

**Free Questions for JN0-683**

**Shared by Sanchez on 04-10-2024**

**For More Free Questions and Preparation Resources**

**Check the Links on Last Page**

# Question 1

---

**Question Type: MultipleChoice**

---

Exhibit.

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

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

You have a sample configuration for connecting two sites through EVPN-VXLAN by exchanging IP prefix routes.

Referring to the exhibit, which two statements regarding the configuration are true? {Choose two.}

### Options:

---

- A-** The advertise direct-next-hop option enables the receiver to resolve the next-hop route using only information carried in the Type 5 route.
- B-** The advertise direct-next-hop 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:

---

A, C

### 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-next-hop 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 2

---

**Question Type:** MultipleChoice

---

You are asked to interconnect two 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?

**Options:**

---

**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 3

---

**Question Type: MultipleChoice**

---

Exhibit.

```
user@switch> ping overlay tunnel-type vxlan vni 100 tunnel-src 192.168.2.10 tunnel-dst 192.168.2.20
mac 00:00:5E:00:53:cc count 1
ping-overlay protocol vxlan
  vni 100
    tunnel src ip 192.168.2.10
    tunnel dst ip 192.168.2.20
    mac address 00:00:5E:00:53:cc
    count 5
    ttl 255

WARNING: following hash-parameters are missing -
        hash computation may not succeed

        end-host smac
        end-host dmac
        end-host src ip
        end-host dst ip
        end-host protocol
        end-host 14-src-port
        end-host 14-dst-port
Request for seq 1, to 192.168.2.20, at 09-24 23:53:54 PDT.089 msec
Response for seq 1, from 192.168.2.20, at 09-24 23:53:54 PDT.089 msec, rtt 6 msec
Overlay-segment present at RVTEP 192.168.2.20
End-System Not Present
```

Referring to the exhibit, which statement is correct?



## Options:

---

- 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 4

---

**Question Type:** MultipleChoice

---

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

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

### Options:

---

- 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:

---

B, C

### **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.

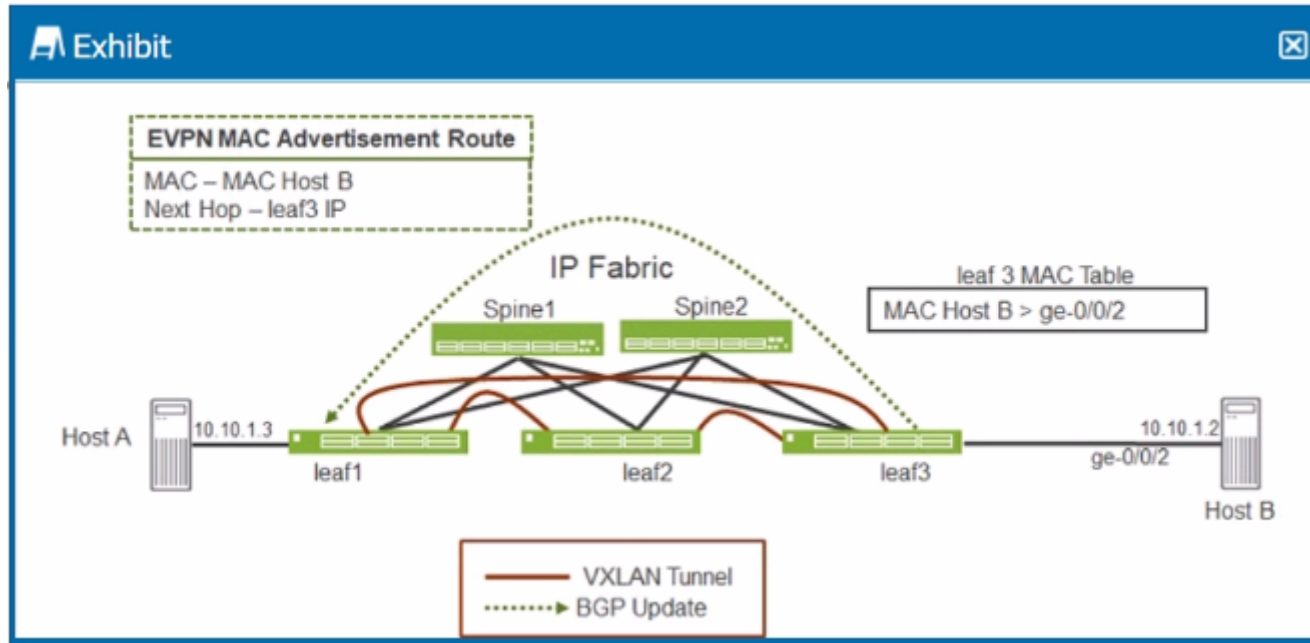
Option B: Correct---Adding spine devices decreases the oversubscription ratio.

## **Question 5**

---

**Question Type: MultipleChoice**

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?

**Options:**

**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 6

---

**Question Type:** MultipleChoice

---

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

```

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?

### Options:

---

- 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 10.1.2.1 is being forwarded into the VXLAN tunnel with the correct VTEP (VXLAN Tunnel Endpoint).

Conclusion:

Option B: Correct---The traffic from Host1 is entering the VXLAN tunnel, as evidenced by the next hop pointing to a VTEP. However, the issue could lie elsewhere, possibly with the remote VTEP, routing configurations, or the receiving leaf/spine devices.

## Question 7

---

**Question Type: MultipleChoice**

---

Exhibit.

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

The exhibit shows the truncated output of the show evpn database command.

Given this output, which two statements are correct about the host with MAC address 40:00:dc:01:00:04? (Choose two.)

## Options:

---

- 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:

---

A, D

## Explanation:

---

Understanding the Output:

The show evpn database command output shows the MAC address, VLAN, active source, timestamp, and IP address associated with various hosts in the EVPN instance.

Analysis of the MAC Address:

Option A: The MAC address 40:00:dc:01:00:04 is associated with the IP address 10.4.4.5, as indicated by the output in the IP address column. This confirms that this host has been assigned the IP 10.4.4.5.

Option D: The active source for the MAC address 40:00:dc:01:00:04 is listed as 00:02:00:00:00:04:00:04:00:00:04:00:04, which indicates that the host is connected via an ESI (Ethernet Segment Identifier) LAG (Link Aggregation Group). This setup is typically used

in multi-homing scenarios to provide redundancy and load balancing across multiple physical links.

Conclusion:

Option A: Correct---The host with MAC 40:00:dc:01:00:04 is assigned IP 10.4.4.5.

Option D: Correct---The host is originating from an ESI LAG, as indicated by the active source value.

## Question 8

---

**Question Type: MultipleChoice**

---

Exhibit.

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

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.)

### Options:

---

- 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:

---

B, D

### 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 9

---

**Question Type:** MultipleChoice

---

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?

**Options:**

---

**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.

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 10

---

**Question Type:** MultipleChoice

---

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

### Options:

---

- 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:

---

B, D

### 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 11

---

**Question Type:** MultipleChoice

---

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

**Options:**

---

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.

**To Get Premium Files for JN0-683 Visit**

**<https://www.p2pexams.com/products/jn0-683>**

**For More Free Questions Visit**

**<https://www.p2pexams.com/juniper/pdf/jn0-683>**

