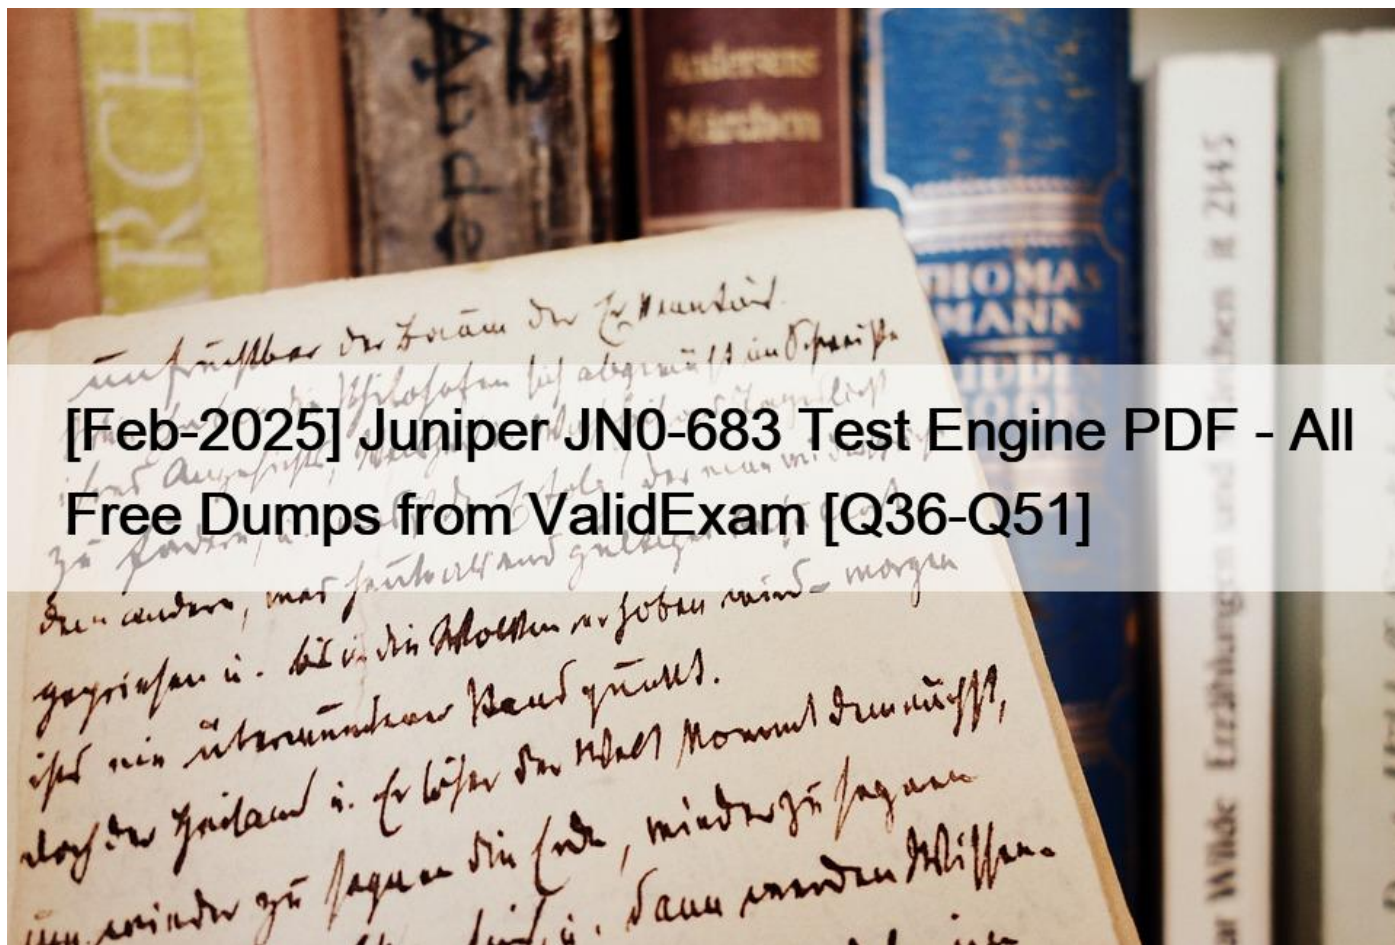


[Feb-2025 Juniper JN0-683 Test Engine PDF - All Free Dumps from ValidExam [Q36-Q51]



[Feb-2025] Juniper JN0-683 Test Engine PDF - All Free Dumps from ValidExam
Get New JN0-683 Certification & Valid Exam Dumps Questions

Juniper JN0-683 Exam Syllabus Topics:

TopicDetailsTopic 1- EVPN-VXLAN Signaling: This section assesses an understanding of Ethernet VPN (EVPN) concepts, including route types, multicast handling, and Multiprotocol BGP (MBGP). It also covers EVPN architectures like CRB and ERB, MAC learning, and symmetric routing.Topic 2- Data Center Interconnect: For Data Center Engineers, this part focuses on interconnecting data centers, covering Layer 2 and Layer 3 stretching, stitching fabrics together, and using EVPN-signaled VXLAN for seamless communication between data centers.Topic 3- Layer 3 Fabrics: This section measures the knowledge of professionals managing IP-based networks in data centers. It covers IP fabric architecture and routing, ensuring candidates understand how the network is structured for scalability and how traffic is routed efficiently.

NEW QUESTION 36

You want to ensure thatVXLAN traffic from the xe-0/0/12 interlace is being encapsulatedby logical vlep.

32770 and sent to a remote leaf device in this scenario, which command would you use to verify that traffic is flowing?

- * monitor traffic interface xe-0/0/12
- * show interface terse vtep.32770
- * show interfaces terse vtep.32770 statistics
- * show interfaces vtep.32770 detail
- * 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 References:

* Monitoring interface statistics is a crucial step in troubleshooting and validating network traffic, particularly in complex overlay environments like EVPN-VXLAN.

NEW QUESTION 37

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.)

- * Design a five-stage Clos IP fabric.
- * Design a three-stage Clos IP fabric.
- * Use EX4300 Series devices as the spine devices.
- * Use OFX5700 Series devices as the super spines.
- * 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:

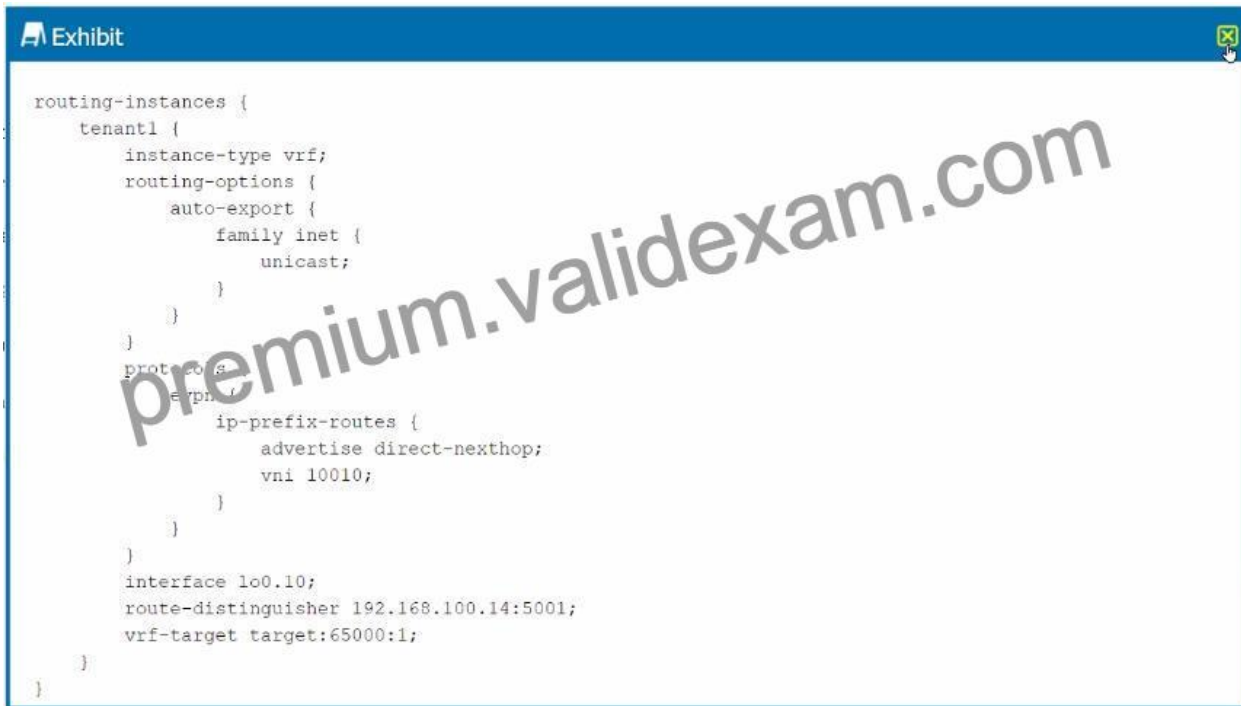
* Option D: Using high-capacity devices like the QFX5700 Series as super spines can handle the increased traffic demands in large data centers and support future growth. These devices provide the necessary bandwidth and scalability for large-scale deployments.

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.

NEW QUESTION 38

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?

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

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.

NEW QUESTION 39

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?

* gNMI

* SNMP

* Native Sensors

* OpenConfig

* 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.

NEW QUESTION 40

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.)

- * Configure a Type 2 EVPN route for each unique prefix.
- * Configure peering for EVPN between border leaf nodes in each data center.
- * Configure a Type 5 EVPN route for each unique prefix.
- * Configure peering for EVPN between all leaf nodes within each data center.
- * 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.

NEW QUESTION 41

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

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

NEW QUESTION 42

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.)

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

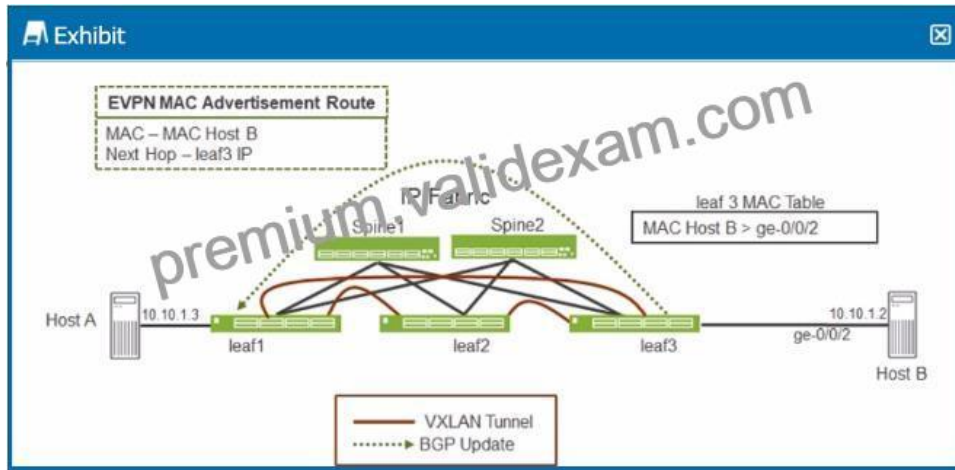
- * A. The inner packet fields are not used in the hash for load balancing: This is incorrect as the inner packet fields are indeed critical for generating the hash used in load balancing.
- * B. The destination port in the UDP header is used to load balance VXLAN traffic: This is incorrect because the 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 References:

- * 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.

NEW QUESTION 43

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?

- * proxy ARP
- * proxy NDP
- * GARP
- * DAD

* 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.

NEW QUESTION 44

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

- * Type 5 EVPN routes require a VXLAN tunnel to the protocol next hop.
- * Type 2 EVPN routes do not require a VXLAN tunnel to the protocol next hop.
- * Type 2 EVPN routes require a VXLAN tunnel to the protocol next hop.

* Type 5 EVPN routes do not require a VXLAN tunnel to the protocol next hop.

* 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.

NEW QUESTION 45

You are implementing VXLAN broadcast domains in your data center environment. Which two statements are correct in this scenario? (Choose two.)

* A VXLAN packet does not contain a VLAN ID.

* The VNI must match the VLAN tag to ensure that the remote VTEP can decapsulate VXLAN packets.

* Layer 2 frames are encapsulated by the source VTEP.

* The VNI is a 16-bit value and can range from 0 through 16.777.215.

* 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.

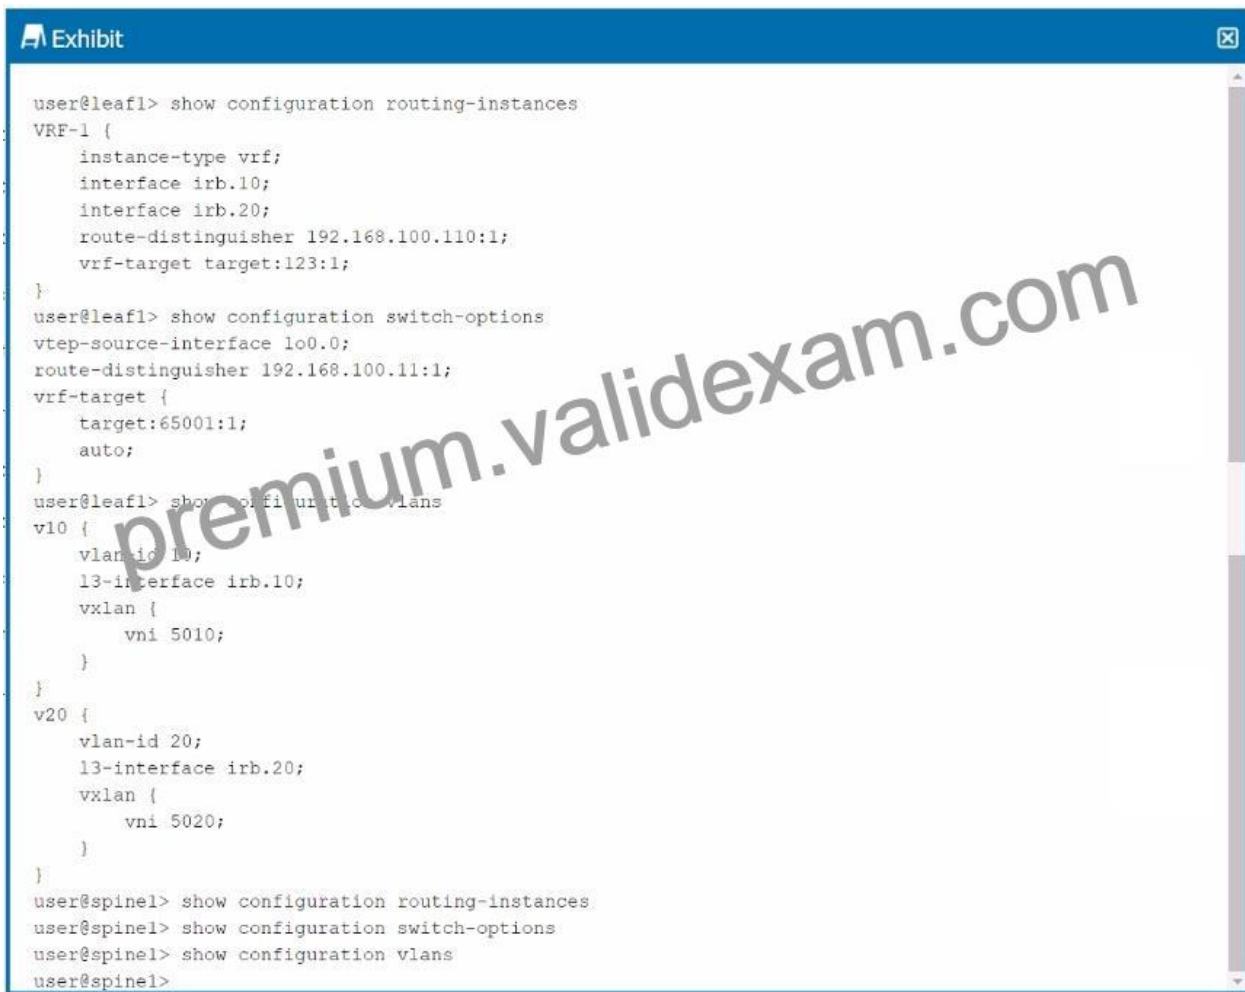
* D. The VNI is a 16-bit value and can range from 0 through 16,777,215: This is incorrect because the VNI is a 24-bit value, allowing for a range of 0 to 16,777,215.

Data Center References:

* VXLAN technology is critical for modern data centers as it enables scalability and efficient segmentation without the constraints of traditional VLAN limits.

NEW QUESTION 46

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;
  13-interface irb.10;
  vxlan {
    vni 5010;
  }
}
v20 {
  vlan-id 20;
  13-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?

- * A PBB-EVPN architecture is being used.
- * An ERB architecture is being used.
- * An OTT architecture is being used.
- * A CRB architecture is being used.
- * 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).

NEW QUESTION 47

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.)

- * the IP address of the FTP server
- * the MAC address for each OFX5120
- * the MAC address of the FTP server
- * the management IP address for each OFX5120
- * 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.

NEW QUESTION 48

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?

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

NEW QUESTION 49

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
-----
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 vxn nat-bits
Instance: default-switch
VLAN  Default-ID  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?

- * The traffic is failing because load balancing is not configured correctly.
- * The traffic is entering the VXLAN tunnel.
- * Host1 and Host2 are directly connected to leaf1.
- * The irb.20 interface is not configured on leaf1.

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.

NEW QUESTION 50

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

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

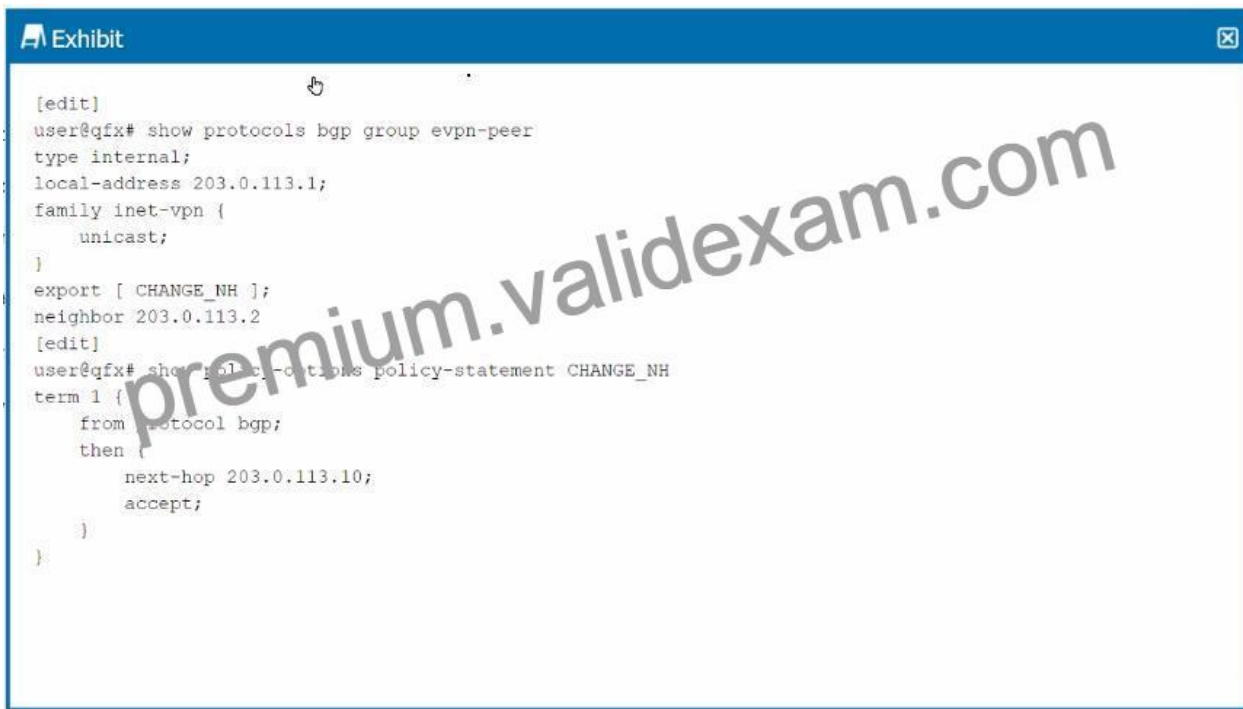
Conclusion:

* 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.

NEW QUESTION 51

Exhibit.



```
[edit]
user@qfx# show protocols bgp group evpn-peer
type internal;
local-address 203.0.113.1;
family inet-vpn {
    unicast;
}
export [ CHANGE_NH ];
neighbor 203.0.113.2
[edit]
user@qfx# show policy-options policy-statement CHANGE_NH
term 1 {
    from protocol bgp;
    then {
        next-hop 203.0.113.10;
        accept;
    }
}
```

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?

- * EVPN routes cannot have the next hop changed.
- * The export policy is incorrectly configured.
- * The vrf-export parameter must be applied.
- * The vpn-apply-export parameter must be applied to this peer.
- * 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 References:

* 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.

100% Passing Guarantee - Brilliant JN0-683 Exam Questions PDF: <https://www.validexam.com/JN0-683-latest-dumps.html>