Vendor: Juniper

Exam Code: JN0-683

Exam Name: Juniper Data Center Professional (JNCIP-DC)

Version: DEMO

QUESTION 1
Exhibit.

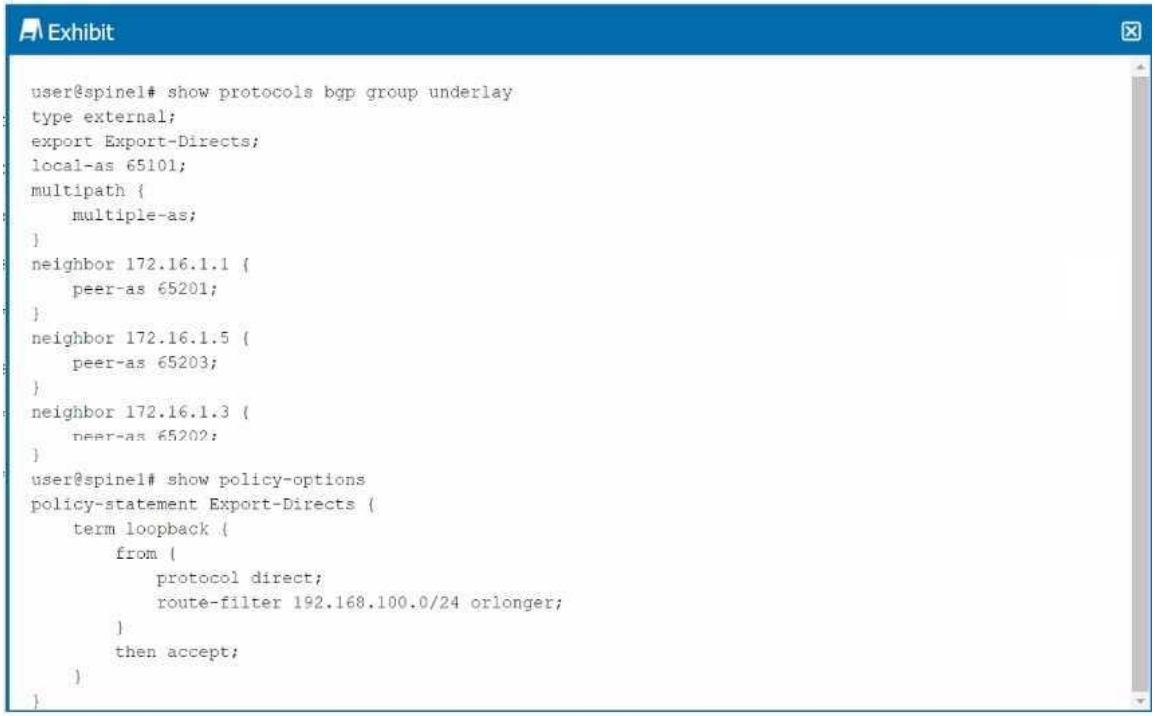

Referring to the exhibit, the spinel device has an underlay BGP group that is configured to peer with its neighbors' directly connected interfaces. Which two statements are true in this scenario? (Choose two.)

A. The multihop statement is not required to establish the underlay BGP sessions.
B. Load balancing for the underlay is not configured correctly.
C. The multihop statement is required to establish the underlay BGP sessions.
D. Load balancing for the underlay is configured correctly.

Answer: AD
Explanation:
Understanding BGP Configuration in the Exhibit:
The exhibit shows a BGP configuration on spine1 with a group named underlay, configured to peer with directly connected interfaces of other devices in the network.
Multipath multiple-as: This statement allows the router to install multiple paths in the routing table for routes learned from different ASes, facilitating load balancing.
A. The multihop statement is not required to establish the underlay BGP sessions: In this case, the BGP peers are directly connected (as indicated by their neighbor IP addresses), so the multihop statement is unnecessary. Multihop is typically used when BGP peers are not directly connected and packets need to traverse multiple hops.
D. Load balancing for the underlay is configured correctly: The multipath { multiple-as; } statement in the configuration enables load balancing across multiple paths from different autonomous systems, which is appropriate for underlay networks in data center fabrics.

QUESTION 2
You manage an IP fabric with an EVPN-VXLAN overlay. You have multiple tenants separated using multiple unique VRF instances. You want to determine the routing information that belongs

in each routing instance's routing table.
In this scenario, which property is used for this purpose?

A. the VRF target community
B. the routing instance type
C. the VRF table label
D. the route distinguisher value

Answer: D
Explanation:
Understanding VRF and Routing Instances:
In an EVPN-VXLAN overlay network, multiple tenants are separated using unique VRF (Virtual Routing and Forwarding) instances. Each VRF instance maintains its own routing table, allowing for isolated routing domains within the same network infrastructure.
Role of Route Distinguisher:
Route Distinguisher (RD): The RD is a unique identifier used in MPLS and EVPN environments to distinguish routes belonging to different VRFs. The RD is prepended to the IP address in the route advertisement, ensuring that routes from different tenants remain unique even if they use the same IP address range.
Correct Property:
D. the route distinguisher value: This is the correct answer because the RD is crucial in determining which routing information belongs to which VRF instance. It ensures that each VRF's routing table only contains relevant routes, maintaining isolation between tenants.
Data Center Reference:
The RD is a key element in MPLS and EVPN-based multi-tenant environments, ensuring proper routing segregation and isolation for different VRFs within the data center fabric.

QUESTION 3
Exhibit.

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

Answer: D
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 Reference:
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.

QUESTION 4
You are asked to automatically provision new Juniper Networks devices in your network with minimal manual intervention Before you begin, which two statements are correct? (Choose two.)

A. You must have a DHCP server that provides the location of the software image and configuration files.
B. You must have a system log (syslog) server to manage system log messages and alerts.
C. You must have an NTP server to perform time synchronization.
D. You must have a file server that stores software image and configuration files.

Answer: AD
Explanation:
Zero-Touch Provisioning (ZTP):
ZTP is a feature that allows for the automatic provisioning of devices with minimal manual intervention. It is widely used in large-scale deployments to quickly bring new devices online.
Key Requirements for ZTP:
A. DHCP Server: A DHCP server is crucial for ZTP as it provides the necessary information to new devices, such as the IP address, the location of the software image, and configuration files.
D. File Server: The file server is where the software image and configuration files are stored. The device downloads these files during the provisioning process.

QUESTION 5
You are selling up an EVPN-VXLAN architecture (or your new data center. this initial deployment will be less than 50 switches: however, it could scale up to 250 switches over time supporting 1024 VLANs. You are still deciding whether to use symmetric or asymmetric routing. In this scenario, which two statements are correct? (Choose two.)

A. Symmetric routing needs an extra VLAN with an IRB interface for each L3 VRF instance.
B. Asymmetric routing is easier lo monitor because of the transit VNI.
C. Symmetric routing supports higher scaling numbers.
D. Asymmetric routing routes traffic on the egress switch.

Answer: CD
Explanation:
Symmetric vs. Asymmetric Routing in EVPN-VXLAN:
Symmetric Routing: Traffic enters and exits the VXLAN network through the same VTEP, regardless of the source or destination. This approach simplifies routing decisions, especially in large networks, and is generally more scalable.
Asymmetric Routing: The routing occurs on the egress VTEP. This method can be simpler to deploy in smaller environments but becomes complex as the network scales, particularly with larger numbers of VNIs and VLANs.
Correct Statements:
C. Symmetric routing supports higher scaling numbers: Symmetric routing is preferred in larger EVPN-VXLAN deployments because it centralizes routing decisions, which can be more easily managed and scaled.

D. Asymmetric routing routes traffic on the egress switch: This is accurate, as asymmetric routing means the routing decision is made at the final hop, i.e., the egress VTEP before the traffic reaches its destination.

QUESTION 6
Your organization is implementing EVPN-VXLAN and requires multiple overlapping VLAN-IDs. You decide to use a routing-instance type mac-vrf to satisfy this request. Which two statements are correct in this scenario? (Choose two.)

A. Host-facing interfaces must be configured using a service-provider style configuration.
B. Host-facing interfaces must be configured using enterprise-style configuration.
C. Spine-facing interfaces must be configured using an enterprise-style configuration.
D. The routing-instance service type can be VLAN-based.

Answer: AD
Explanation:
Understanding the Scenario:
EVPN-VXLAN deployments often involve scenarios where multiple tenants or applications require overlapping VLAN IDs, which can be managed using the mac-vrf routing instance type. This allows you to segregate traffic within the same VLAN ID across different tenants.
Host-facing Interface Configuration:
A. Host-facing interfaces must be configured using a service-provider style configuration: This is correct. In mac-vrf configurations, host-facing interfaces (those connecting end devices) typically follow a service-provider style configuration, where each customer or tenant's traffic is isolated even if overlapping VLAN IDs are used.
B. Host-facing interfaces must be configured using enterprise-style configuration: This is incorrect for mac-vrf instances because enterprise-style configurations are more common in simpler, less segmented networks.
Routing Instance Service Type:
D. The routing-instance service type can be VLAN-based: This is correct. The service type in mac-vrf can indeed be VLAN-based, which is particularly useful in scenarios where VLAN ID overlap is needed between different tenants or services.
Data Center Reference:
The mac-vrf instance type is powerful for handling complex multi-tenant environments in EVPN-VXLAN, especially when dealing with overlapping VLAN IDs across different segments of the network.

Thank You for Trying Our Product

Braindump2go Certification Exam Features:

★ More than **99,900** Satisfied Customers Worldwide.

★ Average **99.9%** Success Rate.

★ **Free Update** to match latest and real exam scenarios.

★ **Instant Download** Access! No Setup required.

★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.

★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.

★ **100%** Guaranteed Success or **100%** Money Back Guarantee.

★ **Fast**, helpful support **24x7**.

View list of all certification exams: http://www.braindump2go.com/all-products.html

10% Discount Coupon Code: **ASTR14**