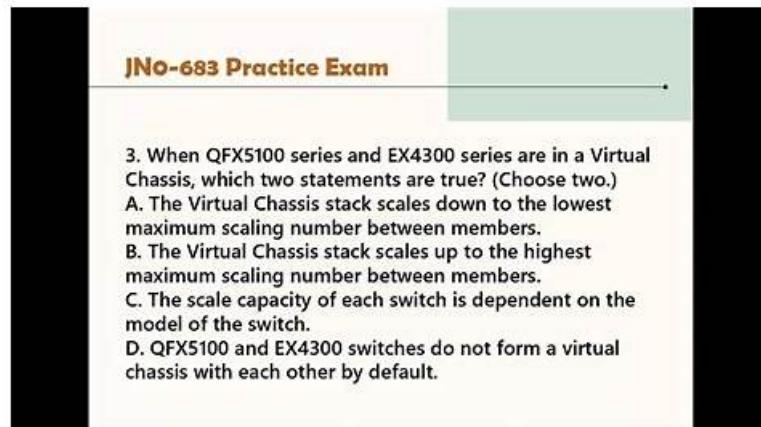


JN0-683 Latest Exam Format | JN0-683 Updated Test Cram



BTW, DOWNLOAD part of PrepAwayExam JN0-683 dumps from Cloud Storage: <https://drive.google.com/open?id=1f7uix1bGT3OQwXE7m0Fw9Q2IaCGlMAdl>

Perhaps now you are one of the candidates of the Juniper JN0-683 exam, perhaps now you are worried about not passing the exam smoothly. Now we have good news for you: our JN0-683 study materials will solve all your worries and help you successfully pass it. With the high pass rate as 98% to 100%, you will find that we have the best Data Center, Professional (JNCIP-DC) JN0-683 learning braindumps which contain the most accurate real exam questions.

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">• 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	<ul style="list-style-type: none">• 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 5	<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 Latest Exam Format <<

JN0-683 Latest Exam Format & Free Download JN0-683 Updated Test Cram Promise You to Purchase Safely and Easily

The Data Center, Professional (JNCIP-DC) (JN0-683) practice questions give you a feeling of a real exam which boost confidence.

Practice under real Data Center, Professional (JNCIP-DC) (JN0-683) exam situations is an excellent way to learn more about the complexity of the Juniper JN0-683 Exam Dumps. You can learn from your Data Center, Professional (JNCIP-DC) (JN0-683) practice test mistakes and overcome them before the actual JN0-683 exam.

Juniper Data Center, Professional (JNCIP-DC) Sample Questions (Q64-Q69):

NEW QUESTION # 64

Exhibit.



A Exhibit

```
routing-instances {
    tenant1 {
        instance-type vrf;
        routing-options {
            auto-export {
                family inet {
                    unicast;
                }
            }
        }
        protocols {
            evpn {
                ip-prefix-routes {
                    advertise direct-next-hop;
                    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?

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

Answer: A

Explanation:

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.

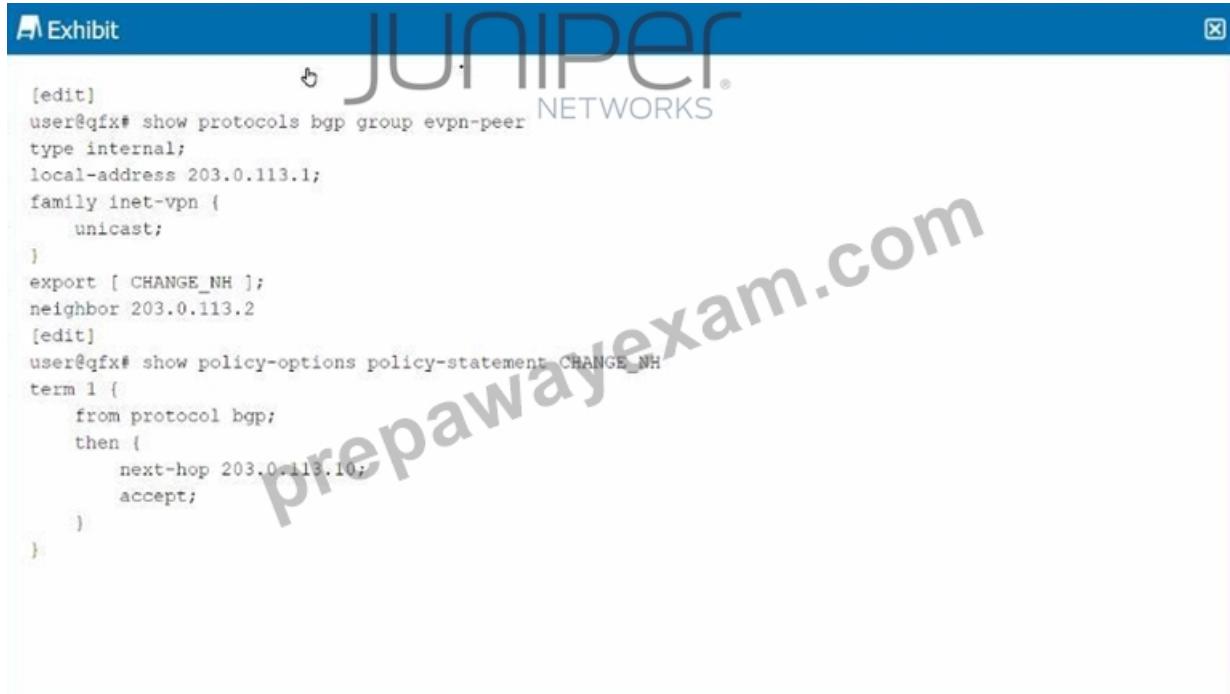
* Option C: Adding VXLAN encapsulation may be necessary for other scenarios but does not directly address the Type 5 route advertisement.

* 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 # 65

Exhibit.



The exhibit shows a Juniper QFX terminal window with the following configuration:

```
[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?

- A. EVPN routes cannot have the next hop changed.
- B. The `vpn-apply-export` parameter must be applied to this peer.**
- C. The export policy is incorrectly configured.
- D. The `vrf-export` parameter must be applied.

Answer: B

Explanation:

* 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.

NEW QUESTION # 66

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.)

- A. Design a five-stage Clos IP fabric.
- B. Use EX4300 Series devices as the spine devices.
- C. Use QFX5700 Series devices as the super spines.
- D. Design a three-stage Clos IP fabric.

Answer: C,D

Explanation:

* 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 # 67

Which two statements are true about IP fabrics using unnumbered BGP? (Choose two.)

- A. Unnumbered BGP requires that family inet is configured on each interface.
- B. Unnumbered BGP peering automatically provisions IPv6 peering.
- C. Unnumbered BGP peering automatically provisions IPv4 peering.
- D. Unnumbered BGP requires that family inet is configured on each interface.

Answer: A,C

Explanation:

* Understanding Unnumbered BGP:

* Unnumbered BGP (Border Gateway Protocol) allows BGP peering between routers without assigning specific IP addresses to the interfaces. Instead, it uses the loopback address or another router identifier for the BGP session, making IP address management more straightforward in large-scale networks.

* Family inet Configuration:

* Option C: The family inet configuration is required on each interface involved in unnumbered BGP peering to support IPv4 address families. This ensures that IPv4 peering sessions can be established between devices.

* Automatic IPv4 Peering:

* Option D: Unnumbered BGP peering automatically provisions IPv4 peering sessions. This simplifies the configuration by eliminating the need to manually assign and manage IP addresses for BGP peering.

Conclusion:

* Option C: Correct - Unnumbered BGP requires the family inet configuration for IPv4.

* Option D: Correct - Unnumbered BGP automatically provisions IPv4 peering, simplifying setup.

NEW QUESTION # 68

You are deploying an IP fabric using EBGP and notice that your leaf devices are advertising and receiving all the routes. However, the routes are not installed in the routing table and are marked as hidden.

Which two statements describe how to solve the issue? (Choose two.)

- A. You need to configure loops 2.
- B. You need to configure as-override.

- C. You need to configure a next-hop self policy.
- D. You need to configure multipath multiple-as.

Answer: A,D

NEW QUESTION # 69

PrepAwayExam JN0-683 exam dumps have been designed with the best possible format, ensuring all necessary information packed in them. Our experts have used only the authentic and recommended sources of studies by the certifications vendors for exam preparation. The information in the JN0-683 Brain Dumps has been made simple up to the level of even an average exam candidate. To ease you in your preparation, each JN0-683 dumps are made into easy English so that you learn information without any difficulty to understand them.

JN0-683 Updated Test Cram: <https://www.prepawayexam.com/Juniper/braindumps.JN0-683.ete.file.html>

What's more, part of that PrepAwayExam JN0-683 dumps now are free: <https://drive.google.com/open?id=1f7uix1bGT3OQwXE7m0Fw9Q2IaCGIMAdl>