DUMPSBOSS.

Data Center Associate (JNCIA-DC)

Juniper JN0-281

Version Demo

Total Demo Questions: 8

Total Premium Questions: 81

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Data Center Architecture	15
Topic 2, Data Center Switching	25
Topic 3, Data Center Routing	40
Topic 4, Data Center Security	1
Total	81

QUESTION NO: 1

Which two statements about EBGP are correct? Choose two.

A. EBGP is used between ASs.

B. EBGP is used within a single AS.

C. EBGP does not require a supporting IGP.

D. EBGP requires a supporting IGP.

ANSWER: A C

Explanation:

EBGP is used between ASs because an external BGP session is, by definition, a peering relationship whose endpoints belong to different autonomous systems. The AS numbers exchanged in the BGP OPEN messages identify the external relationship, and Junos uses the configured peer AS to establish the session type. This model is particularly useful in a leaf-spine IP fabric, where distinct private AS numbers can be assigned to adjacent devices or to groups of devices. It creates explicit routing-policy and AS-path boundaries while allowing each leaf to peer directly with the spine layer.

EBGP does not require a supporting IGP because directly connected EBGP neighbors can establish a session using the connected interface addresses, and BGP can advertise the underlay reachability information needed by the fabric. An IGP may be used by a particular network design, especially when multihop BGP peers need recursive reachability, but it is not an inherent EBGP requirement. In a common fabric deployment, EBGP distributes loopback and infrastructure prefixes while the physical point-to-point links provide direct neighbor reachability. Juniper's BGP documentation describes external peers as belonging to different ASs and documents BGP neighbor configuration and operation: Juniper BGP Overview and Juniper Configuring BGP.

QUESTION NO: 2

You want to advertise the network configured on an interface as an OSPF internal route. You do not want this interface to participate in your OSPF domain. Which OSPF parameter must you use in this scenario?

A. passive

B. interface-type

C. strict-bfd

D. link-protection

ANSWER: A

Explanation:

The `passive` parameter is correct because it enables OSPF to advertise the directly connected network of the specified interface while preventing the interface from actively taking part in OSPF neighbor formation. In Junos, an interface configured as passive is included in the OSPF area configuration, so its connected prefix is originated into OSPF as an intra-area route. However, Junos does not send OSPF Hello packets from that interface and does not establish OSPF adjacencies through it. This is the appropriate design for a network that must be reachable through OSPF but connects only to hosts, servers, management devices, or any segment where no OSPF router neighbor should exist. The configuration is applied under the relevant OSPF area and interface hierarchy, for example, by setting the interface as passive within the area. This preserves route advertisement without exposing the interface to OSPF peering activity. Juniper documents the `passive` statement as preventing OSPF packet transmission on the interface while still advertising the interface address into OSPF. See the Juniper passive statement reference and the OSPF configuration documentation.

QUESTION NO: 3

Which statement is correct about building an IP fabric?

A. Each spine device should have a direct physical connection to every other spine device.

B. Each spine device should have a direct physical connection to every leaf device.

C. Each spine device must have two or more physical connections to each leaf device.

D. Each spine device must have two or more physical connections to every other spine device.

ANSWER: B

Explanation:

Each spine device should have a direct physical connection to every leaf device. This is the defining physical topology of a conventional two-tier leaf-spine IP fabric: leaf switches provide attachment points for servers and other endpoints, while spine switches form the high-bandwidth transit layer. Every leaf connects to every spine, creating multiple equal-cost paths between leaf switches. Routing protocols can advertise these paths and use equal-cost multipath forwarding, allowing traffic to use the available fabric capacity while providing predictable, low-latency connectivity. This arrangement also supports scalable growth: additional leaf switches are connected to all existing spines, and additional spine switches are connected to all leaves. The result is a nonblocking or consistently oversubscribed fabric design, depending on interface speeds and port counts, rather than a hierarchy in which traffic must traverse inter-spine links. This leaf-to-every-spine connectivity model is consistent with the RFC description of a Clos-based data-center fabric and with Juniper's data-center fabric design guidance. See [RFC 7938: Use of BGP for Routing in Large-Scale Data Centers]() and [Juniper Technical Documentation]().

QUESTION NO: 4

What are two available modes when using LACP with an aggregated Ethernet bundle? Choose two.

A. aggressive

B. mixed

C. passive

D. active

ANSWER: C D

Explanation:

The available LACP modes for a Junos aggregated Ethernet interface are **passive** and **active**. LACP is the IEEE 802.1AX protocol used to dynamically establish and maintain member links in an aggregated Ethernet (AE) bundle. With **active** mode, the interface actively transmits LACP protocol data units to initiate negotiation with its directly connected peer. This mode is commonly configured on at least one end of an LACP bundle to ensure that LACP negotiation starts. With **passive** mode, the interface listens for LACP packets and responds after a peer initiates negotiation. Therefore, passive mode interoperates successfully with an active peer, while the overall bundle formation depends on LACP exchanges being initiated by at least one side. In Junos, this behavior is configured under the `aggregated-ether-options lacp` hierarchy for the AE interface. Junos documentation identifies `active` and `passive` as the LACP mode values, and separately supports LACP timing configuration, such as periodic fast or slow transmission. See the Juniper [aggregated-ether-options reference]() and the [Aggregated Ethernet and LACP documentation]().

QUESTION NO: 5

Which feature should be used with a static route that has a secondary next hop with a unique route preference value? Choose one.

A. retain

B. resolve

C. qualified next hop

D. install

ANSWER: C

Explanation:

qualified next hop is the appropriate Junos static-routing feature when a route needs a secondary next hop that uses its own, distinct route preference. A static route can have a primary `next-hop` and one or more `qualified-next-hop` entries. The qualified next hop is associated with the same destination prefix, but it can be configured with an independent preference value. This lets the routing table favor the primary next hop under normal conditions and use the qualified next hop as a backup when its configured preference makes it the best available active path. Junos evaluates route preference with lower numeric values being more preferred, so the primary and backup behavior can be controlled explicitly through these values. The feature is useful for resilient routing designs because it keeps alternate forwarding information attached to the same static destination while allowing the alternate path to be selected only when appropriate. Juniper documents `qualified-next-hop` as a static-route configuration statement and specifically supports a `preference` setting beneath it. See the Junos qualified-next-hop statement reference and the Junos static routes configuration documentation.

QUESTION NO: 6

Which state in the adjacency process do OSPF routers check the MTU size?

A. Init

B. Exchange

C. Done

D. ExStart

ANSWER: D

Explanation:

ExStart is the correct state. After reaching the two-way state, OSPF neighbors enter ExStart to negotiate the master/slave relationship and establish the initial Database Description packet sequence number. Database Description packets include an Interface MTU field. When a router receives one of these packets, it compares the advertised MTU with the size it can accept on the receiving interface without fragmentation. If the peer advertises an MTU that is too large, the packet is rejected, preventing the adjacency from advancing normally.

This check is therefore associated with the Database Description packet exchange initiated during ExStart, even though an operational symptom is often described as a neighbor being stuck in ExStart or failing to complete the subsequent database synchronization. Matching Layer 3 interface MTUs on OSPF neighbors is the normal best practice; Junos also provides an option to disable MTU checking only when there is a deliberate interoperability or migration requirement. The OSPF specification defines the Interface MTU field and the rejection behavior for an unacceptable value in the Database Description packet processing rules. See RFC 2328, Section 10.6 and Juniper's OSPF mtu-ignore statement documentation.

QUESTION NO: 7

What are two challenges of traditional data center architectures? Choose two.

A. inefficient resource usage

B. inefficient network path selection

C. isolated user devices

D. increased latency

ANSWER: A B

Explanation:

Traditional data center designs commonly use a hierarchical access, aggregation, and core topology. In these environments, Layer 2 redundancy is often controlled with Spanning Tree Protocol, which places some physically available links into a blocking state. This produces **inefficient resource usage**, because purchased bandwidth, ports, and redundant paths cannot all forward traffic simultaneously during normal operation. Capacity planning must therefore account for links that are present primarily for failover rather than active load sharing.

Inefficient network path selection is also a core concern in a traditional hierarchical design. East-west traffic between workloads attached to different access switches may need to travel upward through aggregation or core layers before returning downward to its destination. Such topology-driven forwarding can create indirect paths and unnecessary hops rather than using a predictable, direct fabric path. Modern IP leaf-spine fabrics address both concerns by using Layer 3 multipathing, allowing parallel links to forward traffic and providing consistent paths between leaf switches. Juniper describes these fabric characteristics at Juniper: What Is a Data Center Fabric?; the routed Clos-fabric principles are also documented in RFC 7938.

QUESTION NO: 8

What are two advantages of spine and leaf architectures compared to legacy architectures within a data center? Choose two.

A. Containers and hypervisors are easier to manage with spine and leaf architectures.

B. Spine and leaf fabrics demonstrate better scaling features.

C. Better performance is due to low and predictable hop counts.

D. All broadcast domains become reachable, allowing easier host connectivity.

ANSWER: B C

Explanation:

Spine and leaf fabrics demonstrate better scaling features because they use a repeatable, modular two-tier topology. Each leaf switch connects to every spine switch, so capacity can be expanded incrementally: adding leaf switches increases endpoint connectivity, while adding spine switches increases available fabric bandwidth. This model supports horizontal growth without redesigning the network hierarchy, making it well suited to the rapid growth of virtualized and cloud data centers. The parallel leaf-to-spine links also enable equal-cost multipath forwarding, allowing traffic to use multiple available paths efficiently.

Better performance is due to low and predictable hop counts because traffic between devices attached to different leaf switches follows a consistent path through a spine switch. This produces predictable latency and avoids the variable path lengths often found in traditional hierarchical designs. Such consistency is particularly valuable for east-west data center traffic between application tiers, virtual machines, and distributed services. Juniper describes leaf-spine as a scalable architecture that delivers predictable performance through its consistent network structure. See Juniper: What Is a Leaf-Spine Architecture? and Juniper Documentation: Understanding Spine-Leaf Topology.