

Juniper JN0-364


Service Provider Routing and Switching Certification

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/jn0-364>

Latest Version: 6.0

Question: 1

You are troubleshooting a Level 1 IS-IS router that has an adjacency with a Level 1/2 router. Which two statements are correct in this scenario? (Choose two.)

- A. The Level 1/2 router merges Level 1 and Level 2 into one complete topology.
- B. The Level 1 router will learn the full topology of the Level 2 network.
- C. The Level 1/2 router sees the Level 1 network and the Level 2 network as two separate topologies.
- D. The Level 1 router will only learn the topology of the Level 1 network.

Answer: C, D

Explanation:

In the context of Juniper Networks Junos OS and the IS-IS (Intermediate System to Intermediate System) protocol, understanding the hierarchical relationship between router levels is critical for effective troubleshooting and design. IS-IS uses a two-level hierarchy to manage scalability: Level 1 (L1), which represents intra-area routing, and Level 2 (L2), which represents inter-area backbone routing.

When a router is configured as a Level 1/2 (L1/L2) device, it acts as a bridge between the two levels. According to Juniper technical documentation, an L1/L2 router maintains two completely separate Link-State Databases (LSDB)—one for Level 1 and one for Level 2. It does not merge these into a single topology. This separation ensures that local area topology changes (L1) do not necessarily flood into the backbone (L2) unless specific redistribution is configured, and vice versa. Therefore, statement C is correct because the L1/L2 router maintains distinct SPF (Shortest Path First) computations for each level.

Regarding the visibility of the Level 1 router, IS-IS is designed to keep L1 areas "stubby" by default. A Level 1 router only possesses the topology information for its own area (the Level 1 LSDB). It does not receive specific L2 routes or the L2 topology. Instead, the L1/L2 router sets the Attached (ATT) bit in its L1 Link-State PDUs (LSPs) to signal to L1-only routers that it has a connection to the backbone. The L1 router then generates a default route pointing to the L1/L2 router to reach inter-area destinations. This confirms that statement D is correct: the L1 router's knowledge is limited to its local L1 topology.

Conversely, statements A and B are incorrect because merging topologies would violate the hierarchical scaling principles of IS-IS, and L1 routers never learn the full L2 topology without explicit, non-standard route leaking.

Question: 2

By default, which routing table contains a list of all ingress LSPs?

- A. inet.2
- B. inet.3
- C. inet.1
- D. inet.0

Answer: B

Explanation:

In the Juniper Networks Junos operating system, the management of routing information is partitioned into several distinct routing tables (RIBs), each serving a specific architectural purpose. When dealing with Multiprotocol Label Switching (MPLS), understanding the distinction between inet.0 and inet.3 is fundamental for troubleshooting and traffic engineering. The inet.3 routing table is specifically designed to store the egress IPv4 addresses of Label-Switched Paths (LSPs). When an ingress router successfully establishes an LSP (via RSVP or LDP), it places the host address of the egress router (the tail-end) into the inet.3 table. This table is not used for general packet forwarding; instead, it is primarily used by the Border Gateway Protocol (BGP) for next-hop resolution. When BGP receives a route, it checks both inet.0 and inet.3 to resolve the next hop. If a matching entry exists in inet.3, the router knows it can reach that destination via an MPLS tunnel, allowing for the encapsulation of BGP traffic within MPLS. In contrast, inet.0 is the default unicast routing table used for standard IPv4 forwarding and contains routes learned via IGPs (OSPF, IS-IS) or static routing. inet.1 is utilized for multicast forwarding (MBGP), and inet.2 is typically used for Multicast Source Discovery Protocol (MSDP) or RPF checks in multicast environments. By isolating LSP egress points in inet.3, Junos prevents MPLS-specific paths from interfering with standard IGP path selection unless the administrator explicitly chooses to merge them (e.g., using the traffic-engineering bgp-igp command). Therefore, by default, the ingress router maintains its list of reachable LSP endpoints in inet.3.

Question: 3

What are two types of BGP messages exchanged while in the Established state? (Choose two.)

- A. open
- B. request
- C. update
- D. notification

Answer: C, D

Explanation:

In the Border Gateway Protocol (BGP) finite state machine (FSM), the Established state is the final and functional stage of a BGP peering session. According to Juniper Networks technical documentation, once a session reaches this state, the two peers have successfully exchanged

Open messages and agreed upon session parameters (such as AS numbers, hold timers, and BGP identifiers). Only after the session is "Established" can the routers begin the actual exchange of network layer reachability information (NLRI).

The most frequent message type exchanged in the Established state is the UPDATE message. These messages are the heart of BGP operations; they are used to advertise new feasible routes to a peer or to withdraw routes that are no longer reachable. An UPDATE message contains path attributes (like AS-Path, Next-Hop, and Local Preference) and the associated prefixes. In a stable network, UPDATE messages are only sent when there is a change in the topology, adhering to BGP's incremental update philosophy.

The second message type that can be exchanged in this state is the NOTIFICATION message. While ideally, a session stays established, any detected error—such as a hold timer expiration, a malformed update, or a manual "clear" command—will trigger the transmission of a NOTIFICATION message. This message informs the peer of the specific error code and immediately causes the BGP session to transition back to the Idle state, tearing down the TCP connection.

It is important to note that OPEN messages (Option A) are only used during the session initialization phase to transition from the OpenConfirm state to Established. REQUEST (Option B) is not a valid BGP message type defined in the standard (RFC 4271); the closest equivalent in functionality would be a Route-Refresh message, which is a separate extension. Therefore, in the context of standard BGP operations within the Established state, Updates and Notifications are the correct answers.

Question: 4

What is the default route preference for an aggregate route?

- A. 180
- B. 150
- C. 130
- D. 5

Answer: C

Explanation:

In the Junos OS architecture, route preference (often referred to as administrative distance in other vendor platforms) is the primary metric used by the Routing Engine to select the "best" path when multiple protocols provide a route to the same destination. Each routing protocol and route type is assigned a default numeric value; the lower the value, the more preferred the route.

According to Juniper Networks technical documentation, an aggregate route is assigned a default preference of 130. Aggregate routes are a form of static-like route used to group specific routes into a single, broader prefix to reduce the size of routing tables and limit the scope of routing updates. They are "protocol-independent" because they are not learned from a dynamic neighbor but are manually defined by the administrator.

To understand where 130 fits in the hierarchy, it is helpful to compare it with other common Junos preferences:

Directly connected interfaces: 0

Static routes: 5

OSPF Internal: 10

IS-IS Level 1/2: 15/18

Aggregate routes: 130

OSPF AS External: 150

BGP (Internal and External): 170

Generated routes: 150

By setting the aggregate route preference to 130, Junos ensures that specific routes learned via IGPs (like OSPF or IS-IS) are preferred over the aggregate. This is essential because an aggregate route is often used as a "catch-all" or a discard route when more specific path information is missing. If the aggregate had a lower preference (like 5), it might override dynamic routing information, leading to suboptimal routing or black-holed traffic.

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>