DUMPSARENA

Data Center Specialist (JNCIS-DC)

Juniper JN0-480

Version Demo

Total Demo Questions: 8

Total Premium Questions: 86

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Data Center Security	2
Topic 2, Data Center Fabric	30
Topic 3, Data Center Automation and Orchestration	54
Total	86

QUESTION NO: 1

Exhibit.

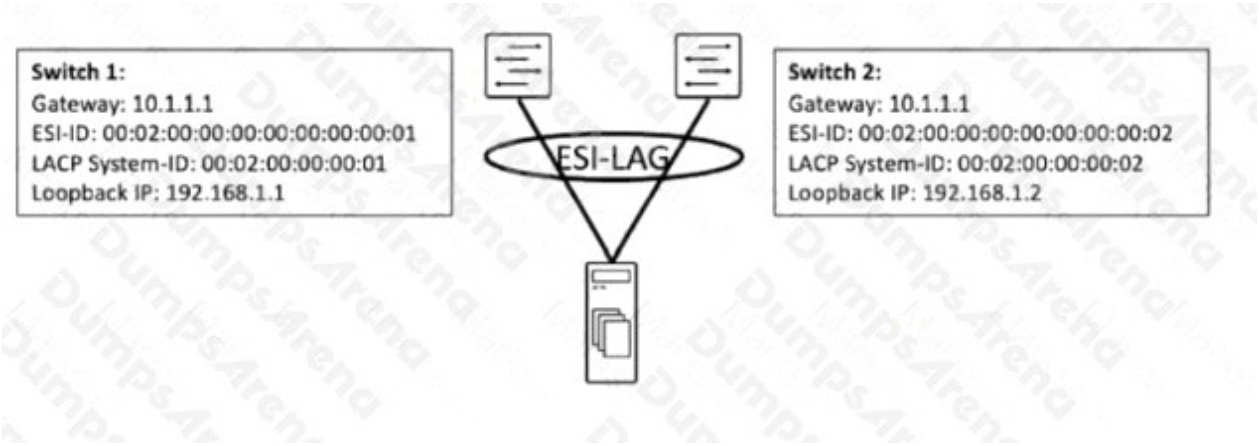

You are working to build an ESI-LAG for a multihomed server. The ESI-LAG is not coming up as multihomed.

Referring to the exhibit, what are two solutions to this problem? (Choose two.)

A. The gateway IP addresses on both devices must be different.

B. The LACP system ID on both devices must be the same.

C. The loopback IP addresses on both devices must be the same.

D. The ESI ID on both devices must be the same.

ANSWER: B D

Explanation:

An ESI-LAG presents links from two EVPN provider edge devices to a multihomed server as one logical LAG. Therefore, **The LACP system ID on both devices must be the same.** Sharing the LACP system ID ensures that the server recognizes both devices as the same LACP partner and can successfully form one LAG across its links to the two devices. This is a fundamental requirement for the server-facing bundle to operate as a multihomed LAG rather than as independent connections.

The ESI ID on both devices must be the same. The Ethernet Segment Identifier identifies the shared Ethernet segment to the EVPN control plane. A common ESI lets both devices advertise and coordinate reachability for the same multihomed segment, including the EVPN signaling used for designated-forwarder election and split-horizon behavior. With matching LACP system identity and ESI values, the data plane and EVPN control plane describe the same dual-homed attachment, allowing the ESI-LAG to become operational as multihomed. Juniper's EVPN documentation describes ESI-LAG multihoming and its shared Ethernet-segment identity requirements: EVPN-VXLAN LAG Multihoming and Ethernet Segment Identifiers in EVPN.

QUESTION NO: 2

What does EVPN use lo identity which remote leaf device advertised the EVPN route?

A. a route distinguisher value

B. a community tag

C. a route target value

D. a VRF target value

ANSWER: A

Explanation:

a route distinguisher value is correct. EVPN carries its reachability information in BGP EVPN Network Layer Reachability Information (NLRI), and the route distinguisher is an 8-byte value included with the EVPN route. The route distinguisher makes routes unique across EVPN instances and participating provider edge devices, even when overlapping customer addressing or identical MAC addresses are present. In a leaf-and-spine EVPN fabric, each leaf acting as a virtual tunnel endpoint originates EVPN routes using a unique route distinguisher, commonly derived from an identifier associated with that device. This uniqueness lets a receiving leaf distinguish otherwise equivalent EVPN advertisements and associate an advertised route with its originating remote leaf context. The route distinguisher is therefore part of the EVPN route identity rather than a policy-only label. RFC 7432 specifies that the route distinguisher must be unique per EVPN instance per provider edge device and describes its use in EVPN route encoding. Juniper's EVPN documentation also describes route distinguishers as values used to keep EVPN routes distinct. See RFC 7432, BGP MPLS-Based Ethernet VPN and Juniper EVPN Overview.

QUESTION NO: 3

What are two system-defined user roles that are available in Juniper Apstra? (Choose two.)

A. authorized

B. root

C. viewer

D. user

ANSWER: C D

Explanation:

viewer and **user** are system-defined roles in Juniper Apstra's role-based access-control model. The viewer role is intended for read-only access, allowing personnel such as operations observers or auditors to inspect Apstra objects and configuration information without making changes. This supports separation of duties while giving stakeholders visibility into the fabric and its managed resources.

The user role provides broader day-to-day access than a viewer role. It is intended for users who need to work with Apstra-managed objects and make permitted configuration edits, while not receiving the unrestricted administrative privileges associated with the administrator role. Apstra also includes system-defined roles such as administrator and device_ztp, but the two requested roles are viewer and user. Role assignment can be applied to locally managed accounts and used in conjunction with supported authentication and authorization workflows, enabling administrators to grant access appropriate to operational responsibilities. Juniper's Apstra documentation portal provides the User Guide and platform administration material, including user and role management information: Juniper Apstra Documentation. For product-level documentation and release-specific resources, see Juniper Documentation Portal.

QUESTION NO: 4

You are deploying a distributed anycast gateway for a virtual network that is extended across multiple leaf devices.

Which two statements are correct in this scenario? (Choose two.)

A. Each participating leaf uses the same gateway IP address for the virtual network.

B. Each participating leaf uses the same virtual gateway MAC address.

C. Each participating leaf uses the same VTEP loopback address.

D. Only one leaf can provide the gateway function for the virtual network.

ANSWER: A B

Explanation:

With a distributed anycast gateway, each participating leaf provides the same default-gateway IP address and virtual MAC address for the virtual network. A locally attached host can therefore use its local leaf as the first-hop gateway, avoiding unnecessary traffic traversal to another leaf for inter-VN routing. The VTEP address is not shared; each leaf requires its own unique VTEP address for VXLAN encapsulation.

QUESTION NO: 5

When working with logical devices, you specify where each port group is connected.

In this scenario, which two Juniper Apstra UI options are available to the operator? {Choose two.)

A. router

B. unused

C. generic

D. firewall

ANSWER: B C

Explanation:

In Juniper Apstra, a logical device defines the physical characteristics and port-group layout of a device model for use in a blueprint. While configuring the connectivity for a port group, the operator assigns a role that tells Apstra how that group should be handled in the fabric design. The available general-purpose choices include **unused** and **generic**. Selecting **unused** leaves the port group outside of Apstra's managed fabric connectivity, which is appropriate when ports are intentionally not participating in the design. Selecting **generic** identifies the port group as a general connection rather than assigning it a fabric-specific role. This supports port groups that need to be represented in the logical-device definition without being categorized for a particular topology function. These role selections are made as part of defining the logical device's port groups, alongside the port layout and supported speed information. Juniper's Apstra documentation describes logical devices as reusable device definitions that model interfaces and port groups used during blueprint design. See the Juniper Apstra logical devices documentation and the Juniper Apstra User Guide.

QUESTION NO: 6

In an EVPN-VXLAN bridged overlay, a new leaf device joins a virtual network and must inform the other VTEPs that it participates in BUM traffic replication for that virtual network.

Which EVPN route type is used for this purpose?

A. EVPN Type 1 Ethernet Auto-Discovery route

B. EVPN Type 2 MAC/IP Advertisement route

C. EVPN Type 3 Inclusive Multicast Ethernet Tag route

D. EVPN Type 5 IP Prefix route

ANSWER: C

Explanation:

An EVPN Type 3 Inclusive Multicast Ethernet Tag route advertises VTEP membership for a bridged EVPN instance. Remote VTEPs use this information to determine the VTEPs that should receive BUM traffic when ingress replication is used.

A Type 2 route advertises MAC and IP reachability, while a Type 5 route advertises IP prefixes. A Type 1 route provides Ethernet segment auto-discovery information and is not the route used to advertise inclusive multicast membership.

QUESTION NO: 7

In Juniper Apstra. which statement is correct?

A. VMware anomaly detection is on by default.

B. VMware anomaly detection requires a vCenter server configured under External Systems

C. VMware anomaly detection requires a VMware hypervisor with exports enabled.

D. VMware anomaly detection requires an Apstra server running on VMware.

ANSWER: B

Explanation:

VMware anomaly detection requires a vCenter server configured under External Systems. Apstra obtains the virtual-infrastructure information needed for VMware visibility and anomaly analysis by integrating with VMware vCenter. The vCenter integration provides inventory and relationship data for components such as ESXi hosts, virtual machines, distributed virtual switches, port groups, and host interfaces. Apstra can then correlate that VMware-side information with the physical connectivity and configuration it manages in the data-center blueprint.

After the vCenter server is defined as an external system, the virtual infrastructure can be associated with the relevant blueprint and committed. This enables Apstra to discover the VMware environment and validate expected relationships between the virtual and physical networks. The resulting service anomalies help operators identify conditions that could affect workload connectivity, including inconsistencies in attachment, VLAN-related configuration, or visibility between virtual and physical interfaces. The external-system configuration is therefore a prerequisite for Apstra to access the vCenter-managed environment and perform this VMware-specific analysis. Juniper documents the vCenter connection and virtual-infrastructure workflow in its VMware virtual infrastructure configuration guide, and describes related operational validation in the Apstra anomalies documentation.

QUESTION NO: 8

Exhibit.

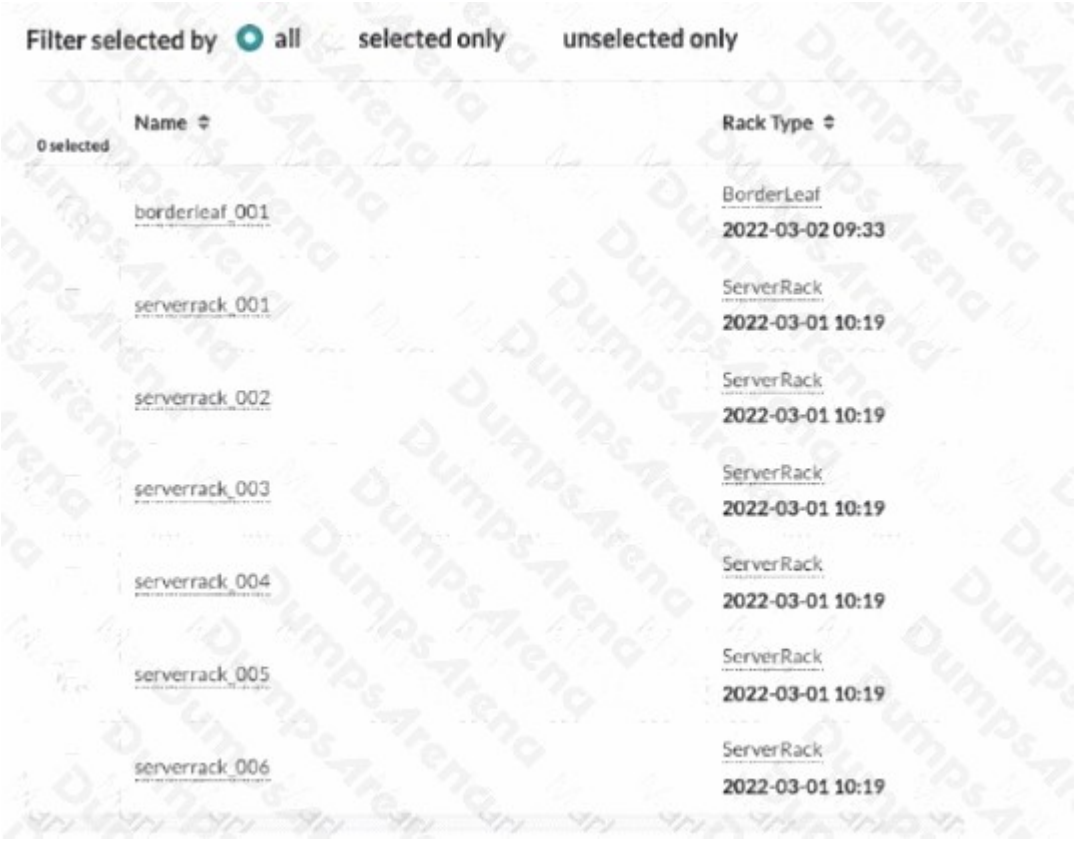

Referring to the exhibit, how many tack types ate used in the staged blueprint?

A. six

B. three

C. seven

D. two

ANSWER: D

Explanation:

The correct answer is **two**. In the staged blueprint's Racks table, the Rack Type column identifies the rack-type resource assigned to each individual rack. Although the exhibit contains seven rack entries, those entries use only two distinct rack-type names: **BorderLeaf** and **ServerRack**. Counting unique values in the Rack Type column, rather than counting total rack rows, gives a result of two rack types.

In Juniper Apstra, a rack type defines the device composition and logical role of a rack build. It can specify the leaf devices, access switches, and generic systems that make up a particular rack design. A staged blueprint can instantiate the same rack type multiple times, which is why multiple rack rows may share the ServerRack type while still representing only one distinct rack type. The BorderLeaf entry represents the other distinct type used by the blueprint. This distinction between rack instances and reusable rack-type definitions is central to interpreting the staged Racks view correctly.

For further detail, see Juniper's documentation on rack types and staged racks.