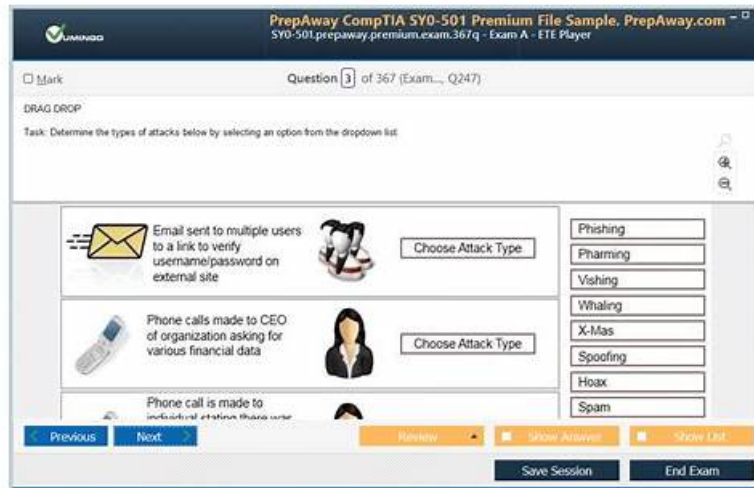


JN0-683 Test Labs & Test JN0-683 Dates



What's more, part of that ExamDumpsVCE JN0-683 dumps now are free: <https://drive.google.com/open?id=1JkjpgHYLFow90rMJQnMx8vrPh5bk-wBm>

In the major environment, people are facing more job pressure. So they want to get Juniper certification rise above the common herd. How to choose valid and efficient JN0-683 guide torrent should be the key topic most candidates may concern. So now, it is right, you come to us. Our company is famous for its high-quality JN0-683 Exam Questions in this field especially for Juniper certification exams. It has been accepted by thousands of candidates who practice our JN0-683 study materials for their exam

Juniper JN0-683 Exam Syllabus Topics:

Topic	Details
Topic 1	<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 2	<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 3	<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.
Topic 4	<ul style="list-style-type: none"> • 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.
Topic 5	<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.

>> JN0-683 Test Labs <<

Authoritative JN0-683 Test Labs – 100% Accurate Test Data Center, Professional (JNCIP-DC) Dates

The language which is easy to be understood and simple, JN0-683 exam questions are suitable for any learners no matter he or she

is a student or the person who have worked for many years with profound experiences. So it is convenient for the learners to master the JN0-683 Guide Torrent and pass the exam in a short time. The amount of the examinee is large. For the office workers, they are both busy in their job and their family life; for the students, they possibly have to learn or do other things.

Juniper Data Center, Professional (JNCIP-DC) Sample Questions (Q56-Q61):

NEW QUESTION # 56

A local VTEP has two ECMP paths to a remote VTEP.

Which two statements are correct when load balancing is enabled in this scenario? (Choose two.)

- A. The source port in the UDP header is used to load balance VXLAN traffic.
- B. The inner packet fields are not used in the hash for load balancing.
- C. The destination port in the UDP header is used to load balance VXLAN traffic.
- D. The inner packet fields are used in the hash for load balancing.

Answer: A,D

Explanation:

The source port in the UDP header is used to load balance VXLAN traffic: In an ECMP (Equal-Cost Multi-Path) scenario with VXLAN, the source port in the UDP header is one of the fields used to hash the traffic and determine which path to use. This helps distribute traffic across multiple equal-cost paths between VTEPs.

The inner packet fields are used in the hash for load balancing: For VXLAN traffic, the inner packet fields (such as the inner source and destination IP addresses, and sometimes the inner VLAN) are included in the hash calculation to determine the forwarding path. This allows for more effective load balancing across the available paths.

NEW QUESTION # 57

Why is a designated forwarder required in a multihomed CE-to-PE VXLAN environment using EVPN signalling?

- A. The designated forwarder is required to prevent a traffic storm from being received on multihomed hosts.
- B. The designated forwarder is required to prevent flooding of MAC addresses to multihomed hosts.
- C. The designated forwarder is required to prevent duplicate packets from being received on multihomed hosts.
- D. The designated forwarder is required to prevent packets from looping between the PEs.

Answer: C

Explanation:

* Understanding Multihomed CE-to-PE VXLAN Environment:

* In a VXLAN environment using EVPN signaling, multiple PEs (Provider Edge devices) can be connected to the same CE (Customer Edge device). This setup is referred to as multihoming, where a CE device has multiple connections to the network to ensure redundancy and load balancing.

* Role of the Designated Forwarder:

* The designated forwarder (DF) is a mechanism used in EVPN to manage the forwarding of broadcast, unknown unicast, and multicast (BUM) traffic in a multihomed environment. The DF is selected to ensure that only one of the PEs forwards this type of traffic to the CE, preventing loops and unnecessary duplicate packets.

* Avoiding Duplicate Packets:

* Without a designated forwarder, all PEs connected to a multihomed CE could potentially forward the same packet to the CE, resulting in duplicate packets. This duplication can cause issues with packet processing on the CE, leading to inefficiencies and potential network problems.

Conclusion:

* Option D: Correct- The designated forwarder is essential to prevent duplicate packets from being received on multihomed hosts, ensuring that only one PE forwards BUM traffic to the CE.

NEW QUESTION # 58

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

Answer: A

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 # 59

Click the Exhibit button. The exhibit shows the truncated output of the show evpn database command.

```
user@leaf1> show evpn database
Instance: evpn-1
VLAN DomainId MAC address Active source Timestamp
address
10001 00:1c:73:00:00:01 irb.4000 Apr 16 11:46:14
10.4.4.1
10001 40:00:dc:01:00:01 00:02:00:00:00:00:04:00:00:04 Apr 16 11:46:14
10.4.4.2
10001 40:00:dc:01:00:02 00:02:00:00:00:00:04:00:00:04 Apr 16 11:46:14
10.4.4.3
10001 40:00:dc:01:00:03 00:02:00:00:00:00:04:00:00:04 Apr 16 11:46:14
10.4.4.4
10001 40:00:dc:01:00:04 00:02:00:00:00:00:04:00:00:04 Apr 16 11:46:14
10.4.4.5
10001 40:00:dc:01:00:05 00:02:00:00:00:00:04:00:00:04 Apr 16 11:46:14
10.4.4.6
10001 44:11:01:00:00:01 00:02:00:00:00:00:04:00:00:04 Apr 16 11:46:14
10001 44:11:01:00:00:02 00:02:00:00:00:00:04:00:00:04 Apr 16 11:46:14
10001 44:11:01:00:00:03 00:02:00:00:00:00:04:00:00:04 Apr 16 11:46:14
10001 44:11:01:00:00:04 00:02:00:00:00:00:04:00:00:04 Apr 16 11:46:14
10001 44:11:01:00:00:05 00:02:00:00:00:00:04:00:00:04 Apr 16 11:46:14
10001 44:12:01:00:00:01 00:02:00:00:00:00:03:00:00:03 Apr 16 11:46:14
10001 44:12:01:00:00:02 00:02:00:00:00:00:03:00:00:03 Apr 16 11:46:14
10001 44:12:01:00:00:03 00:02:00:00:00:00:03:00:00:03 Apr 16 11:46:14
10001 44:12:01:00:00:04 00:02:00:00:00:00:03:00:00:03 Apr 16 11:46:14
10001 44:12:01:00:00:05 00:02:00:00:00:00:03:00:00:03 Apr 16 11:46:14
10002 00:1c:73:00:00:01 irb.300 Apr 16 11:46:14
10.3.3.1
10002 30:00:dc:01:00:01 00:02:00:00:00:00:01:00:00:01 Apr 16 11:46:14
10002 30:00:dc:01:00:02 00:02:00:00:00:00:01:00:00:01 Apr 16 11:46:14
10002 30:00:dc:01:00:03 00:02:00:00:00:00:01:00:00:01 Apr 16 11:46:14
10002 30:00:dc:01:00:04 00:02:00:00:00:00:01:00:00:01 Apr 16 11:46:14
```

Given this output, which two statements are correct about the host with MAC address 40:00:dc:01:00:04? (Choose two.)

- A. The host is located on VNI 10002.
- B. The host is assigned IP address 10.4.4.5.
- C. The host is originating from an ESI LAG.
- D. The host is originating from irb.300.

Answer: B,C

Explanation:

The host is assigned IP address 10.4.4.5: The MAC address 40:00:dc:01:00:04 is listed alongside the IP address 10.4.4.5 in the output, indicating that this host is associated with the IP address 10.4.4.5.

The host is originating from an ESI LAG: The Active source value 02:02:00:00:00:04:00:04 is in the Ethernet Segment Identifier (ESI) format, which indicates the host is behind an EVPN multihomed ESI LAG.

NEW QUESTION # 60

Host A is connected to vlan 100 on leaf. Host B is connected to vlan 200 on leaf1. Host A and Host B are unable to communicate. You have reviewed the routing and your hosts have the correct default route (.1) Referring to the exhibit, which two commands will solve the problem? (Choose two.)

```
Exhibit
user@leaf1> show configuration
...
interfaces {
  ge-0/0/0 {
    description "facing_spine1:ge-0/0/1";
    speed 10g;
    mtu 9192;
    unit 0 {
      family inet {
        mtu 9170;
        address 172.16.0.9/31;
      }
    }
  }
  ge-0/0/1 {
    description "facing_spine2:ge-0/0/1";
    speed 10g;
    mtu 9192;
    unit 0 {
      family inet {
        mtu 9170;
        address 172.16.0.11/31;
      }
    }
  }
  irb {
    unit 200 {
      family inet {
        address 192.168.200.1/24;
      }
    }
  }
}
vlans {
  vn100 {
    vlan-id 100;
    description "BLUE";
  }
  vn200 {
    description RED;
    vlan-id 200;
    13-interface irb.200;
  }
}
}
```

- A. set vlans vn100 13-interface irb.100
- B. delete vlans vn200 13-interface irb.200
- C. set interfaces irb unit 100 family inet address 192-168.100.1
- D. set routing-options static route 0.0.0.0/0 next-hop 192.168.200.10

Answer: A,D

Explanation:

VLAN 100 (BLUE) does not have an IRB interface configured, meaning there is no Layer 3 gateway for hosts in VLAN 100.

Adding an IRB.100 interface with an IP address will allow routing between VLAN 100 and VLAN 200.

Even after configuring the IRB interface for VLAN 100, you need to bind it to VLAN 100 so that the switch correctly associates VLAN 100 traffic with the IRB. This ensures that the VLAN can route traffic through the correct Layer 3 gateway.

NEW QUESTION # 61

.....

Do you feel bored about current jobs and current life? Go and come to obtain a useful certificate! JN0-683 study guide is the best product to help you achieve your goal. If you pass exam and obtain a certification with our JN0-683 study materials, you can apply

