Free Questions for JN0-683 by go4braindumps

Shared by Mendez on 06-09-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Exhibit.

Exhibit

```
user@leaf1> show configuration routing-instances
VRF-1 {
    instance-type vrf;
    interface irb.10;
    interface irb.20;
    route-distinguisher 192.168.100.110:1;
    vrf-target target:123:1;
}
user@leaf1> show configuration switch-options
vtep-source-interface lo0.0;
route-distinguisher 192.168.100.11:1;
vrf-target {
    target:65001:1;
    auto;
}
user@leaf1> show configuration vlans
v10 {
    vlan-id 10;
    l3-interface irb.10;
    vxlan {
        vni 5010;
    }
}
v20 {
    vlan-id 20;
    l3-interface irb.20;
    vxlan {
        vni 5020;
    }
}
user@spine1> show configuration routing-instances
user@spine1> show configuration switch-options
user@spine1> show configuration vlans
user@spine1>
```

Referring to the exhibit, which statement Is true?

Options:

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).

Question 2

Question Type: MultipleChoice

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

Options:

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 3

Question Type: MultipleChoice

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

Options:

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:

B, D, E

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 4

Question Type: **MultipleChoice**

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?

Options:

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 5

Question Type: MultipleChoice

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.)

Options:

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:

B, D

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 6

Question Type: **MultipleChoice**

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?

Options:

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.