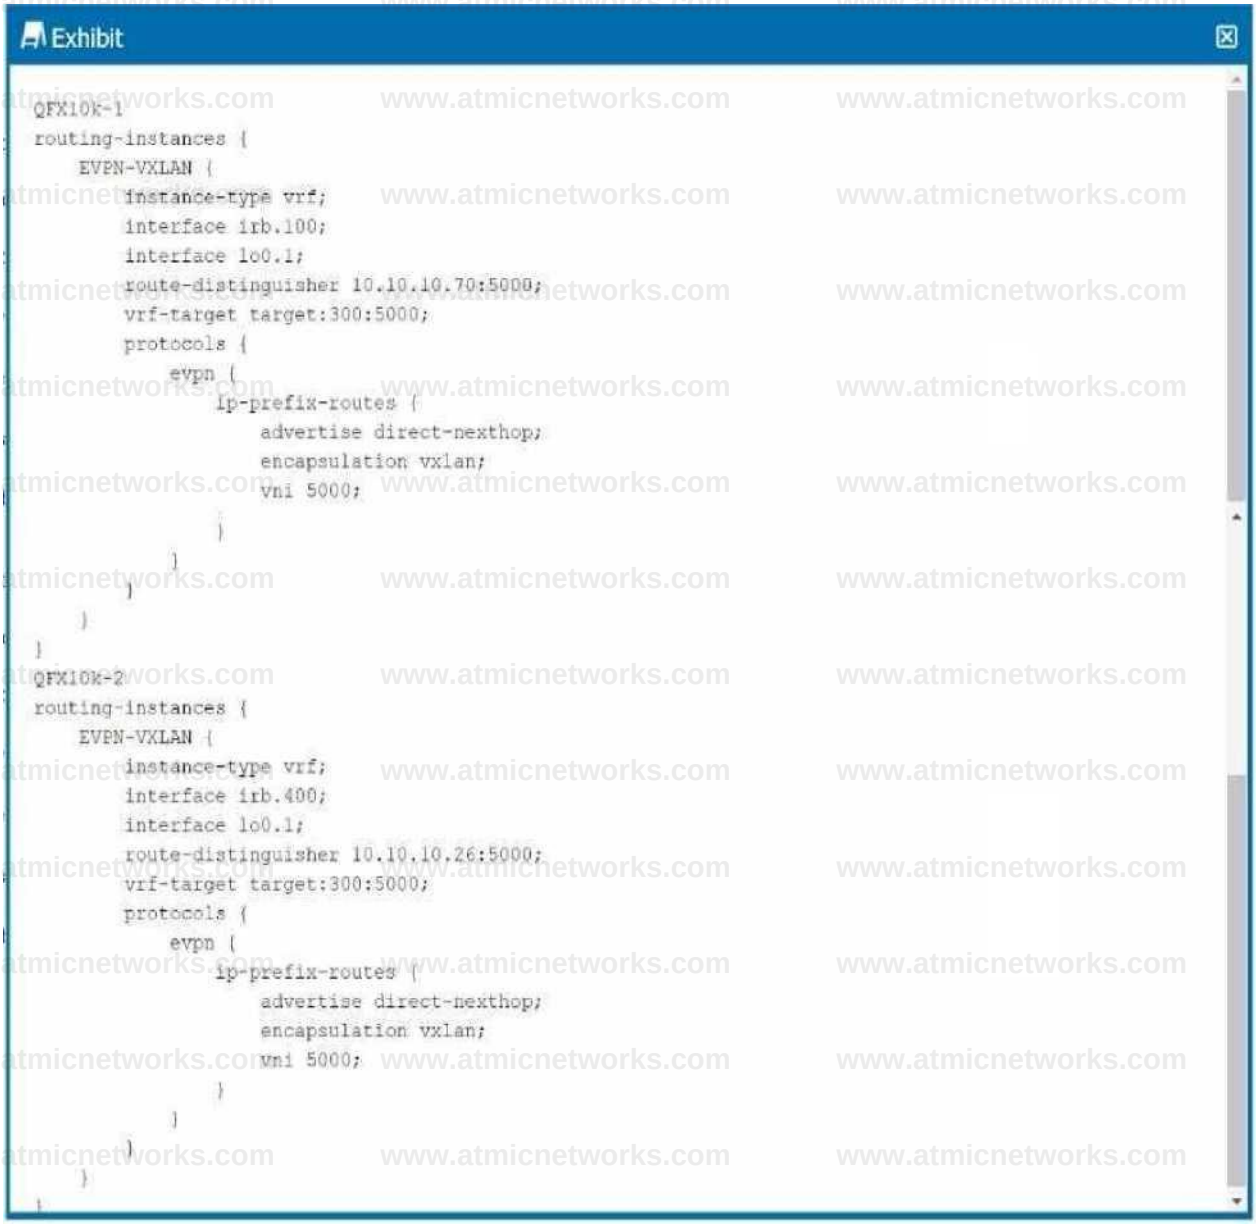

You have a sample configuration for connecting two sites through EVPN-VXLAN by exchanging IP prefix routes. Referring to the exhibit, which two statements regarding the configuration are true? {Choose two.)

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.

Answer: AC

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: 60

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.

Answer: B

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: 61

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.

Answer: BD

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: 62

You are asked for TX and RX traffic statistics for each interface to which an application server is attached. The statistics need to be reported every five seconds. Using the Junos default settings, which telemetry method would accomplish this request?

A. gNMI
B. SNMP
C. Native Sensors
D. OpenConfig

Answer: C

Explanation:

Telemetry Methods in Junos:
Telemetry is used to collect and report data from network devices. For high-frequency statistics reporting, such as every five seconds, you need a telemetry method that supports this level of granularity and real-time monitoring.
Junos Native Sensors:
Option C: Native Sensors in Junos provide detailed, high-frequency telemetry data, including TX and RX traffic statistics for interfaces. They are designed to offer real-time monitoring with customizable sampling intervals, making them ideal for the five-second reporting requirement.
Conclusion:
Option C: Correct—Native Sensors in Junos are capable of providing the required high-frequency telemetry data every five seconds.

Question: 63

Which three statements are correct about VXLAN control planes? (Choose three.)

A. EVPN is inefficient and does not scale well.
B. Both multicast and EVPN can facilitate MAC learning.
C. Multicast is not agile and requires manual VNI mapping.
D. EVPN enables fast convergence and updates.
E. Multicast does not require as many resources.

Answer: BDE

Explanation:

VXLAN Control Planes:
VXLAN (Virtual Extensible LAN) uses different control planes to handle MAC learning and traffic forwarding. The control planes include multicast and EVPN (Ethernet VPN).
Multicast and EVPN Comparison:
Option B: Both multicast and EVPN can be used for MAC learning in a VXLAN environment. Multicast is a more traditional approach, while EVPN is more advanced and supports distributed MAC learning. Option D: EVPN offers benefits such as fast convergence and rapid updates, making it more efficient and scalable for modern data

center environments.

Option E: Multicast does not require as many resources because it relies on traditional Layer 3 multicast mechanisms to distribute broadcast, unknown unicast, and multicast (BUM) traffic. However, it can be less flexible and less scalable compared to EVPN.

Conclusion:

Option B: Correct—Both control planes facilitate MAC learning.

Option D: Correct—EVPN provides fast convergence and updates.

Option E: Correct—Multicast is resource-efficient but less flexible.

Question: 64

You are asked to set up an IP fabric that supports AI or ML workloads. You have chosen to use lossless Ethernet in this scenario, which statement is correct about congestion management?

A. The switch experiencing the congestion notifies the source device.
B. Only the source and destination devices need ECN enabled.
C. ECN marks packets based on WRED settings.
D. ECN is negotiated only among the switches that make up the IP fabric for each queue.

Answer: A

Explanation:

Understanding Lossless Ethernet and Congestion Management:

Lossless Ethernet is crucial for AI and ML workloads, where packet loss can significantly degrade performance. To implement lossless Ethernet, congestion management protocols like ECN (Explicit Congestion Notification) are used.

Role of ECN in Congestion Management:

Option A: In an IP fabric that supports lossless Ethernet, when a switch experiences congestion, it can mark packets using ECN. This marking notifies the source device of the congestion, allowing the source to reduce its transmission rate, thereby preventing packet loss.

Conclusion:

Option A: Correct—The switch experiencing congestion notifies the source device via ECN marking.

Question: 65

Exhibit.

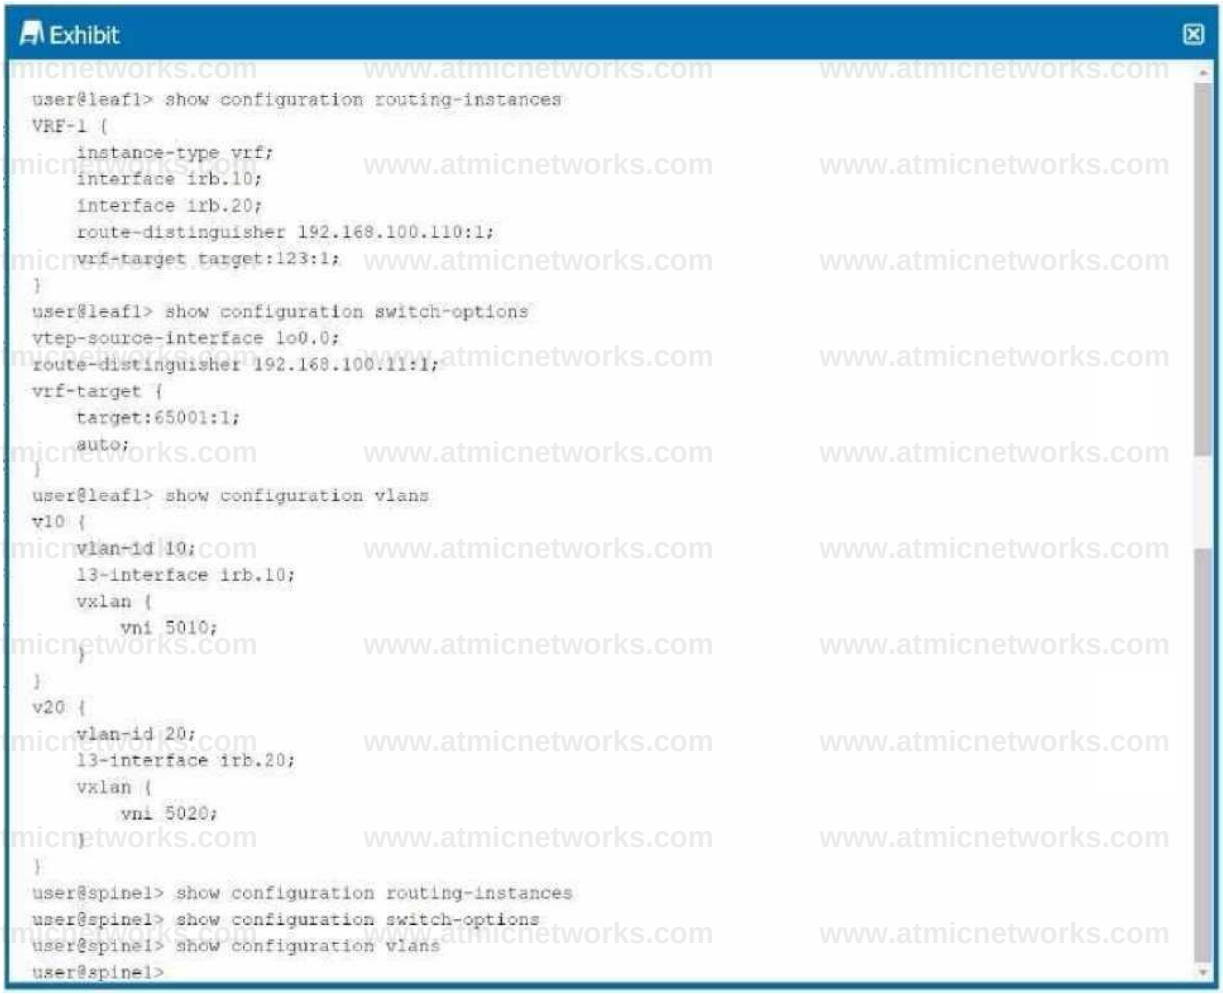

Referring to the exhibit, which statement Is true?

A. A PBB-EVPN architecture is being used.
B. An ERB architecture is being used.
C. An OTT architecture is being used.
D. A CRB architecture is being used.

Answer: B

Explanation:

Understanding Network Architectures:

ERB (Edge Routed Bridging) architecture involves routing at the network's edge (leaf nodes), while traffic between leaf nodes is switched. This is commonly used in VXLAN-EVPN setups.

Analysis of the Exhibit:

The exhibit shows configurations related to routing instances, VXLAN, and VLANs, with VNIs being used for each VLAN. This setup is characteristic of an ERB architecture where each leaf device handles Layer 3 routing for its connected devices.

Conclusion:

Option B: Correct—The configuration shown corresponds to an ERB architecture where routing occurs at the network's edge (leaf devices).