

# Juniper JN0-683 Exam | JN0-683 Learning Materials - 10 Years of Excellence of JN0-683 Formal Test

**Exams Test Knowledge, But Life Tests Confidence**

Give your child the skills that matter beyond the classroom

**ENROLL NOW**

**Workshop Highlights**

- 23 Dec (Tue) | 26 Dec (Fri)
- 10:00am - 5:00pm
- Fee : RM 280 per child | RM 520 for 2 children (Early Bird promotion available)

**Join Our 1 day Public Speaking workshop**

2A-3, Block 2, Jalan Jati Jaya 2, Jati Link, Jalan 1/155B, Bukit Jalil, 57000 Kuala Lumpur, Wilayah Persekutuan Kuala Lumpur

017-765 2388 | JohanBukitJalil

BONUS!!! Download part of TroytecDumps JN0-683 dumps for free: <https://drive.google.com/open?id=1KWA0zSJKnt5-coT5QUiYbrB0kWaNXxFU>

Will you feel nervous while facing the real exam? Choose us, since we will help you relieve your nerves. JN0-683 Soft test engine can stimulate the real exam environment, so that you can know the procedure of the exam, and your confidence for the exam will be strengthened. In addition, JN0-683 exam dumps are edited by professional experts, who are quite familiar with the exam center, therefore the quality can be guaranteed. We offer you free demo for JN0-683 to have a try before buying. And you will receive the downloading link and password within ten minutes for JN0-683 exam materials, so that you can start your learning immediately.

## Juniper JN0-683 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"> <li>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.</li> </ul>
Topic 2	<ul style="list-style-type: none"> <li>VXLAN: This part requires knowledge of VXLAN, particularly how the control plane manages communication between devices, while the data plane handles traffic flow. Demonstrate knowledge of how to configure, Monitor, or Troubleshoot VXLAN.</li> </ul>

Topic 3	<ul style="list-style-type: none"> <li>• <b>Data Center Multitenancy and Security:</b> 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.</li> </ul>
Topic 4	<ul style="list-style-type: none"> <li>• <b>Data Center Interconnect:</b> 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.</li> </ul>

>> JN0-683 Learning Materials <<

## JN0-683 Formal Test, Free JN0-683 Updates

With JN0-683 test guide, you only need a small bag to hold everything you need to learn. In order to make the learning time of the students more flexible, JN0-683 exam materials specially launched APP, PDF, and PC three modes. With the APP mode, you can download all the learning information to your mobile phone. In this way, whether you are in the subway, on the road, or even shopping, you can take out your mobile phone for review. JN0-683 study braindumps also offer a PDF mode that allows you to print the data onto paper so that you can take notes as you like and help you to memorize your knowledge. At the same time, regardless of which mode you use, JN0-683 test guide will never limit your download times and the number of concurrent users. For the same information, you can use it as many times as you want, and even use together with your friends.

## Juniper Data Center, Professional (JNCIP-DC) Sample Questions (Q25-Q30):

### NEW QUESTION # 25

Exhibit.

```

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

**Answer: A**

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 # 26

Exhibit.

```
user@Leaf-1> show configuration switch-options
service-id 1;
route-distinguisher 192.168.100.51:1;
vrf-target target:65000:55;
user@Leaf-2> show configuration switch-options
vtep-source-interface lo0.0;
route-distinguisher 192.168.100.51:2;
vrf-target target:65000:54;
```

Connections between hosts connected to Leaf 1 and Leaf-2 are not working correctly.

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: 65000:55 parameter on Leaf-2.
- B. Configure the set switch-options service-id 1 parameter on Leaf-2.
- C. Configure the set switch-options vtep-source-interface 100.0 parameter on Leaf-1.
- D. Configure the set switch-options route-distinguisher 192.168.100.51:2 parameter on Leaf-1.

**Answer: A,B**

Explanation:

\* Review of the Exhibit:

\* The exhibit shows the switch configuration for Leaf 1 and Leaf-2. The configurations include route distinguishers, VRF targets, and service IDs, all of which are crucial for ensuring proper operation in an EVPN-VXLAN environment.

\* Service-ID Consistency:

\* The service ID must be consistent across all participating leaf devices in the same EVPN instance to ensure that they are part of the same VXLAN overlay network.

\* VRF Target Consistency:

\* The vrf-target parameter must also be consistent across devices to ensure that VRFs (Virtual Routing and Forwarding instances)

are correctly imported and exported between leaf nodes.

Conclusion:

- \* Option B:Correct-Setting the same service-id on Leaf-2 ensures that it is part of the same VXLAN overlay as Leaf-1.
- \* Option D:Correct-The vrf-target on Leaf-2 should match Leaf-1 to ensure consistent routing policies and proper route exchange.

### NEW QUESTION # 27

What are two ways in which an EVPN-signaled VXLAN is different from a multicast-signaled VXLAN?

(Choose two.)

- A. An EVPN-signaled VXLAN features slower and more complete convergence.
- **B. An EVPN-signaled VXLAN can perform autodiscovery of VTEPs using BGP.**
- **C. An EVPN-signaled VXLAN is less resource intensive.**
- D. An EVPN-signaled VXLAN can perform autodiscovery of VTEPs using IS-IS.

**Answer: B,C**

Explanation:

\* Multicast-Signaled VXLAN:

\* In traditional multicast-signaled VXLAN, VTEPs (VXLAN Tunnel Endpoints) use multicast to flood and learn about remote VTEPs. This method relies on multicast in the underlay network to distribute BUM (Broadcast, Unknown unicast, and Multicast) traffic.

\* This approach can be resource-intensive due to the need for multicast group management and increased network traffic, especially in large deployments.

\* EVPN-Signaled VXLAN:

\* EVPN-signaled VXLAN uses BGP (Border Gateway Protocol) to signal the presence of VTEPs and distribute MAC address information. BGP is used for VTEP autodiscovery and the distribution of endpoint information.

\* This method is more efficient because it reduces the reliance on multicast, instead using BGP control-plane signaling to handle VTEP discovery and MAC learning, which reduces the overhead on the network and improves scalability.

\* Correct Statements:

\* B. An EVPN-signaled VXLAN can perform autodiscovery of VTEPs using BGP:This is correct because EVPN uses BGP for VTEP autodiscovery, making it more efficient and scalable compared to multicast-based methods.

\* C. An EVPN-signaled VXLAN is less resource-intensive:This is correct because it eliminates the need for multicast flooding in the underlay, instead using BGP for signaling, which is less demanding on network resources.

\* Incorrect Statements:

\* A. An EVPN-signaled VXLAN can perform autodiscovery of VTEPs using IS-IS:This is incorrect because EVPN relies on BGP, not IS-IS, for VTEP discovery and signaling.

\* D. An EVPN-signaled VXLAN features slower and more complete convergence:This is incorrect; EVPN with BGP typically provides faster convergence due to its use of a control plane rather than relying on data plane learning.

Data Center References:

\* EVPN-VXLAN is widely adopted in modern data center designs due to its scalability, efficiency, and reduced resource consumption compared to multicast-based VXLAN solutions. It leverages the strengths of BGP for control-plane-driven operations, resulting in more efficient and scalable networks.

### NEW QUESTION # 28

Exhibit.

```
user@device> show configuration routing-instances
Customer_B {
  instance-type vrf;
  routing-options {
    graceful-restart;
    multipath;
    auto-export;
  }
  protocols {
    evpn {
      irb-symmetric-routing {
        vni 10006;
      }
      ip-prefix-routes {
        advertise direct-nexthop;
        encapsulation vxlan;
        vni 10006;
        export export_policy;
      }
    }
  }
  interface irb.400;
  interface irb.800;
  interface lo0.3;
  route-distinguisher 172.16.0.2:20;
  vrf-target target:10006:1;
}
Customer_A {
  instance-type vrf;
  routing-options {
    graceful-restart;
    multipath;
    auto-export;
  }
  protocols {
    evpn {
      irb-symmetric-routing {
        vni 10000;
      }
      ip-prefix-routes {
        advertise direct-nexthop;
      }
    }
  }
  instance-type vrf;
  routing-options {
    graceful-restart;
    multipath;
    auto-export;
  }
  protocols {
    evpn {
      irb-symmetric-routing {
        vni 10000;
      }
      ip-prefix-routes {
        advertise direct-nexthop;
        encapsulation vxlan;
        vni 10000;
        export export_policy;
      }
    }
  }
  interface et-0/0/51.5;
  interface irb.3;
  interface irb.300;
  interface irb.1000;
  interface irb.2000;
  interface irb.4000;
  interface lo0.2;
  route-distinguisher 172.16.0.2:2;
  vrf-target target:10000:1;
}
```

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).

**Answer: A,B**

Explanation:

\* Understanding the Configuration:

\* 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.

\* Communication Between VLANs and Routing Instances:

\* 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.

Conclusion:

\* Option B:Correct-There is no direct communication between devices in irb.400 (Customer\_B) and routing instance Customer\_A.

\* Option D:Correct-Devices in VLAN 400 and VLAN 800 can communicate within the Customer\_B routing instance over the fabric.

#### NEW QUESTION # 29

Which two statements are correct about an IP fabric? (Choose two.)

- A. All leaf devices can use the same AS number in an IP fabric without making any adjustments to the EBGp configuration
- B. FBGP is only required to route mostrouting information to external devices outside the fabric.
- C. The multipath multiple-as statement is required to enable ECMP if every device has a different AS number.
- D. Only a single point to point EBGp session is required between peers in an IP fabric.

**Answer: A,C**

Explanation:

\* BGP in IP Fabric:

\* In an IP fabric, Border Gateway Protocol (BGP) is used to manage the routing between leaf and spine devices. Each device can have the same or different Autonomous System (AS) numbers depending on the network design.

\* Multipath Multiple-AS:

\* Option B:If every device in the fabric has a different AS number, then enabling Equal-Cost Multi-Path (ECMP) routing requires the multipath multiple-as statement. This configuration allows BGP to consider multiple paths across different AS numbers as equal cost, enabling efficient load balancing across the network.

\* Same AS Number Configuration:

\* Option A:It's possible for all leaf devices to use the same AS number in an IP fabric, which simplifies the configuration. EBGp (External BGP) will still function correctly in this setup because BGP considers the peering relationship rather than strictly enforcing different AS numbers in this specific use case.

Conclusion:

\* Option B:Correct-This statement is essential for enabling ECMP in a multi-AS environment.

\* Option A:Correct-Leaf devices can share the same AS number without needing special EBGp configuration.

#### NEW QUESTION # 30

.....

We offer 24 - hour, 365 – day online customer service to every user on our JN0-683 study materials. Our service staff will help you solve the problem about the JN0-683 training materials with the most professional knowledge and enthusiasm. We believe that can completely dispel your worries on JN0-683 Exam Braindumps. So please feel free to contact us if you have any trouble on our JN0-683 practice questions.

**JN0-683 Formal Test:** <https://www.troytecdumps.com/JN0-683-troytec-exam-dumps.html>

