

DUMPSBOSS.

Service Provider Routing and Switching Specialist (JNCIS-SP)

Juniper JN0-364

Version Demo

Total Demo Questions: 8

Total Premium Questions: 82

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, OSPF	12
Topic 2, IS-IS	12
Topic 3, BGP	19
Topic 4, Routing Policy	3
Topic 5, MPLS	17
Topic 6, Layer 2 VPNs	2
Topic 7, Layer 3 VPNs	2
Topic 8, IPv6	4
Topic 9, Mix Questions	11
Total	82

QUESTION NO: 1

Which two events cause a static route to be removed from a routing table? (Choose two.)

- A. The route is manually removed.
- B. The outbound interface becomes unavailable.
- C. The route has no traffic for 30 days.
- D. Hosts two hops away become unreachable.

ANSWER: A B

Explanation:

A static route is present because it is explicitly configured under the routing options hierarchy. Therefore, when an administrator deletes the static-route configuration and commits the change, Junos removes that route from the relevant routing table because the route no longer has a configured source. Juniper documents static routes as routes configured manually using the `routing-options static` hierarchy; see the [Junos static-route configuration documentation](#).

The outbound interface becoming unavailable can also cause a static route using that interface, or a next hop dependent on it, to lose a usable next hop. Junos route selection requires a resolvable next hop for a route to be active. If the directly connected path is no longer available, the static route is no longer an active route eligible for installation in the forwarding table and is withdrawn from the active routing-table entry. This behavior follows Junos next-hop resolution and route state processing, described in the [route resolution documentation](#).

Accordingly, the applicable events are manual removal of the configured route and loss of the outbound interface needed to reach its next hop.

QUESTION NO: 2

You are configuring BGP on a Juniper router to peer with an external provider. After committing the configuration, the BGP session remains in the Idle state. Which configuration issue would prevent the BGP session from progressing beyond the Idle state?

- A. The peer IP address is unreachable.
- B. The local AS number is higher than the peer's AS number.
- C. The peer is configured with a different router ID.
- D. The BGP group type is set to internal instead of external.

ANSWER: A

Explanation:

The peer IP address is unreachable. is correct because BGP must first establish a TCP connection to its configured neighbor before it can exchange BGP OPEN messages and advance through the BGP finite-state machine. On Junos, the routing table must contain an active route that resolves the configured peer address, and the corresponding forwarding path must permit connectivity to TCP port 179. If the router has no usable route to that address, it cannot send the initial TCP SYN packet to the neighbor. As a result, the BGP process cannot establish the underlying transport session and remains at, or repeatedly returns to, Idle while it waits for the conditions needed to initiate a connection.

For an external BGP peer, validating IP reachability is therefore an essential first troubleshooting step. Administrators should verify that the peer address has an active route in the relevant routing instance, confirm that the selected source address is reachable in both directions, and ensure that filters or security policy do not block TCP/179. Juniper's BGP troubleshooting guidance describes checking peer reachability and the BGP state, while Juniper's BGP overview explains that BGP uses TCP as its reliable transport protocol. See [Juniper BGP troubleshooting documentation](#) and [Juniper BGP overview documentation](#).

QUESTION NO: 3

Referring to the exhibit, which protocol would automatically create a full mesh of label-switched paths between MPLS-enabled routers?

- A. LDP
- B. BFD
- C. BGP
- D. RSVP

ANSWER: A

Explanation:

LDP is correct because it distributes MPLS label bindings automatically for IGP-reachable routes. After LDP is enabled on the MPLS core interfaces, adjacent routers discover one another and establish LDP sessions. Each participating router advertises labels for its local and learned forwarding-equivalence classes, while the IGP supplies the underlying next-hop reachability and shortest paths. Consequently, an ingress router can build an LDP-signaled LSP to every destination router loopback or other advertised prefix in the MPLS domain. When all core routers participate, this produces the transport LSP connectivity commonly described as a full mesh, without individually configuring an ingress-to-egress tunnel for every router pair.

In Junos, LDP is typically deployed with an IGP such as IS-IS or OSPF and MPLS enabled on the same core interfaces. The IGP determines the route toward an egress, and LDP supplies the label operations needed to forward traffic along that route. This automatic, topology-driven behavior makes LDP suitable as the MPLS transport foundation for services such as Layer 3 VPNs. Juniper's [LDP overview](#) describes LDP as the protocol that distributes labels for MPLS forwarding, and its [MPLS configuration documentation](#) explains its use with core routing protocols.

QUESTION NO: 4

You are asked to configure a new network environment that will be based on IPv6 and use OSP

F.

In this scenario, which two statements correctly identify configuration task considerations? (Choose two.)

- A. Participating interfaces must be configured with both IPv4 and IPv6 protocol families and addresses.
- B. The router ID used must be based on a 128-bit identifier value.
- C. The router ID used must be based on a 32-bit identifier value.
- D. Participating interfaces are only required to be configured with the IPv6 protocol family and address.

ANSWER: C D

Explanation:

OSPFv3 retains the OSPF router ID as a 32-bit value, represented in dotted-decimal notation, even though OSPFv3 is used to exchange IPv6 routing information. The router ID is a protocol identifier rather than an IPv6 address, so it is not derived from a 128-bit IPv6 value. In an IPv6-only Junos deployment, configure an explicit router ID under `routing-options` when one cannot be selected from an IPv4 interface address. This ensures the OSPFv3 process has the identifier required to originate LSAs and form a functioning routing domain.

OSPFv3 is enabled on interfaces configured for the IPv6 address family, `family inet6`. Neighbor discovery and OSPFv3 adjacency establishment use IPv6 link-local addresses, and OSPFv3 packets are sent to IPv6 link-local multicast addresses. Consequently, an interface does not need an IPv4 protocol family or an IPv4 address merely to participate in IPv6 OSPFv3. This supports an IPv6-only routing environment while preserving the protocol's required local IPv6 operation. Juniper's OSPFv3 documentation describes configuring OSPFv3 under the IPv6 protocol family, and its `routing-options` reference documents the router-ID setting: [Juniper OSPFv3 configuration documentation](#) and [Juniper router-id CLI reference](#).

QUESTION NO: 5

What happens if an IS-IS router receives a link-state PDU with a higher sequence number than the one in its database?

- A. It ignores the link-state PDU.
- B. It updates its database and floods the link-state PDU.
- C. It sends a CSNP to request confirmation from the source of the link-state PDU.
- D. It resets the adjacency with the source of the link-state PDU.

ANSWER: B

Explanation:

It updates its database and floods the link-state PDU. IS-IS maintains a link-state database containing the most current Link-State PDU (LSP) originated by each system. An LSP's sequence number identifies its version: an originating router increments the sequence number whenever it generates an updated instance of its LSP. Therefore, a received LSP with a higher sequence number is newer than the locally stored instance and must be accepted. The receiving router installs that newer LSP in its link-state database, allowing its shortest-path-first calculation to use the current topology and reachability information. It also floods the accepted LSP to eligible adjacent neighbors, subject to normal IS-IS flooding rules, so that all routers in the level and area can converge on the same database contents. Reliable flooding uses acknowledgments and retransmission mechanisms to ensure propagation. This sequence-number comparison is a core part of IS-IS database synchronization and prevents older topology information from replacing a newer version. Juniper's IS-IS protocol overview describes LSP flooding and the link-state database behavior: [Juniper IS-IS Protocol Overview](#). The IS-IS specification defines LSP sequence numbers and the processing of newer LSPs in its update process: [RFC 1195](#).

QUESTION NO: 6

You must advertise a summary route for 10.40.0.0/16 when either of two contributing /17 routes is active. The summary must remain active if one contributing route fails, but it should be withdrawn if both contributing routes are unavailable. Which Junos route type satisfies this requirement?

- A. an aggregate route
- B. a generated route
- C. a static route
- D. a discard route

ANSWER: A

Explanation:

An aggregate route satisfies this requirement. Junos installs an aggregate route when at least one contributing route is active. Therefore, the 10.40.0.0/16 aggregate can remain present while either 10.40.0.0/17 or 10.40.128.0/17 is available, and it is withdrawn when no contributing routes remain.

A generated route is unsuitable because it requires all configured contributing routes to be active before it is installed. A static route does not track the presence of the component routes in this manner, and a discard route only defines forwarding behavior rather than creating conditional route aggregation.

QUESTION NO: 7

Which two protocols would be used for dynamic routing in IPv6 environments? (Choose two.)

- A. IGMP

- B. IS-IS
- C. OSPFv2
- D. BGP

ANSWER: B D

Explanation:

IS-IS and BGP are both valid dynamic-routing protocols for IPv6 in Junos OS. IS-IS was extended for IPv6 by adding IPv6-specific Type-Length-Value elements, allowing one IS-IS routing process to advertise and calculate reachability for both IPv4 and IPv6 prefixes. This makes IS-IS particularly useful in dual-stack service-provider networks, where a common link-state protocol can support the two address families without operating separate routing processes. Junos supports IPv6 operation with IS-IS through the `family inet6` configuration hierarchy and IPv6 interface addressing.

BGP supports IPv6 through Multiprotocol BGP extensions. These extensions use address-family capabilities to exchange IPv6 Network Layer Reachability Information over a BGP peering session. In Junos, IPv6 unicast BGP is enabled using the `family inet6 unicast` hierarchy. BGP can therefore distribute IPv6 Internet routes as well as IPv6 VPN routes when the applicable address family is configured. Together, IS-IS and BGP provide commonly deployed interior and exterior routing capabilities for IPv6 service-provider designs. Juniper documents IPv6 support for IS-IS and BGP in its routing protocol documentation: [IS-IS for IPv6](#) and [BGP for IPv6](#).

QUESTION NO: 8

You are configuring BGP for IPv6 operations. In this scenario, which two statements are correct? (Choose two.)

- A. The Autonomous System Number (ASN) must be a 64-bit value.
- B. The router ID uses a 128-bit identifier value.
- C. The router ID uses a 32-bit identifier value.
- D. The Autonomous System Number (ASN) can be either a 16-bit or 32-bit value.

ANSWER: C D

Explanation:

For IPv6 routing, Multiprotocol BGP extends BGP by carrying IPv6 Network Layer Reachability Information through multiprotocol extensions; it does not change the fundamental size or purpose of the BGP router ID. Therefore, "The router ID uses a 32-bit identifier value" is correct. BGP identifies a speaker with a four-octet (32-bit) BGP Identifier, conventionally displayed in IPv4 dotted-decimal notation. On Junos devices, explicitly configuring a router ID is particularly important when no suitable IPv4 address is available, such as in an IPv6-only deployment. The identifier remains distinct from the IPv6 addresses advertised and used for BGP peering.

"The Autonomous System Number (ASN) can be either a 16-bit or 32-bit value" is also correct. BGP originally used two-octet AS numbers, and the protocol was extended to support four-octet AS numbers for the expanded ASN space. This ASN capability is independent of whether the address family being exchanged is IPv4 or IPv6; IPv6 MP-BGP sessions can use either legacy two-octet or modern four-octet ASNs, subject to peer capability and configuration. The IETF defines the BGP Identifier as a four-octet value in [RFC 4271](#) and specifies four-octet AS number support in [RFC 6793](#).