# Juniper JN0-683 Practice Questions

**Shared by Cooley on 19-10-2026**

**For More Free Questions and Preparation Resources**

Check the Links on Last Page

## Question 1

Question Type: MultipleChoice

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

### Options:

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:

B, C

### 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 2

Question Type: MultipleChoice

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?

### 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.

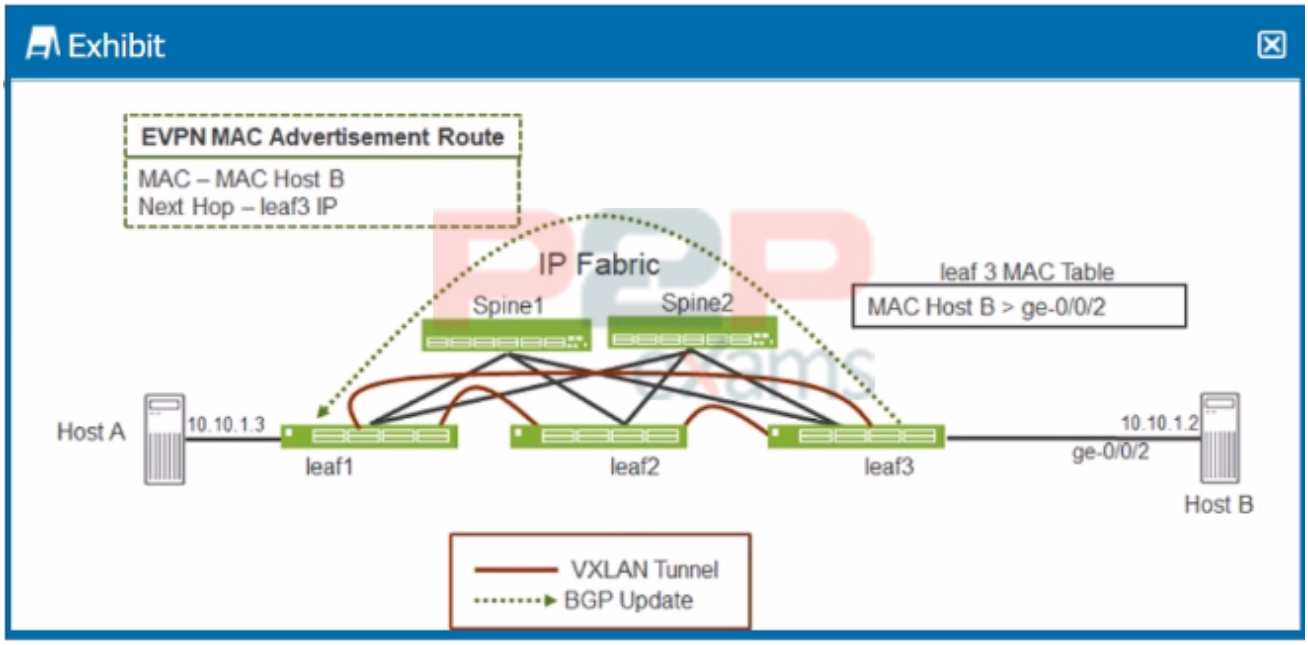
Conclusion:

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.




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.

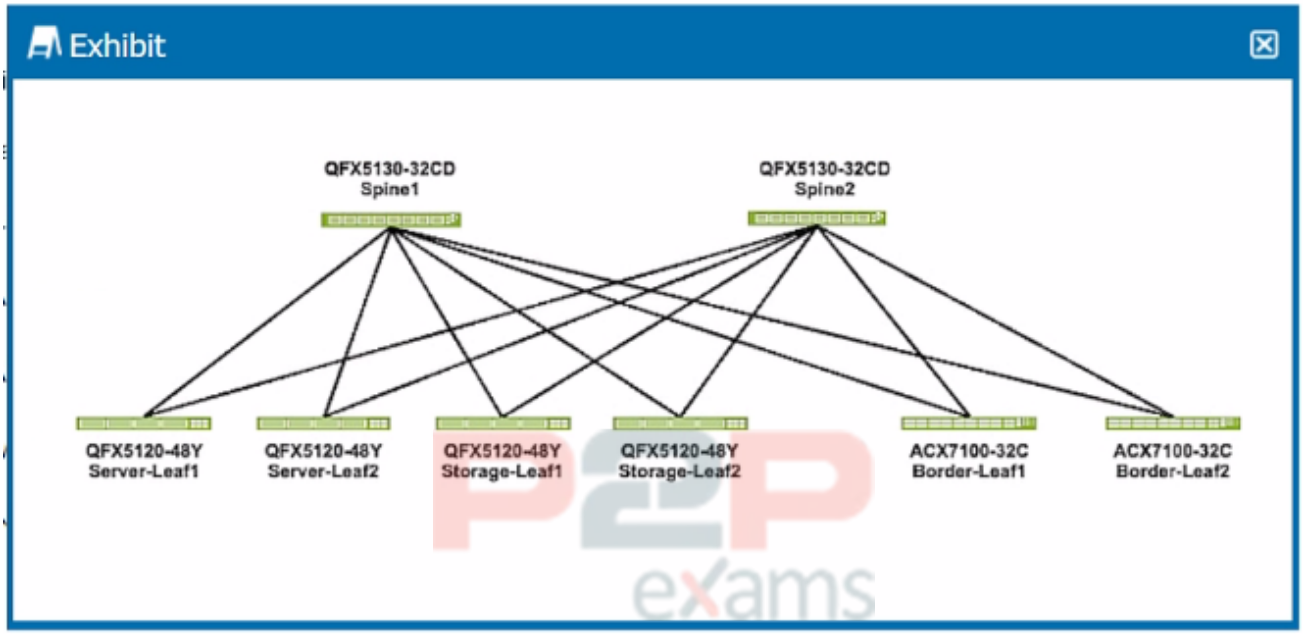
Conclusion:

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 4

Question Type: MultipleChoice

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.)

### Options:

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:

B, C, E

## 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 5

Question Type: MultipleChoice

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

## Options:

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:

B, C

## Explanation:

Multicast-Signaled VXLAN:

In traditional multicast-signaled VXLAN, VTEPs (VXLAN Tunnel Endpoints) use multicast to flood and learn about remote VTEPs. This method relies on multicast in the underlay network to distribute BUM (Broadcast, Unknown unicast, and Multicast) traffic.

This approach can be resource-intensive due to the need for multicast group management and increased network traffic, especially in large deployments.

EVPN-Signaled VXLAN:

EVPN-signaled VXLAN uses BGP (Border Gateway Protocol) to signal the presence of VTEPs and distribute MAC address information. BGP is used for VTEP autodiscovery and the distribution of endpoint information.

This method is more efficient because it reduces the reliance on multicast, instead using BGP control-plane signaling to handle VTEP discovery and MAC learning, which reduces the overhead on the network and improves scalability.

Correct Statements:

B . An EVPN-signaled VXLAN can perform autodiscovery of VTEPs using BGP: This is correct because EVPN uses BGP for VTEP autodiscovery, making it more efficient and scalable compared to multicast-based methods.

C . An EVPN-signaled VXLAN is less resource-intensive: This is correct because it eliminates the need for multicast flooding in the underlay, instead using BGP for signaling, which is less demanding on network resources.

Incorrect Statements:

A . An EVPN-signaled VXLAN can perform autodiscovery of VTEPs using IS-IS: This is incorrect because EVPN relies on BGP, not IS-IS, for VTEP discovery and signaling.

D . An EVPN-signaled VXLAN features slower and more complete convergence: This is incorrect; EVPN with BGP typically provides faster convergence due to its use of a control plane rather than relying on data plane learning.

Data Center Reference:

EVPN-VXLAN is widely adopted in modern data center designs due to its scalability, efficiency, and reduced resource consumption compared to multicast-based VXLAN solutions. It leverages the strengths of BGP for control-plane-driven operations, resulting in more efficient and scalable networks.

## Question 6

Question Type: MultipleChoice

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.)

### Options:

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:

B, C

### 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 7

Question Type: MultipleChoice

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

### Options:

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:

A, D

### 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:

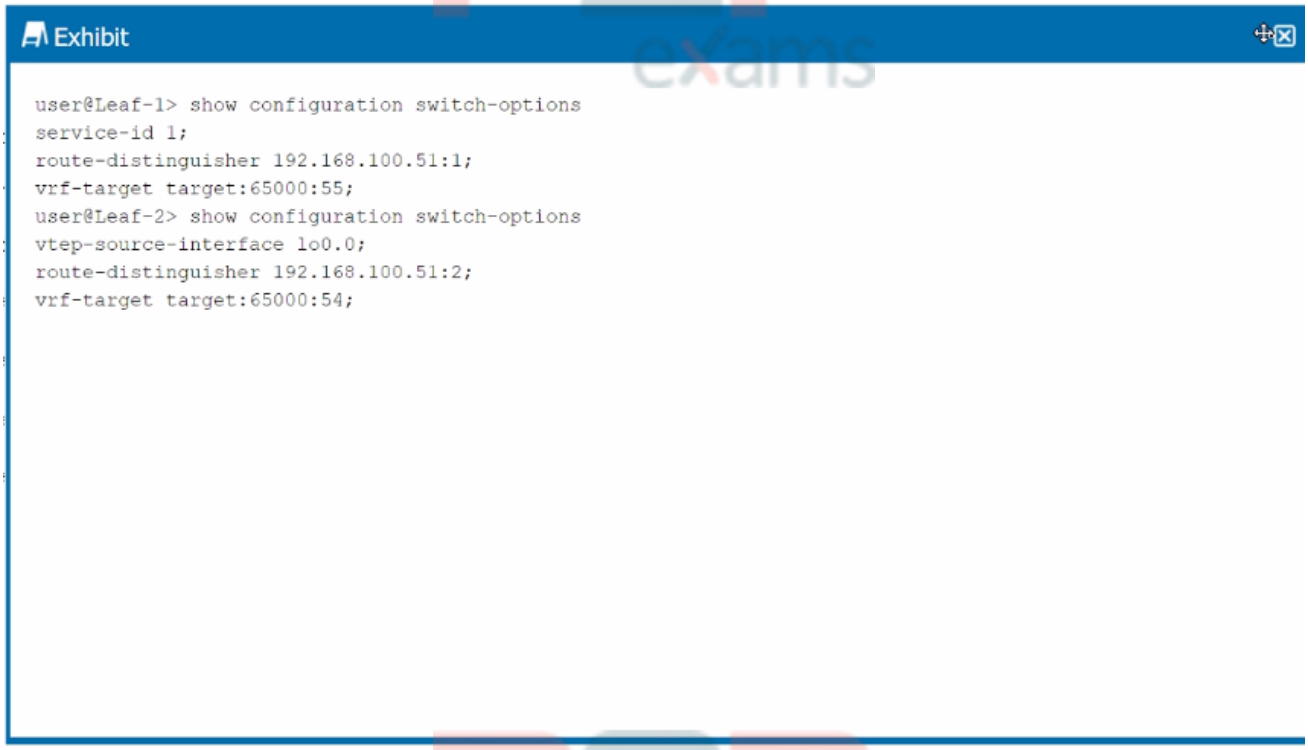
Option A: Correct---Devices in an IP fabric function as Layer 3 routers.

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

## Question 8

Question Type: MultipleChoice

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.)

### Options:

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:

B, D

## 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 9

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 10

Question Type: MultipleChoice

A local VTEP has two ECMP paths to a remote VTEP

Which two statements are correct when load balancing is enabled in this scenario? (Select two.)

## Options:

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:

C, D

## 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 11

Question Type: MultipleChoice

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?

## Options:

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 12

Question Type: MultipleChoice

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

## Options:

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:

B, C

## 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 13

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

P2P exams

## 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.

P2P exams

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

