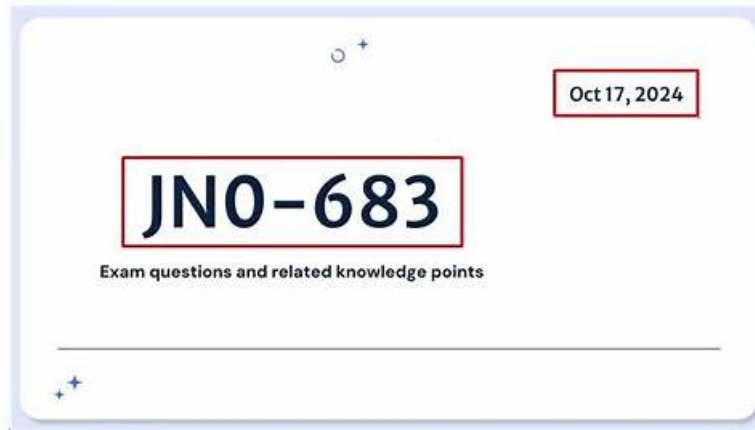


JN0-683 Reliable Test Topics - Valid Test JN0-683 Braindumps



P.S. Free & New JN0-683 dumps are available on Google Drive shared by VCE4Dumps: https://drive.google.com/open?id=1esB9iGJoWgh3qIeV5_XyoTiMdm9as8KF

Some people are not good at operating computers. So you might worry about that the JN0-683 certification materials are not suitable for you. Try to believe us. Our experts have taken your worries seriously. They have made it easy to operate for all people. Even if you know little about computers, you can easily begin to do exercises of the JN0-683 real exam dumps. Also, we have invited for many volunteers to try our study materials. The results show our products are suitable for them. In addition, the system of our JN0-683 test training is powerful. You will never come across system crashes. The system we design has strong compatibility. High speed running completely has no problem at all.

Juniper JN0-683 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none">• 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 2	<ul style="list-style-type: none">• Data Center Deployment and Management: This section assesses the expertise of data center networking professionals like architects and engineers, focusing on key deployment concepts. Topics include Zero-touch provisioning (ZTP), which automates device setup in data centers without manual input.
Topic 3	<ul style="list-style-type: none">• Data Center Multitenancy and Security: This section tests knowledge of single-tenant and multitenant data center setups. Candidates such as Data Center Professionals are evaluated on ensuring tenant traffic isolation at both Layer 2 and Layer 3 levels in shared infrastructure environments.

>> JN0-683 Reliable Test Topics <<

Valid Test JN0-683 Braindumps & JN0-683 Reliable Dumps Sheet

When we are in some kind of learning web site, often feel dazzling, because web page design is not reasonable, put too much information all rush, it will appear desultory. Believe it or not, we face the more intense society, and we should prompt our competitiveness and get a JN0-683 certification to make our dreams come true. Although it is not an easy thing to achieve it, once you choose our JN0-683 prepare torrent, we will send the new updates for one year long, which is new enough to deal with the exam for you and guide you through difficulties in your exam preparation.

Juniper Data Center, Professional (JNCIP-DC) Sample Questions (Q20-Q25):

NEW QUESTION # 20

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. distributed symmetric routing
- C. bridged overlay
- **D. CRB**

Answer: D

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.

* Explanation:

* 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.

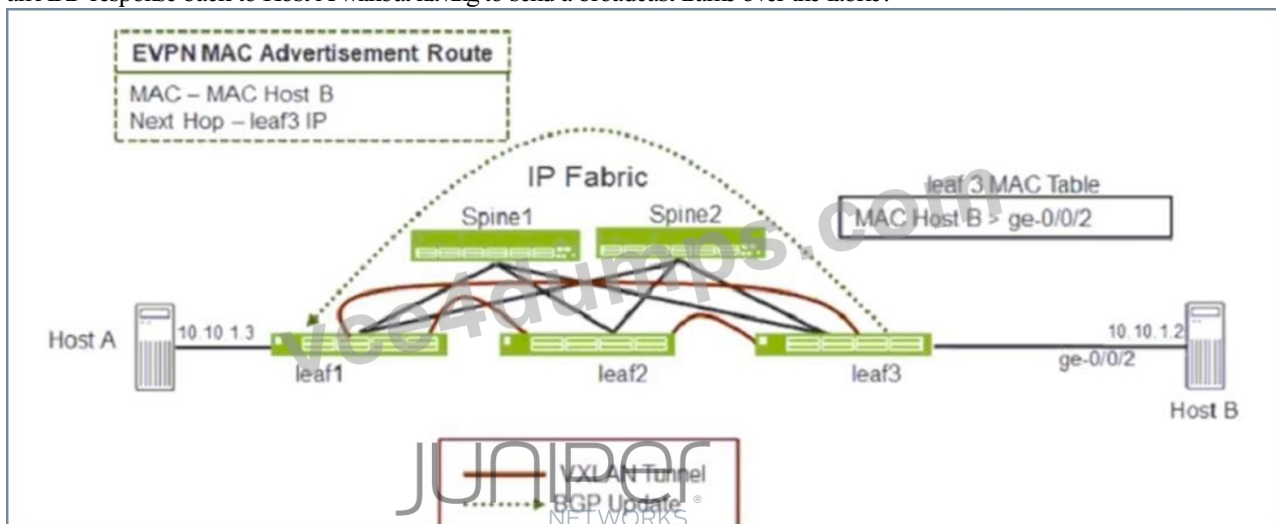
* This model is particularly advantageous in scenarios where centralized management and routing control are desired, reducing the complexity and configuration burden on the leaf switches.

Data Center References:

* 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.

NEW QUESTION # 21

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?



- A. GARP
- B. DAD
- **C. proxy ARP**
- D. proxy NDP

Answer: C

Explanation:

proxy ARP: In this scenario, when Host A sends an ARP request for Host B's IP address, leaf1 can use proxy ARP to respond to the ARP request without needing to forward it as a broadcast over the fabric. Proxy ARP allows leaf1 to reply with Host B's MAC address even though Host A and Host B are not directly connected on the same network segment. The leaf switch acts as the proxy for the destination IP.

NEW QUESTION # 22

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
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
-----
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?

- A. Host1 and Host2 are directly connected to leaf1.
- **B. The traffic is entering the VXLAN tunnel.**
- C. The traffic is failing because load balancing is not configured correctly.
- 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.

NEW QUESTION # 23

Click the Exhibit button. Connections between hosts connected to Leaf-1 and Leaf-2 are not working correctly.

```
user@Leaf-1> show configuration switch-options
service-id 1;
vtep-source-interface lo0.0;
route-distinguisher 192.168.100.51:1;
vrf-target target:65000:1;
user@Leaf-2> show configuration switch-options
vtep-source-interface lo0.0;
route-distinguisher 192.168.100.51:1;
vrf-target target:65000:2;
```

Referring to the exhibit, which two configuration changes are required to solve the problem?
(Choose two.)

- A. Configure the set switch-options vrf-target target:65000:1 parameter on Leaf-2.
- B. Configure the set switch-options route-distinguisher 192.168.100.51:1 parameter on Leaf-1.
- C. Configure the set switch-options service-id 1 parameter on Leaf-2.
- D. Configure the set switch-options vtep-source-interface irb.0 parameter on Leaf-1.

Answer: A,C

Explanation:

Configure the set switch-options vrf-target target:65000:1 parameter on Leaf-2: The vrf-target parameter on Leaf-2 must match the vrf-target on Leaf-1 to ensure that both leaves use the same routing information for their respective VRFs. In the configuration, Leaf-1 has vrf-target target:65000:1, while Leaf-2 has vrf-target target:65000:2. These must be consistent to allow proper communication and routing between the leaves.

Configure the set switch-options service-id 1 parameter on Leaf-2: The service-id configuration should be consistent across all leaf nodes to ensure that they are part of the same VXLAN service. Leaf-1 is configured with service-id 1, so Leaf-2 should be configured with the same service-id 1 to ensure consistency in the VXLAN deployment.

NEW QUESTION # 24

What are two supported methods for exporting data when using the Junos telemetry interface?
(Choose two.)

- A. using REST
- B. using SNMP
- C. using gRPC
- D. using UDP

Answer: C,D

Explanation:

The Junos Telemetry Interface (JTI) supports streaming real-time data from devices using modern, efficient protocols: using UDP

JTI supports UDP-based streaming, often with Juniper's JTIs streaming protocol (for low-latency telemetry).

using gRPC

JTI supports gRPC/gNMI for structured, reliable telemetry export.

NEW QUESTION # 25

