検証する-更新するJN0-683資格取得試験-試験の準備方法JN0-683ミシュレーション問題

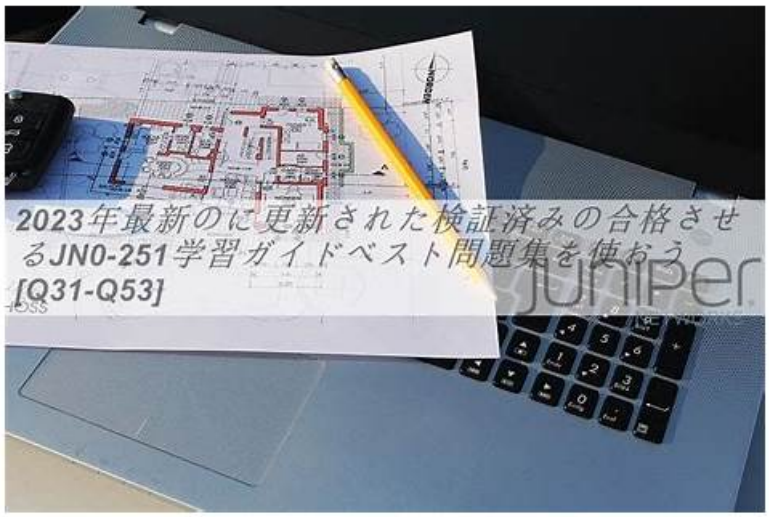

P.S. JPTestKingがGoogle Driveで共有している無料かつ新しいJN0-683ダンプ：https://drive.google.com/open?id=10bRci0vX2wj0SbWxf6gdl8QDQNgBit7K

弊社のJN0-683問題集は三種類の版を提供いたします。PDF版、ソフト版とオンライン版があります。PDF版のJN0-683日本語問題集は印刷されることができ、ソフト版のJN0-683日本語問題集はいくつかのパソコンでも使われることもでき、オンライン版の問題集はパソコンでもスマホでも直接に使われることができます。お客様は自分の愛用する版が選べます。

JN0-683試験問題の更新を1年以内にクライアントに無料で提供し、1年後にクライアントは50％の割引を受けることができます。クライアントが古いクライアントの場合、一定の割引を享受できます。当社Juniperの専門家は、毎日JN0-683ガイドトレントを更新し、JN0-683スタディガイドの最新の更新をクライアントに提供します。私たちはクライアントに割引を提供し、彼らがより少ないお金を使うようにします。あなたが古いクライアントである場合、あなたは特別割引を享受することができますので、お金を節約することができます。したがって、JN0-683テストトレントを購入することは非常に価値があります。

>> JN0-683資格取得 <<

Juniper JN0-683ミシュレーション問題、JN0-683日本語認定

質問と回答のみを提供するPDFバージョンの機能に満足できない場合は、JN0-683試験の教材のAPPバージョンでさらに多くを提供できます。 APPバージョンは、実際のテストシーンをシミュレートするだけでなく、間違いを指摘し、何度も練習することに気付くことができます。 Juniper JN0-683試験の教材のこのバージョンはかなり強力です。 あなたが喜んでいるなら、あなたは毎日あなたのパフォーマンスをマークし、比較的あなたの勉強と準備を調整することができます。 JN0-683試験の教材は、お客様の要求を満たすために最善を尽くします。

Juniper JN0-683 認定試験の出題範囲:

トピック	出題範囲
トピック 1	• データ センターのマルチテナントとセキュリティ:このセクションでは、シングルテナントおよびマルチテナント データ センターの設定に関する知識がテストされます。データ センター プロフェッショナルなどの候補者は、共有インフラストラクチャ環境におけるレイヤー 2 レベルとレイヤー 3 レベルの両方でテナント トラフィックの分離を確実に行えるかどうかが評価されます。

トピック 2	• データ センターの展開と管理: このセクションでは、主要な展開コンセプトに焦点を当てて、アーキテクトやエンジニアなどのデータ センター ネットワーク プロフェッショナルの専門知識を評価します。トピックには、手動入力なしでデータ センターのデバイス セットアップを自動化するゼロ タッチ プロビジョニング (ZTP) が含まれます。
トピック 3	• EVPN-VXLAN シグナリング: このセクションでは、ルート タイプ、マルチキャスト処理、マルチプロトコル BGP (MBGP) などのイーサネット VPN (EVPN) の概念に関する理解を評価します。また、CRB や ERB、MAC 学習、対称ルーティングなどの EVPN アーキテクチャについても説明します。

Juniper Data Center, Professional (JNCIP-DC) 認定 JN0-683 試験問題 (Q48-Q53):

質問 # 48
You are implementing seamless stitching between two data centers and have a proposed configuration for a border leaf device.
In this scenario, which two statements are correct? {Choose two.)

- A. The translation-vni must match in both data centers.
- B. The ESI must match in both data centers.
- C. The ESI must be different in each data center.
- D. The translation-vni must be different in each data center.

正解：B、D

解説：
When implementing seamless VXLAN stitching between two data centers, the Ethernet Segment Identifier (ESI) must match in both data centers to ensure that the same multi-homed segment is recognized consistently across the environments. This allows seamless failover and redundancy. However, the translation VNI (Virtual Network Identifier) must be different in each data center because VXLAN stitching involves mapping different VNIs to enable interconnectivity between distinct VXLAN domains.
If the same translation VNI is used in both data centers, there would be no differentiation between network segments, leading to potential routing and forwarding issues. Keeping them different ensures proper traffic isolation and mapping.

質問 # 49
Exhibit.
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 PROVIDER export LOOPBACKS
- B. delete protocols bgp group OVERLAY accept-remote-nexthop
- C. delete protocols bgp group UNDERLAY advertise-external
- D. set protocols bgp group overlay export loopbacks

正解：B

解説：
* 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.

質問 # 50
Click the Exhibit button.
Referring to the configuration shown in the exhibit, assume that there is no external router present and that the configuration is fabric-only.
Which two statements are true about the example configuration? (Choose two.)

- A. Devices in irb.400 (vlan 400) are not able to communicate directly with devices in routing instance Customer A.
- B. Devices in irb.400 (vlan 400) and irb.800 (vlan 800) are able to communicate over the fabric.
- C. Devices in routing instance Customer A are able to communicate with devices in routing instance Customer B.
- D. VNI 10006 is assigned to vlan 800 (irb.800).

正解：A、B

解説：
The exhibit shows configurations for two VRFs (Customer_A and Customer_B) with specific VLANs and VNIs assigned. Each VRF has interfaces (IRBs) associated with particular VLANs.
Option B: VLAN 400 (irb.400) is part of Customer_B, and there is no direct connection or routing between Customer_A and Customer_B in the configuration provided. Therefore, devices in irb.400 cannot communicate directly with devices in the Customer_A routing instance.
Option D: Since irb.400 (VLAN 400) and irb.800 (VLAN 800) are part of the same routing instance (Customer_B), they can communicate over the fabric using VXLAN encapsulation.

質問 # 51
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 as an export policy to your BGP
- C. The load-balance policy must be applied to the forwarding table under the routing-options hierarchy.
- D. The multipast multiple -as configuration must be configured for each peer in the BGP spine group.

正解：C

解説：
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.

質問 # 52
You are deploying an IP fabric with an oversubscription ratio of 3:1.
In this scenario, which two statements are correct? (Choose two.)

- A. The oversubscription ratio remains the same when you add leaf devices.
- B. The oversubscription ratio increases when you remove leaf devices.
- C. The oversubscription ratio decreases when you add leaf devices.
- D. The oversubscription ratio remains the same when you remove leaf devices.

正解：A、B

解説：
Adding leaf devices increases the edge bandwidth (workload-facing ports), which lowers the oversubscription ratio because the amount of bandwidth facing the workloads increases relative to the fabric bandwidth.
However, the oversubscription ratio itself remains the same when you add leaf devices if the fabric design and spine capacity remain unchanged because the ratio is a function of edge bandwidth to spine bandwidth. Adding leaf devices without changing spine devices keeps the ratio at the same level.
Removing leaf devices would increase the oversubscription ratio because the amount of edge bandwidth decreases while spine bandwidth remains the same.

質問 # 53
.....

- JN0-683資格勉強 □ JN0-683資格勉強 □ JN0-683最新な問題集 □ ▹ www.goshiken.com ◃を開き、➡ JN0-683 □□□を入力して、無料でダウンロードしてくださいJN0-683無料ダウンロード
- 最高-実際的なJN0-683資格取得試験-試験の準備方法JN0-683ミシュレーション問題 □ 今すぐ【 www.passtest.jp 】を開き、【 JN0-683 】を検索して無料でダウンロードしてくださいJN0-683試験解答
- JN0-683認証資格 □ JN0-683日本語版 □ JN0-683無料模擬試験 □ □ www.goshiken.com □で➽ JN0-683 □ □を検索し、無料でダウンロードしてくださいJN0-683模擬資料
- JN0-683資格取得を使用して - Data Center, Professional (JNCIP-DC)に別れを告げる □ ➡ www.goshiken.com □□□にて限定無料の【 JN0-683 】問題集をダウンロードせよJN0-683最新な問題集
- JN0-683資格勉強 □ JN0-683資格勉強 □ JN0-683認証資格 □ □ JN0-683 □の試験問題は《 www.goshiken.com 》で無料配信中JN0-683無料模擬試験
- JN0-683復習資料 ۞ JN0-683テスト難易度 □ JN0-683出題内容 □ ⇒ www.passtest.jp ⇐は、□ JN0-683 □を無料でダウンロードするのに最適なサイトですJN0-683資格勉強
- 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, myportal.utt.edu.tt, www.stes.tyc.edu.tw, myportal.utt.edu.tt, 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, Disposable vapes

2026年JPTestKingの最新JN0-683 PDFダンプおよびJN0-683試験エンジンの無料共有：https://drive.google.com/open?id=10bRci0vX2wj0SbWxf6gdl8QDQNgBit7K