

DUMPS ARENA

Data Center, Associate (JNCIA-DC)

Juniper JN0-280

Version Demo

Total Demo Questions: 8

Total Premium Questions: 81

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 Networking Fundamentals	4
Topic 2, Data Center Switching and Routing	60
Topic 3, Data Center Fabric	14
Topic 4, Data Center Security	3
Total	81

QUESTION NO: 1

What is the primary purpose of an IRB Layer 3 interface?

- A. to provide load balancing
- B. to provide a default VLAN ID
- C. to provide inter-VLAN routing
- D. to provide port security


ANSWER: C

Explanation:

The primary purpose of an IRB Layer 3 interface is **to provide inter-VLAN routing**. IRB, meaning Integrated Routing and Bridging, is a logical Layer 3 interface that is associated with a VLAN or bridge domain. It gives that Layer 2 broadcast domain an IP interface, which can act as the default gateway for hosts in the VLAN. Once separate VLANs or bridge domains have their own IRB interfaces and IP subnets, the Junos routing table can route traffic between them. For example, a host in one VLAN sends traffic destined for a different IP subnet to its configured IRB gateway; the switch then performs the Layer 3 forwarding decision and sends the packet toward the destination VLAN or another routed network. This integrated model allows a Juniper switch to retain Layer 2 switching within each VLAN while also delivering gateway and routing services at Layer 3. Juniper documentation describes configuring an IRB interface as the Layer 3 interface for a VLAN, enabling communication beyond the VLAN's local Layer 2 segment. See Juniper's [IRB and bridging documentation](#) and its [IRB interface configuration documentation](#).

QUESTION NO: 2

Exhibit:



R2 received an OSPF update from R1, and it received the same update from R3.

Referring to the exhibit, what will R2 do?

- A. R2 ignores the update from R1.
- B. R2 does nothing with R3's update.
- C. R2 ignores the update from R3.

D. R2 acknowledges R3 and discards it.

ANSWER: D

Explanation:

R2 acknowledges R3 and discards it. OSPF routers maintain a link-state database containing the current instance of each link-state advertisement (LSA). An LSA instance is identified and compared using its LSA type, link-state ID, advertising router, sequence number, checksum, and age. After R2 has accepted the LSA received from R1, the matching LSA received from R3 represents an identical instance already present in R2's link-state database. R2 therefore does not install it again or run another SPF calculation simply because it arrived through a second neighbor.

OSPF reliability still requires R2 to acknowledge the LSA received from R3. The acknowledgement confirms to R3 that R2 has received that LSA, allowing R3 to clear it from its retransmission list and preventing needless retransmissions. Thus, the duplicate LSA is discarded after receipt processing while its arrival is acknowledged. This behavior is defined in the OSPF LSA-receipt procedure in [RFC 2328, Section 13](#), which specifies that an identical LSA is discarded and acknowledged when appropriate. Junos implements standard OSPF link-state flooding and database synchronization behavior; see Juniper's [OSPF overview documentation](#).

QUESTION NO: 3

What are three correct layer names used in legacy hierarchical network design? (Choose three.)

- A. Access layer
- B. Modular layer
- C. Aggregation layer
- D. Core layer
- E. Function layer

ANSWER: A C D

Explanation:

Legacy hierarchical network design separates network responsibilities into the access, aggregation, and core layers. The **Access layer** provides the point at which servers, endpoints, and other edge devices attach to the network. It focuses on reliable device connectivity and can enforce edge-facing network policies. The **Aggregation layer**, also commonly called the distribution layer, consolidates connectivity from multiple access switches. It is the policy and services boundary in the traditional model, where designs commonly apply functions such as route summarization, traffic control, and redundant uplinks between the access and core portions of the network. The **Core layer** is the high-capacity transport portion of the hierarchy. Its purpose is to move traffic quickly and reliably among aggregation blocks and toward shared network services, while maintaining resilient paths. This layered approach makes a conventional network easier to scale because device attachment, aggregation/policy, and high-speed backbone transport have distinct roles. Juniper describes modern data-center architectures and their evolution from conventional hierarchical designs in its [Data Center Networking solutions overview](#); its [technical documentation portal](#) provides platform and design documentation supporting these network roles.

QUESTION NO: 4

Which signaling protocol is used for EVPN?

- A. OSPF
- B. PIM
- C. IS-IS
- D. BGP

ANSWER: D

Explanation:

BGP is the signaling protocol used by EVPN. EVPN defines a BGP address family, BGP EVPN, through which provider edge devices advertise and receive EVPN Network Layer Reachability Information. This control-plane information includes MAC/IP reachability, Ethernet segment discovery, inclusive multicast Ethernet-tag routes, and IP prefix routes, depending on the EVPN service design. Using BGP gives EVPN a standards-based and scalable method to distribute endpoint and multihoming information between VTEPs across either an MPLS or VXLAN data-center fabric. In a Junos EVPN-VXLAN deployment, BGP EVPN operates as the overlay control plane: it distributes MAC and IP bindings and VXLAN-related reachability information, while the underlay routing protocol provides IP transport between VTEPs. Juniper documentation describes configuring BGP EVPN peers to exchange EVPN routes and identifies BGP as the EVPN control-plane protocol. See Juniper's [EVPN Overview](#) and the IETF's [RFC 7432: BGP MPLS-Based Ethernet VPN](#).

QUESTION NO: 5

A router has two usable routes to 192.0.2.0/24: a static route with preference 200 and an OSPF route with preference 10. Which two statements are correct? (Choose two.)

- A. The OSPF route is selected as the active route.
- B. The static route remains available as a backup route.
- C. The static route is selected because static routes are always preferred.
- D. The two routes are installed as equal-cost multipath routes.

ANSWER: A B

Explanation:

Junos selects the OSPF route as active because a lower route preference value is more preferred. The OSPF route has preference 10, which is lower than the static route preference of 200. The static route remains as an eligible backup route while its next hop remains resolvable.

Route preference is evaluated before protocol-specific metrics from different protocols can be compared. The static route does not win merely because it is statically configured; its explicitly configured preference makes it less preferred than the OSPF route in this scenario.

QUESTION NO: 6

By default, which two statements are correct about BGP advertisements? (Choose two.)

- A. BGP peers advertise routes received from EBGp peers to other IBGP peers.
- B. BGP peers advertise routes received from IBGP peers to other IBGP peers.
- C. BGP peers advertise routes from EBGp peers to other IBGP peers using its own address as the next hop.
- D. BGP peers advertise routes from IBGP peers to EBGp peers using its own address as the next hop.

ANSWER: A D

Explanation:

BGP advertises routes learned from an external BGP neighbor to internal BGP peers. This behavior distributes externally learned reachability throughout the local autonomous system, allowing internal BGP-speaking routers to select and use those destinations. When an externally learned route is advertised through an internal BGP session, its original BGP next-hop attribute is normally preserved; network designs commonly ensure that internal routers can resolve that next hop or explicitly apply a next-hop policy where needed.

BGP also advertises routes learned from an internal BGP peer to external BGP peers. For an advertisement sent over an external BGP session, the advertising router sets the NEXT_HOP attribute to its own interface address used for that external peer relationship. This gives the external neighbor a directly reachable next hop for forwarding toward the advertised destination. These default advertisement behaviors are fundamental to using BGP at an autonomous-system boundary while propagating externally learned routes within the AS. The NEXT_HOP processing rules are defined in [RFC 4271, section 5.1.3](#); Junos BGP configuration and operational concepts are documented in the [Juniper BGP configuration guide](#).

QUESTION NO: 7

You want to minimize topology disruptions in your network when the rpd process restarts on a device. Which service would accomplish this task?

- A. Bidirectional Forwarding Detection (BFD)
- B. link aggregation groups
- C. graceful restart (GR)
- D. Virtual Chassis

ANSWER: C

Explanation:

graceful restart (GR) is the service designed to minimize topology disruption when the Junos routing protocol daemon (`rpdd`) restarts. Normally, `rpdd` maintains routing-protocol adjacencies and installs routes into the forwarding plane. A restart could therefore cause neighboring devices to detect lost protocol sessions, withdraw routes, and recalculate the topology. With graceful restart, the device preserves forwarding-state information while the control plane recovers. The forwarding plane can continue using routes that were previously installed, avoiding an immediate traffic interruption while `rpdd` restarts and routing adjacencies are reestablished.

Graceful restart also uses protocol support from neighboring devices, often called helper behavior. Neighbors retain the restarting router's routes for a limited restart interval rather than immediately treating them as unavailable. This prevents unnecessary routing churn and limits reconvergence across the network. In Junos, graceful restart is an important high-availability capability for planned maintenance and certain routing-process recovery events, because it separates temporary control-plane recovery from the ongoing forwarding of traffic. Juniper documents graceful restart as a mechanism that enables uninterrupted packet forwarding during routing-engine or routing-protocol restart scenarios. See the [Juniper Graceful Restart documentation](#) and the [Junos routing protocol graceful restart overview](#).

QUESTION NO: 8

An aggregated Ethernet interface has two active member links and is carrying traffic between a leaf device and a spine device.

Which two statements are correct if one member link fails while the other member remains operational? (Choose two.)

- A. The aggregated Ethernet interface can remain operational.
- B. Traffic can continue through the remaining active member link.
- C. The remaining member link must be administratively disabled.
- D. The aggregate must create a separate logical interface for the remaining member.

ANSWER: A B

Explanation:

An aggregated Ethernet interface represents multiple physical member links as one logical interface. If at least one eligible member remains active, the aggregate can remain operational rather than causing the entire logical connection to fail.

Traffic that can use the remaining active member continues to be forwarded through the aggregate. Per-flow load balancing normally selects a member link by using a forwarding hash, so flows that had been using the failed member are redirected to an available member after the failure is recognized.