

EXAMGOOD

QUESTION & ANSWER

Exam Good provides update free of charge in one year!

Accurate study guides
High passing rate!

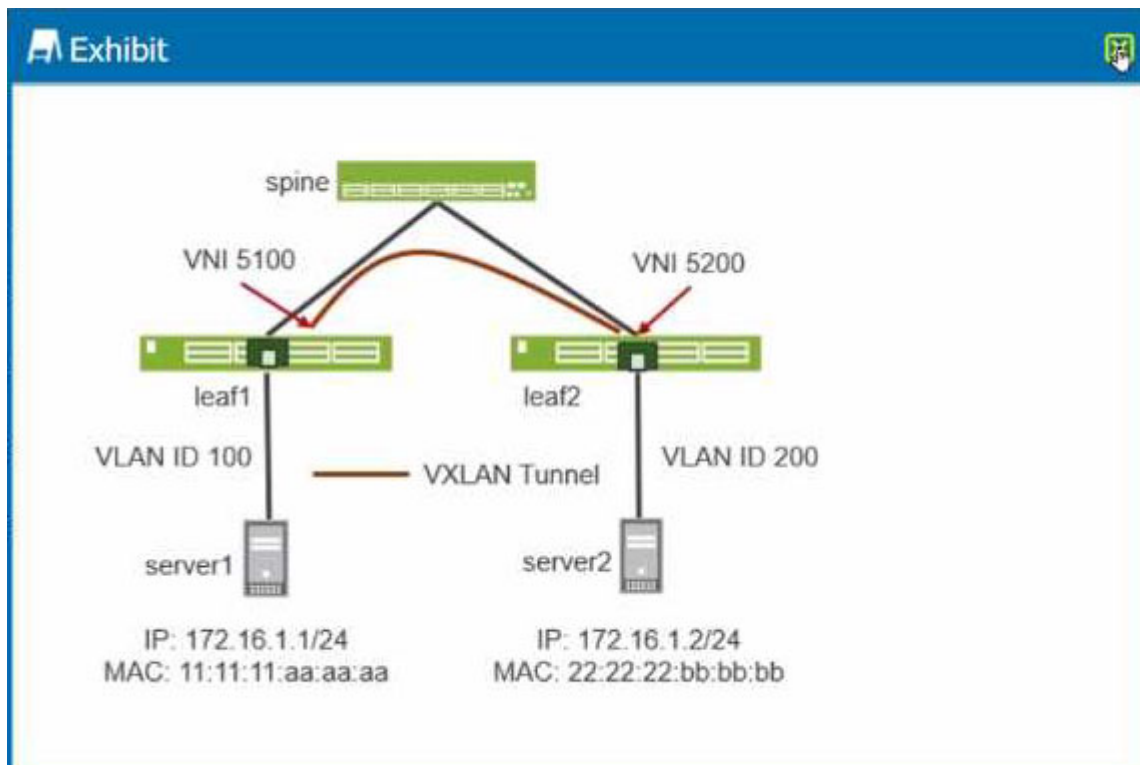
<http://www.examgood.com>

Exam : **JN0-683**

Title : Data Center Professional
(JNCIP-DC) Exam

Version : DEMO

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.

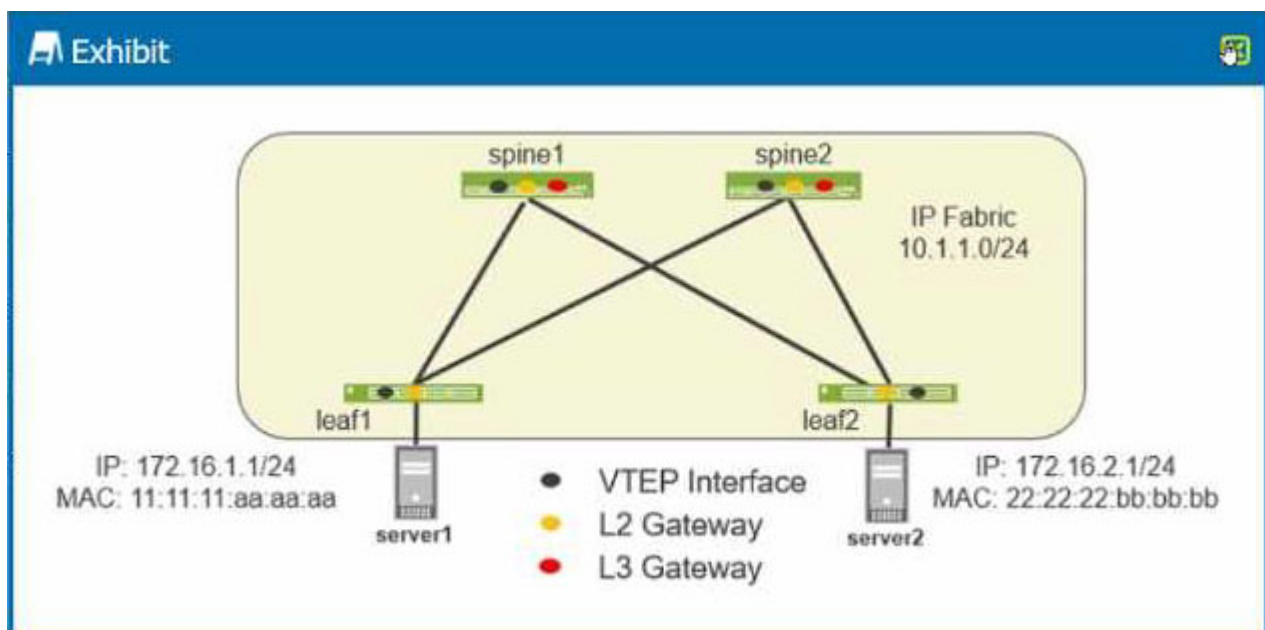
Data Center

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.

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 spine1 and spine2.
- B. Traffic from server1 to server2 will transit a VXLAN tunnel to spine1 or spine2. then a VXLAN tunnel from spine1 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 inter-subnet routing for traffic between Server1 and Server2.

D. Traffic from server1 to server2 will transit the VXLAN tunnel between leaf1 and leaf2: This describes the correct VXLAN operation where the traffic is encapsulated by leaf1 and decapsulated by leaf2.

Data Center

Reference: In EVPN-VXLAN architectures, the leaf switches often handle both Layer 2 switching and Layer 3 routing via IRB interfaces. This allows for efficient routing within the data center fabric without the need to involve the spine switches for every routing decision.

The described traffic flow aligns with standard EVPN-VXLAN designs, where direct VXLAN tunnels between leaf switches enable seamless and scalable communication across a data center network.

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.

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.

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.

5. You want to ensure that VXLAN traffic from the xe-0/0/12 interface is being encapsulated by logical vtep.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.

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.