JN0-683 Discount & Reliable JN0-683 Braindumps Files

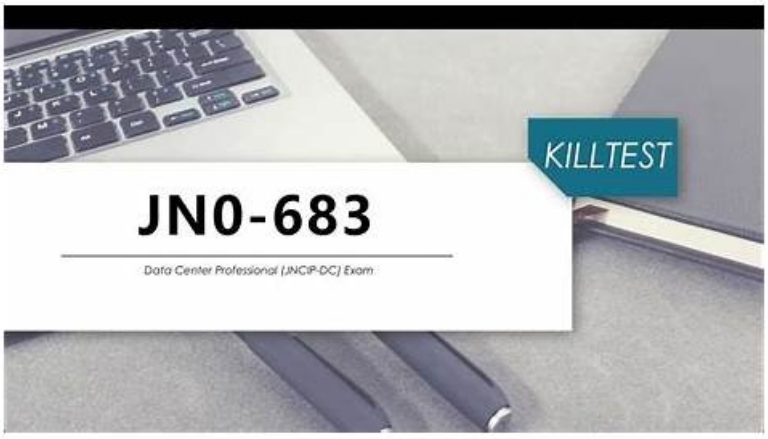

BTW, DOWNLOAD part of TrainingQuiz JN0-683 dumps from Cloud Storage: https://drive.google.com/open?id=1qRXs31qkSVp0ET74Pu9IW8LABPSJhXtO

The language in our JN0-683 test guide is easy to understand that will make any learner without any learning disabilities, whether you are a student or a in-service staff, whether you are a novice or an experienced staff who has abundant experience for many years. Our JN0-683 Exam Questions are applicable for everyone in all walks of life which is not depends on your educated level. Therefore, it should be a great wonderful idea to choose our JN0-683 guide torrent for sailing through the difficult test and pass it.

Juniper JN0-683 Exam Syllabus Topics:

Topic	Details
Topic 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 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	• 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.
Topic 4	• 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 Discount <<

Reliable JN0-683 Braindumps Files, Latest JN0-683 Exam Pdf

Our Juniper JN0-683 free demo provides you with the free renewal in one year so that you can keep track of the latest points happening in the world. As the questions of our Juniper JN0-683 Exam Dumps are involved with heated issues and customers who prepare for the Juniper JN0-683 exams must haven't enough time to keep trace of JN0-683 exams all day long.

Juniper Data Center, Professional (JNCIP-DC) Sample Questions (Q44-Q49):

NEW QUESTION # 44
Exhibit.

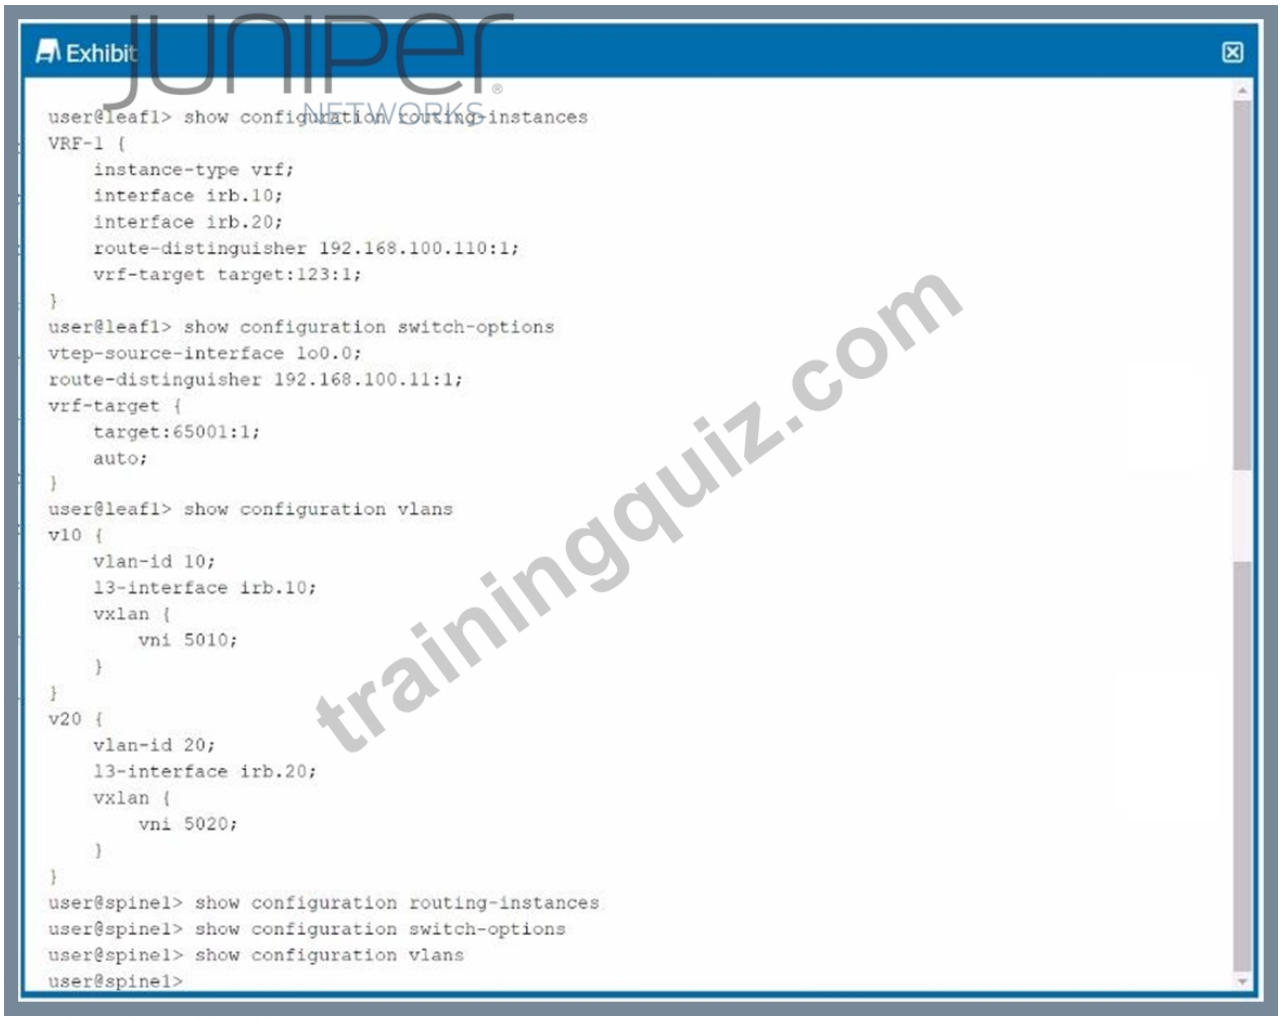

Referring to the exhibit, which statement Is true?

- A. A CRB architecture is being used.
- B. An ERB architecture is being used.
- C. An OTT architecture is being used.
- D. A PBB-EVPN 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).

NEW QUESTION # 45
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 leafdevices. What will satisfy these requirements?

- A. IP VPN
- B. EVPN Type 2 stretch
- C. over the top full-mesh interconnect
- 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 stitchingis 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.

NEW QUESTION # 46
You manage an IP fabric with an EVPN-VXLAN overlay. You have multiple tenants separated using multiple unique VRF instances. You want to determine the routing information that belongs in each routing instance's routing table.
In this scenario, which property is used for this purpose?

- A. the VRF table label
- B. the route distinguisher value
- C. the routing instance type
- D. the VRF target community

Answer: B

Explanation:
* Understanding VRF and Routing Instances:
* In an EVPN-VXLAN overlay network, multiple tenants are separated using unique VRF (Virtual Routing and Forwarding) instances. Each VRF instance maintains its own routing table, allowing for isolated routing domains within the same network infrastructure.
* Role of Route Distinguisher:
* Route Distinguisher (RD):The RD is a unique identifier used in MPLS and EVPN environments to distinguish routes belonging to different VRFs. The RD is prepended to the IP address in the route advertisement, ensuring that routes from different tenants remain unique even if they use the same IP address range.
* Correct Property:
* D. the route distinguisher value:This is the correct answer because the RD is crucial in determining which routing information belongs to which VRF instance. It ensures that each VRF' s routing table only contains relevant routes, maintaining isolation between tenants.
Data Center References:
* The RD is a key element in MPLS and EVPN-based multi-tenant environments, ensuring proper routing segregation and isolation for different VRFs within the data center fabric.

NEW QUESTION # 47
Exhibit.

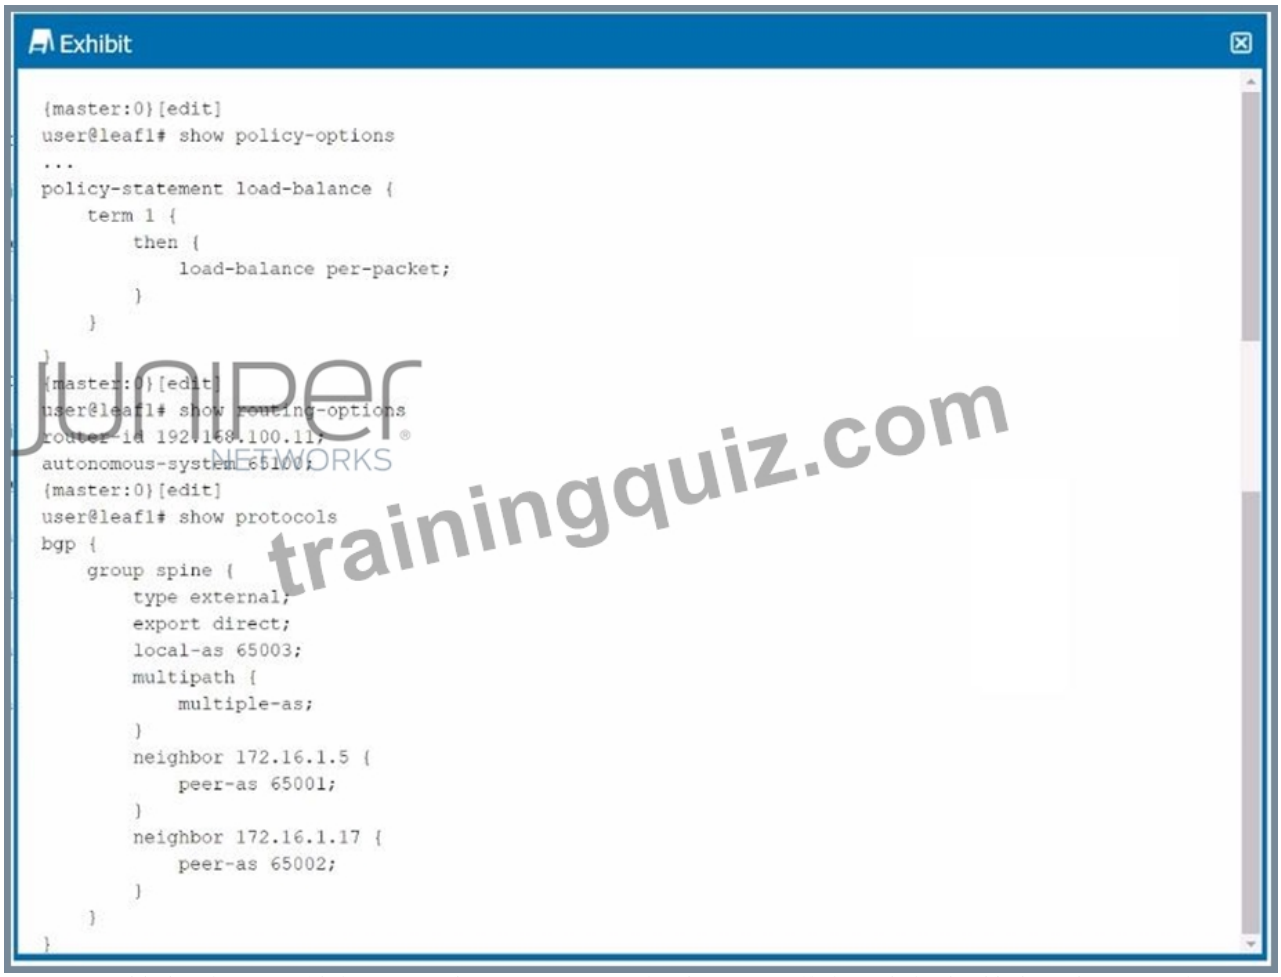

You are troubleshooting an IP fabric (or your data center. You notice that your traffic is not being load balanced to your spine devices from your leaf devices. Referring to the configuration shown in the exhibit, what must be configured to solve this issue?

- A. The load-balance policy must have a from statement that matches on protocol bgp.
- B. The load-balance policy must be applied to the forwarding table under the routing-options hierarchy.
- C. The multipast multiple -as configuration must be configured for each peer in the BGP spine group.
- D. The load-balance policy must be applied as an export policy to your BGP

Answer: B

Explanation:
Step 1: Understand the Configuration in the Exhibit
The exhibit provides three configuration snippets from a leaf device (user@leaf#):
* Policy Options:
user@leaf# show policy-options
policy-statement load-balance {
term 1 {
then {
load-balance per-packet;
}
}
}
* A policy named load-balance is defined, which applies the load-balance per-packet action. In Juniper terminology, per-packet actually means per-flow load balancing (a common point of confusion). This policy is intended to enable load balancing across multiple paths.
* Routing Options:
user@leaf# show routing-options
router-id 192.168.100.11;
autonomous-system 65100;
* The router ID is set to 192.168.100.11, and the autonomous system (AS) number is 65100. There' s no mention of applying the load-balance policy here, which is a clue to the issue.

* BGP Configuration:
user@leaf# 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;
}
}
}
* BGP is configured with an external group spine, where the leaf device (local AS 65003) peers with spine devices (AS 65001 and 65002).
* The multipath multiple-as statement is enabled, which allows BGP to install multiple paths for the same prefix in the routing table, even if the paths come from different AS numbers. This is a prerequisite for load balancing in a multi-AS environment like an IP fabric.
* The export direct policy is applied, which likely exports directly connected routes to the spine devices.
Step 2: Identify the Problem
The issue is that traffic from the leaf to the spine devices is not being load-balanced, despite the presence of a load-balance policy and BGP multipath. For load balancing to work in this scenario:
* BGP multipath ensures multiple paths are installed in the routing table.
* The load-balance per-packet policy is meant to distribute traffic across those paths.
* However, the load-balance policy is defined but not applied anywhere in the configuration shown. For load balancing to take effect, the policy must be applied in the correct context.
Step 3: Evaluate the Options
Let's go through each option to determine the correct solution:
* A. The load-balance policy must be applied to the forwarding table under the routing-options hierarchy.
* In Junos, to enable load balancing across multiple paths for forwarding, the load-balance policy must be applied at the forwarding table level. This is done under the routing-options hierarchy using the forwarding-table export statement. For example:
set routing-options forwarding-table export load-balance
* This ensures that the load-balancing policy is applied to the forwarding table, allowing traffic to be distributed across multiple equal-cost paths installed by BGP.
* B. The multipath multiple-as configuration must be configured for each peer in the BGP spine group.
* The multipath multiple-as statement is already configured under the spine group, and it applies to all neighbors in that group (172.16.1.5 and 172.16.1.17). There's no need to configure it per peer, as the group-level configuration is sufficient. This option is incorrect because the required setting is already in place.
* C. The load-balance policy must be applied as an export policy to your BGP.
* Applying the load-balance policy as a BGP export policy (e.g., export load-balance under the BGP group) would affect the routes advertised to the spine devices. However, the load-balance per-packet action is a forwarding action, not a route advertisement action. Applying it as a BGP export policy would not achieve the desired load balancing for traffic forwarding and is incorrect.
* D. The load-balance policy must have a from statement that matches on protocol bgp.
* The load-balance policy currently applies the load-balance per-packet action unconditionally (no from statement). Adding a from protocol bgp condition would make the policy apply only to BGP routes, but this is unnecessary in this context. The policy needs to be applied to the forwarding table to affect traffic, not modified with a from statement. This option doesn't address the core issue of applying the policy.
Step 4: Determine the Correct Answer
The key issue is that the load-balance policy is defined but not applied. For load balancing to work, it must be applied to the forwarding table under routing-options. This matchesOption A:
* A. The load-balance policy must be applied to the forwarding table under the routing-options hierarchy.
Step 5: Provide Official Juniper Documentation Reference
Since I don't have direct access to Juniper's proprietary documents, I can provide an explanation based on standard Junos documentation practices and publicly available resources, such as the Juniper TechLibrary, which is the official source for Junos configuration guides.
In Juniper's official documentation, specifically in theJunos OS Routing Protocols and Policies Configuration Guide, the process for

enabling load balancing is described as follows:
* Load Balancing in Junos: To enable per-flow load balancing across multiple paths, you must define a policy with the load-balance per-packet action and apply it to the forwarding table. The relevant configuration hierarchy is:
routing-options {
forwarding-table {
export <policy-name>;
}
}
* Explanation from Documentation: The load-balance per-packet action (which performs per-flow balancing) requires the policy to be applied at the forwarding-table level to influence how traffic is distributed across multiple paths in the forwarding table. Without this, even if BGP installs multiple paths (via multipath), the forwarding engine will not load-balance traffic.
This aligns with the JNCIP-DC exam objectives, which include understanding how to configure and troubleshoot load balancing in an IP fabric, such as applying policies for traffic distribution.

NEW QUESTION # 48
Exhibit.

Exhibit

```
user@Border-Leaf-1> show configuration protocols bgp
group UNDERLAY {
    type external;
    export LOOPBACKS;
    local-as 65205;
    multipath {
        multiple-as;
    }
    neighbor 172.16.1.5 {
        peer-as 65102;
    }
}
group OVERLAY {
    type external;
    local-address 192.168.100.4;
    family evpn {
        signaling;
    }
    local-as 65101;
    neighbor 192.168.100.1 {
        peer-as 65102;
    }
    neighbor 192.168.100.22 {
        description Border-Leaf-2;
        peer-as 65222;
    }
    accept-remote-nexthop;
}
group PROVIDER {
    type external;
    peer-as 65001;
    local-as 65002;
    neighbor 172.16.1.224;
}
```


You are troubleshooting a DCI connection to another data center The BGP session to the provider is established, but the session to Border-Leaf-2 is not established. Referring to the exhibit, which configuration change should be made to solve the problem?

- A. set protocols bgp group overlay export loopbacks
- B. delete protocols bgp group OVERLAY accept-remote-nexthop
- C. set protocols bgp group PROVIDER export LOOPBACKS
- D. delete protocols bgp group UNDERLAY advertise-external

Answer: B

Explanation:
* Understanding the Configuration:

* The exhibit shows a BGP configuration on a Border-Leaf device. The BGP group UNDERLAY is used for the underlay network, OVERLAY for EVPN signaling, and PROVIDER for connecting to the provider network.
* The OVERLAY group has the accept-remote-nexthop statement, which is designed to accept the next-hop address learned from the remote peer as is, without modifying it.
* Problem Identification:
* The BGP session to Border-Leaf-2 is not established. A common issue in EVPN-VXLAN environments is related to next-hop reachability, especially when accept-remote-nexthop is configured.
* In typical EVPN-VXLAN setups, the next-hop address should be reachable within the overlay network. However, the accept-remote-nexthop can cause issues if the next-hop IP address is not directly reachable or conflicts with the expected behavior in the overlay.
* Corrective Action:
* D. delete protocols bgp group OVERLAY accept-remote-nexthop:Removing this command will ensure that the device uses its own IP address as the next-hop in BGP advertisements, which is standard practice in many EVPN-VXLAN setups. This change should help establish the BGP session with Border-Leaf-2.
Data Center References:
* Proper handling of BGP next-hop attributes is critical in establishing and maintaining stable BGP sessions, especially in complex multi-fabric environments like EVPN-VXLAN. Removing accept- remote-nexthop aligns with best practices in many scenarios.

NEW QUESTION # 49
.....

If you want to quickly study JN0-683 exam questions, printed in the manuscripts to convenient their record at any time, you can choose to PDF model of JN0-683 guide torrent Simulated test, of course, if you want to achieve online, real-time test their learning effect, our JN0-683 study quiz will provide you the Software model, it can make you better in the real test environment to exercise your ability to solve the problem and speed. Finally, if you think that you want to practice with other eletronic devices, you can choose the JN0-683 practice materials by using Online version.

learning.benindonesia.co.id, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, Disposable vapes

P.S. Free & New JN0-683 dumps are available on Google Drive shared by TrainingQuiz: https://drive.google.com/open?id=1qRXs31qkSVp0ET74Pu9IW8LABPSJhXtO