

JN0-364 - Service Provider Routing and Switching, Specialist (JNCIS-SP)–High-quality Valid Test Book



As you can see on our website, there are versions of the PDF, Software and APP online. PDF version of our JN0-364 study materials- it is legible to read and remember, and support customers' printing request. Software version of our JN0-364 exam questions-It support simulation test system and times of setup has no restriction. Remember this version support Windows system users only. App online version of JN0-364 Practice Engine -Be suitable to all kinds of equipment or digital devices.

FreeCram's Juniper JN0-364 exam questions pdf is formed in a proper way that gives candidates the necessary asthenic unformatted data required to pass the Juniper exam. The study materials highlight a few basic and important questions that are repeatedly seen in past Juniper exam paper sheets. The Juniper JN0-364 Practice Questions are easy to access and can be downloaded anytime on your mobile, laptop, or MacBook.

[**>> Valid JN0-364 Test Book <<**](#)

Free PDF Quiz 2026 Juniper JN0-364: Service Provider Routing and Switching, Specialist (JNCIS-SP) – Trustable Valid Test Book

Generally speaking, JN0-364 certification has become one of the most authoritative voices speaking to us today. Let us make our life easier by learning to choose the proper JN0-364 study materials, pass the exam, obtain the certification, and be the master of your own life, not its slave. There are so many of them that they make you believe that their product is what you are looking for. With one type of JN0-364 Study Materials are often shown one after another so that you are confused as to which product you should choose.

Juniper Service Provider Routing and Switching, Specialist (JNCIS-SP) Sample Questions (Q10-Q15):

NEW QUESTION # 10

You have configured an MPLS LSP that begins on R1 and terminates on R5 using the Junos default settings. Referring to the exhibit, which router will perform only label swap operations?

- A. R1
- B. R3
- C. R5
- D. R4

Answer: B

Explanation:

In an MPLS network, routers are categorized by their role along a Label Switched Path (LSP). In this scenario, the LSP originates on R1 (Ingress LER) and terminates on R5 (Egress LER). Between these two endpoints are the Provider (P) routers, also known as Transit Label Switching Routers (LSRs), which include R2, R3, and R4.

To identify which router performs only label swap operations, we must look at the standard Junos data plane behavior:

* R1 (Ingress LER): Performs a Push operation. It receives native IP traffic from Networks 1 or 2, looks up the destination, and imposes (pushes) an MPLS label onto the packet before sending it to R2.

* R2 and R3 (Transit LSRs): These routers perform a Swap operation. They receive a labeled packet, look up the incoming label in their Label Forwarding Information Base (LFIB), replace it with an outgoing label provided by the downstream neighbor, and forward it.

* R4 (Penultimate Hop): By default, Junos uses Penultimate Hop Popping (PHP). Because R4 is the second-to-last router before the egress (R5), the egress router R5 advertises an "implicit-null" label (Label 3) to R4. This instructs R4 to perform a Pop operation—removing the MPLS label entirely—and sending the native IP packet to R5.

* R5 (Egress LER): Receives the packet (which is already unlabeled due to PHP) and performs a standard IP route lookup to reach the final destination in Network 3 or 4.

Among the options provided, R3 is the only router that is a transit LSR but not the penultimate hop. While R2 also performs a swap, it is not an option. R4 performs a Pop (due to PHP), R1 performs a Push, and R5 performs an IP lookup. Therefore, R3 is the correct answer as it solely performs the label swap operation.

NEW QUESTION # 11

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. RSVP
- D. BGP

Answer: A

Explanation:

In Juniper Networks Junos OS, the Label Distribution Protocol (LDP) is specifically designed to automate the creation of Label Switched Paths (LSPs) based on the information provided by the underlying Interior Gateway Protocol (IGP), such as OSPF or ISIS. When LDP is enabled on a set of interfaces within an OSPF area (as shown in the exhibit with Area 0.0.0.0), it automatically discovers neighbors and exchanges label mappings for all known unicast routes in the routing table.

The defining characteristic of LDP in this context is its "topology-driven" nature. Unlike RSVP (Resource Reservation Protocol), which typically requires the manual configuration of each LSP ingress point and destination, LDP follows the IGP's shortest path tree to automatically build a full mesh of LSPs between all participating routers. This means that every Provider Edge (PE) and Provider (P) router in the exhibit—PE1, PE2, PE3, P1, P2, and P3—will establish label-switched connectivity to every other router without the administrator having to define individual tunnels.

LDP accomplishes this through a downstream-unsolicited label distribution mode by default in Junos. Each router assigns a local label for its loopback address and other prefixes and advertises these to its neighbors.

Because every router is performing this action for every reachable prefix in the OSPF domain, a complete fabric of label-switched paths is formed. While RSVP is more robust for traffic engineering and bandwidth reservation, LDP is the preferred protocol for creating a simple, scalable full mesh of LSPs for applications like Layer 3 VPNs or internal BGP tunneling where complex path

manipulation is not required. BFD is a failure detection protocol, and BGP is used for service signaling, making LDP the only correct choice for automatic mesh creation.

NEW QUESTION # 12

Exhibit:

Referring to the exhibit, R1 is advertising prefix 203.0.113.0/24 to R2 over EBGP. R2 is configured to advertise this prefix into IBGP. R3 receives the 203.0.113.0/24 route, however the route is hidden.

Which configuration statement do you need to add to R2 to solve this problem?

- A. set policy-options policy-statement export-to-ibgp then local-preference 50
- B. set protocols bgp group EBGP export export-to-ibgp
- **C. set policy-options policy-statement export-to-ibgp then next-hop self**
- D. set policy-options policy-statement export-to-ibgp from route-filter 203.0.113.0/24 orlonger

Answer: C

Explanation:

In Juniper Networks Junos OS, a "hidden" route in the BGP table typically signifies that the router has received the prefix but cannot install it into the active routing table because the BGP next hop is unreachable.

This is a common occurrence in service provider environments when transitioning between External BGP (EBGP) and Internal BGP (IBGP).

According to Juniper technical documentation, when an EBGP speaker (R1) advertises a prefix to its peer (R2), it sets the next hop to its own interface IP address (\$172.16.10.1\$). By default, when R2 re-advertises that prefix to its IBGP peer (R3), it preserves the original EBGP next-hop address. Unless R3 has a specific route in its Interior Gateway Protocol (IGP) or a static route to reach the \$172.16.10.1\$ subnet, it will mark the route as unusable (hidden).

In the exhibit, the show route output on R3 explicitly shows the nexthop for \$203.0.113.0/24\$ as \$172.16.10.1\$. Since this route is marked "hidden," we can conclude R3 does not know how to reach R2's external peering link.

To resolve this, the network administrator must modify the next-hop attribute before the route is sent to R3.

By adding the statement `set policy-options policy-statement export-to-ibgp then next-hop self` (Option B) on router R2, R2 will replace the external next-hop (\$172.16.10.1\$) with its own internal peering address (\$172.16.20.1\$) before advertising the route to R3. Because R3 already has a direct or IGP connection to R2's internal address, it will successfully resolve the next hop, and the route will transition from "hidden" to

"active."

Option A is unnecessary because the route is already being exported; Option C is redundant as the policy is already applied to the IBGP group; and Option D changes path preference but does not solve the underlying reachability problem.

NEW QUESTION # 13

Which IS-IS adjacency state indicates that hello packets have been exchanged but the adjacency is not yet fully established?

- A. up
- B. two-way
- C. loading
- **D. initializing**

Answer: D

Explanation:

In the IS-IS (Intermediate System to Intermediate System) protocol, the process of forming an adjacency between two neighbors follows a specific sequence of states. While OSPF uses states like "Init," "Two-Way," and "Full," IS-IS uses a slightly different nomenclature within its state machine.

According to Juniper Networks technical documentation, when a router first sends an IS-IS Hello (IIH) PDU and receives one back from a neighbor, but has not yet confirmed that the neighbor "sees" it back, the adjacency enters the `initializing` state. Specifically, on a point-to-point link, the state transitions from `Down` to `initializing` as soon as the first PDU is received. On a broadcast network (like Ethernet), the `initializing` state indicates that the local router has received a Hello PDU from the neighbor, but the local router's own System ID is not yet listed in the neighbor's list of "seen" neighbors (the neighbor's Hello PDU does not yet contain the local router's MAC address).

The adjacency only moves to the `Up` state (Option C) once bi-directional communication is confirmed- meaning both routers have seen each other's System IDs in the incoming Hello PDUs.

Why other options are incorrect:

- * Loading (Option A): This is an OSPF state, not an IS-IS state. In IS-IS, database synchronization happens after the adjacency is Up.
- * Two-Way (Option D): While functionally similar to the state IS-IS is achieving, "Two-Way" is the specific terminology for OSPF. In IS-IS, the intermediate step between knowing a neighbor exists and having a fully functional adjacency is strictly called **Initializing**.

NEW QUESTION # 14

In IS-IS, what would you use to control which external routes are installed in the routing table?

- A. route preference
- B. interface metric
- **C. import policy**
- D. export policy

Answer: C

Explanation:

In Junos OS, the flow of routing information is managed by policies that sit between the protocol's database (the RIB-In/LSDB) and the main routing table (inet.0). Understanding the direction of these policies is critical for correct configuration.

An import policy (Option B) is used to control the movement of routes from a routing protocol into the routing table. According to Juniper Service Provider documentation, even though IS-IS is a link-state protocol that requires all routers in an area to have an identical Link-State Database (LSDB), an import policy can be used to filter which of those validated routes are actually placed into `inet.0` for forwarding. For external routes (routes leaked into IS-IS from other areas or protocols), an import policy allows an administrator to selectively accept or reject prefixes based on specific criteria like prefix-lists or community tags.

It is important to distinguish this from an export policy (Option A). In Junos, an export policy is used to take routes already in the routing table and push them out to a protocol to be advertised to neighbors. For example, you would use an export policy to redistribute static routes into IS-IS. Route preference (Option C) is a global value used to select between different protocols for the same prefix, and the interface metric (Option D) is used by the SPF algorithm to calculate the shortest path within the IS-IS database itself. Therefore, to specifically control which learned external routes are "installed" into the forwarding table, the import policy is the correct tool.

NEW QUESTION # 15

.....

The JN0-364 study braindumps are compiled by our professional experts who have been in this career for over ten years. Carefully written and constantly updated content of our JN0-364 exam questions can make you keep up with the changing direction of the exam, without aimlessly learning and wasting energy. In addition, there are many other advantages of our JN0-364 learning guide. Hope you can give it a look and you will love it for sure!

JN0-364 New Question: <https://www.freecram.com/Juniper-certification/JN0-364-exam-dumps.html>

Our Juniper JN0-364 training materials will help you save money, energy and time. Our FreeCram JN0-364 New Question can always help you solve this problem quickly. Juniper Valid JN0-364 Test Book You can practice our sample questions for free, so you just need to knock the keyboard without any loss and in return for the opportunity for success. The FreeCram recognizes that Juniper JN0-364 aspirants are constantly juggling multiple responsibilities, so questions are ideal for quick preparation.

Don't Want to Purchase from Apple, A major change to the Office XP application New JN0-364 Exam Vce environment is the introduction of task panes, which are multipurpose windowpanes that appear on the right side of an Office application window.

JN0-364 free questions & JN0-364 torrent vce & JN0-364 dumps torrent

Our Juniper JN0-364 Training Materials will help you save money, energy and time. Our FreeCram can always help you solve this problem quickly. You can practice our sample questions for free, so JN0-364 you just need to knock the keyboard without any loss and in return for the opportunity for success.

The FreeCram recognizes that Juniper JN0-364 aspirants are constantly juggling multiple responsibilities, so questions are ideal for quick preparation. Please do not worry;

- Real JN0-364 Questions Valid JN0-364 Study Notes Learning JN0-364 Materials Search for JN0-364 and download it for free on "www.easy4engine.com" website Latest JN0-364 Mock Test

