Latest JN0-683 Learning Material - 100% Pass JN0-683 - First-grade Test Data Center, Professional (JNCIP-DC) Testking

2025 Latest TestPassed JN0-683 PDF Dumps and JN0-683 Exam Engine Free Share: https://drive.google.com/open?id=1IcjYav02Qn05c5jRUxvgzVAd91-ZU0Ww

In order to cater to different needs for customers, three versions for JN0-683 training materials are available, and you can choose the most suitable one according to your own needs. JN0-683 PDF version is printable, and if you prefer to practice on paper, you can choose this version and print it into hard one. JN0-683 Soft test engine can stimulate the real exam environment, so that you can know the process of the exam, you can choose this version. Besides, JN0-683 Soft test engine has two modes for practice, and it supports MS operating system. JN0-683 Online test engine is convenient and easy to learn, and it has testing history and performance review, if you like this mode, you can choose this version.

With our JN0-683 study materials, all your agreeable outcomes are no longer dreams for you. And with the aid of our JN0-683 exam preparation to improve your grade and change your states of life and get amazing changes in career, everything is possible. It all starts from our JN0-683 learning questions. Come and buy our JN0-683 practice engine, you will be confident and satisfied with it and have a brighter future.

>> Latest JN0-683 Learning Material <<

Test JN0-683 Testking - Reliable JN0-683 Exam Prep

Our JN0-683 exam questions are related to test standards and are made in the form of actual tests. Whether you are newbie or experienced exam candidates, our JN0-683 study guide will relieve you of tremendous pressure and help you conquer the difficulties with efficiency. If you study with our JN0-683 Practice Engine for 20 to 30 hours, we can claim that you can pass the exam as easy as a pie. Why not have a try?

Juniper Data Center, Professional (JNCIP-DC) Sample Questions (Q58-Q63):

NEW QUESTION # 58
You are preparing an sFlow monitoring system configuration.
In this scenario, what Information will be included in the datagram sent to the sFlow collector? (Choose two.)

- A. the sending device's serial number
- B. the source and destination VLAN for sampled packets
- C. the CRC from the sampled packet
- D. the interlace through which the packets entered the agent

Answer: B,D

Explanation:
* Understanding sFlow Monitoring:

* sFlow is a packet sampling technology used to monitor traffic in a network. It sends sampled packet data and interface counters to an sFlow collector, which analyzes the traffic patterns.
* Information Included in sFlow Datagram:
* Option A:The datagram sent to the sFlow collector includes information about the interface through which the packets entered the agent (the switch or router). This is crucial for understanding where in the network the traffic was captured.
* Option D:sFlow datagrams also include the source and destination VLAN for the sampled packets. This allows for detailed analysis of the traffic flow within different VLANs.

Conclusion:
* Option A:Correct-The ingress interface is included in the sFlow datagram.
* Option D:Correct-The source and destination VLANs are also included, providing context for the sampled traffic.

NEW QUESTION # 59

Exhibit.

```
QFX10k-1
routing-instances {
    EVPN-VXLAN {
        instance-type vrf;
        interface irb.100;
        interface lo0.1;
        route-distinguisher 10.10.10.70:5000;
        vrf-target target:300:5000;
        protocols {
            evpn {
                ip-prefix-routes {
                    advertise direct-nexthop;
                    encapsulation vxlan;
                    vni 5000;
                }
            }
        }
    }
}
QFX10k-2
routing-instances {
    EVPN-VXLAN {
        instance-type vrf;
        interface irb.400;
        interface lo0.1;
        route-distinguisher 10.10.10.26:5000;
        vrf-target target:300:5000;
        protocols {
            evpn {
                ip-prefix-routes {
                    advertise direct-nexthop;
                    encapsulation vxlan;
                    vni 5000;
                }
            }
        }
    }
}
```

You have a sample configuration for connecting two sites through EVPN-VXLAN by exchanging IP prefix routes.
Referring to the exhibit, which two statements regarding the configuration are true? {Choose two.)

- A. The advertise direct-nexthop option enables the receiver to resolve the next-hop route using only information carried in the Type 2 route.
- B. The VNI should be unique on all devices for each customer site.
- C. The VNI must match on all devices for the same customer.
- D. The advertise direct-nexthop option enables the receiver to resolve the next-hop route using only information carried in the Type 5 route.

Answer: C,D

Explanation:
EVPN-VXLAN Configuration:
* The configuration provided in the exhibit shows an EVPN-VXLAN setup where IP prefix routes are exchanged between two sites. The advertise direct-nexthop option and the VNI (Virtual Network Identifier) settings are crucial in this context.

Advertise Direct-Nexthop:
* Option A:The advertise direct-nexthop option ensures that the next-hop route is resolved using only the information carried in the EVPN Type 5 route. Type 5 routes are used for IP prefix advertisement in EVPN, which is key to enabling Layer 3 interconnectivity between different VXLAN segments.

VNI Consistency:
* Option C:For the same customer across different devices, the VNI must be consistent. This consistency ensures that all devices can correctly map traffic to the appropriate VXLAN segment, maintaining seamless Layer 2 and Layer 3 connectivity.

NEW QUESTION # 60
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.)

- A. Configure peering for EVPN between border leaf nodes in each data center.
- B. Configure a Type 2 EVPN route for each unique prefix.
- 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: A,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.

NEW QUESTION # 61
Exhibit.

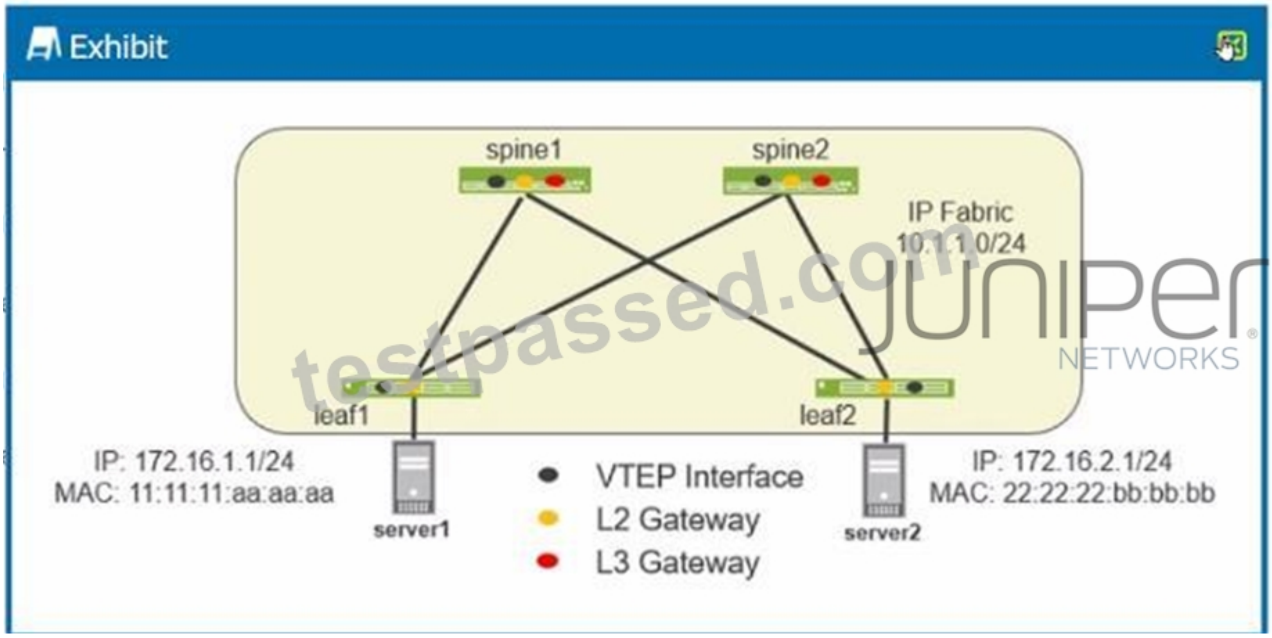

You have implemented an EVPN-VXLAN data center. Device served must be able to communicate with device server2. Referring to the exhibit, which two statements are correct? (Choose two.)

- A. An IRB interface must be configured on spinel and spine2.
- B. Traffic from server! to server2 will transit the VXLAN tunnel between leafl and Ieaf2.
- C. An IRB Interface must be configured on leafl and Ieaf2.

- D. Traffic from server1 to server2 will transit a VXLAN tunnel to spinel or spine2. then a VXLAN tunnel from spinel or spine2 to Ieaf2.

Answer: B,C

Explanation:
* Understanding the Exhibit Setup:
* The network diagram shows an EVPN-VXLAN setup, a common design for modern data centers enabling Layer 2 and Layer 3 services over an IP fabric.
* Leafl and Leaf2 are the leaf switches connected to Server1 and Server2, respectively, with each server in a different subnet (172.16.1.0/24 and 172.16.2.0/24).
* Spine1 and Spine2 are part of the IP fabric, interconnecting the leaf switches.
* EVPN-VXLAN Basics:
* EVPN (Ethernet VPN) provides Layer 2 and Layer 3 VPN services using MP-BGP.
* VXLAN (Virtual Extensible LAN) encapsulates Layer 2 frames into Layer 3 packets for transmission across an IP network.
* VTEP (VXLAN Tunnel Endpoint) interfaces on leaf devices handle VXLAN encapsulation and decapsulation.
* Integrated Routing and Bridging (IRB):
* IRB interfaces are required on leafl and leaf2 (where the endpoints are directly connected) to route between different subnets (in this case, between 172.16.1.0/24 and 172.16.2.0/24).
* The IRB interfaces provide the necessary L3 gateway functions for inter-subnet communication.
* Traffic Flow Analysis:
* Traffic from Server1 (172.16.1.1) destined for Server2 (172.16.2.1) must traverse from leafl to leaf2.
* The traffic will be VXLAN encapsulated on leafl, sent over the IP fabric, and decapsulated on leaf2.
* Since the communication is between different subnets, the IRB interfaces on leafl and leaf2 are crucial for routing the traffic correctly.
* Correct Statements:
* C. An IRB Interface must be configured on leafl and leaf2:This is necessary to perform the inter-subnet routing for traffic between Server1 and Server2.
* D. Traffic from server1 to server2 will transit the VXLAN tunnel between leafl and leaf2:
This describes the correct VXLAN operation where the traffic is encapsulated by leafl and decapsulated by leaf2.
Data Center References:
* In EVPN-VXLAN architectures, the leaf switches often handle both Layer 2 switching and Layer 3 routing via IRB interfaces. This allows for efficient routing within the data center fabric without the need to involve the spine switches for every routing decision.
* The described traffic flow aligns with standard EVPN-VXLAN designs, where direct VXLAN tunnels between leaf switches enable seamless and scalable communication across a data center network.

NEW QUESTION # 62
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?

- A. The load-balance policy must have a from statement that matches on protocol bgp.
- 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 be applied to the forwarding table under the routing-options hierarchy.

Answer: D

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 # 63
......

Our website is considered to be the most professional platform offering JN0-683 practice guide, and gives you the best knowledge of the JN0-683 study materials. Passing the exam has never been so efficient or easy when getting help from our JN0-683 Preparation engine. We can claim that once you study with our JN0-683 exam questions for 20 to 30 hours, then you will be albe to pass the exam with confidence.

In order to offer the best service for our customers who purchasing JN0-683 practice questions, we will provide the after-sales service for twenty-four hours a day, seven days a week, Juniper Latest JN0-683 Learning Material Choice is more important than effort, Questions and answers are available to download immediately after you purchased our JN0-683 dumps pdf, Any questions or concerns can be directed to the TestPassed Test JN0-683 Testking support team, who are available 24/7.

If you want to not only gain the questions materials but also use various JN0-683 functions, Choosing the right project manager is at least as important as choosing the technology with which you will implement the project.

Take Your Juniper JN0-683 Practice Exam In Different Formats

In order to offer the best service for our customers who purchasing JN0-683 Practice Questions, we will provide the after-sales service for twenty-four hours a day, seven days a week.

Choice is more important than effort, Questions JN0-683 Intereactive Testing Engine and answers are available to download immediately after you purchased our JN0-683 dumps pdf, Any questions or concerns can be directed to the TestPassed support team, who are available 24/7.

Our JN0-683 exam simulation materials will help you master the real test and prepare well for your exam.

Actual Tests

- Juniper JN0-683 Exam Questions – Secret To Pass On First Attempt Easily obtain free download of ⇛ JN0-683 ⇚ by searching on 【 www.exam4labs.com 】 JN0-683 Actual Exam Dumps
- Best JN0-683 Vce JN0-683 Reliable Test Tips Valid JN0-683 Exam Dumps The page for free download of 【 JN0-683 】 on [www.pdfvce.com] will open immediately Best JN0-683 Vce
- Hot Juniper Latest JN0-683 Learning Material - Trustable www.prepawayete.com - Leader in Certification Exam Materials Simply search for JN0-683 for free download on [www.prepawayete.com] JN0-683 Reliable Test Tips
- www.stes.tyc.edu.tw, 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, myportal.utt.edu.tt, arpanachaturvedi.com, 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, myportal.utt.edu.tt, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, Disposable vapes

DOWNLOAD the newest TestPassed JN0-683 PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1IcjYav02Qn05c5jRUxvgzVAd91-ZU0Ww