

JN0-364 Latest Exam Pdf, Exam JN0-364 Tips

Download the latest JN0-636 dumps pdf to pass your exam successfully

Exam : JN0-636

Title : Security,Professional
(JNCIP-SEC)

<https://www.passcert.com/JN0-636.html>

1 / 6

2026 Latest ExamsLabs JN0-364 PDF Dumps and JN0-364 Exam Engine Free Share: <https://drive.google.com/open?id=1FaQby82yJdKl3hdFOd7OuVchW9HgpQjA>

We learned that a majority of the candidates for the exam are office workers or students who are occupied with a lot of things, and do not have plenty of time to prepare for the JN0-364 exam. Taking this into consideration, we have tried to improve the quality of our JN0-364 training materials for all our worth. Now, I am proud to tell you that our JN0-364 Exam Questions are definitely the best choice for those who have been yearning for success but without enough time to put into it. Just buy them and you will pass the exam by your first attempt!

JN0-364 certifications are thought to be the best way to get good jobs in the high-demanding market. There is a large range of JN0-364 certifications that can help you improve your professional worth and make your dreams come true. Our JN0-364 Certification Practice materials provide you with a wonderful opportunity to get your dream certification with confidence and ensure your success by your first attempt.

>> JN0-364 Latest Exam Pdf <<

Exam JN0-364 Tips & Exam JN0-364 Guide

Our JN0-364 practice questions and answers are created according to the requirement of the certification center and the latest exam information. Our JN0-364 real dumps cover the comprehensive knowledge points and latest practice materials that enough to help you Clear JN0-364 Exam tests. You will get our valid JN0-364 dumps torrent and instantly download the exam pdf after payment.

Juniper Service Provider Routing and Switching, Specialist (JNCIS-SP) Sample Questions (Q21-Q26):

NEW QUESTION # 21

What prevents routing loops in a single-area OSPF network?

- A. The Dijkstra algorithm
- B. Routing policies
- C. The Bellman-Ford algorithm
- D. Forwarding policies

Answer: A

Explanation:

In OSPF, loop prevention within a single area is achieved through the fundamental nature of its link-state architecture. Unlike distance-vector protocols that rely on "routing by rumor," OSPF ensures that every router within an area maintains an identical Link-State Database (LSDB). This database acts as a complete map of the network topology.

Once the LSDB is synchronized, each router independently executes the Shortest Path First (SPF) algorithm, which is formally known as the Dijkstra algorithm. This mathematical process treats the local router as the "root" of a tree and calculates the shortest path to every other node (router) and prefix in the area based on the cumulative interface costs. Because every router uses the same synchronized map (the LSDB) and the same deterministic algorithm, they all arrive at a consistent, loop-free view of the best paths.

According to Juniper Networks technical documentation, the Dijkstra algorithm is superior to the Bellman-Ford algorithm (used by distance-vector protocols like RIP) in this regard. Bellman-Ford is susceptible to

"count-to-infinity" problems and loops because routers only know the distance and direction to a destination provided by their neighbors, rather than the full topology. In OSPF, even if a link fails, the updated Link-State Advertisement (LSA) is flooded rapidly, and the Dijkstra algorithm is re-run to find a new loop-free path.

Routing policies (Option B) are used to manipulate path selection or filter routes but are not the primary mechanism for fundamental loop prevention in OSPF. Similarly, forwarding policies (Option D) govern how traffic is handled at the data plane level rather than determining the control plane's loop-free topology.

NEW QUESTION # 22

During OSPF neighbor establishment, which packet type is used to describe the contents of the link-state database?

- A. Link-State PDU (LSP)
- B. Link-State Request (LSR)
- C. Database Description (DBD)
- D. Hello packet

Answer: C

Explanation:

In the OSPF (Open Shortest Path First) protocol, ensuring that all routers within an area have a synchronized Link-State Database (LSDB) is fundamental to building a consistent loop-free topology. During the adjacency formation process—specifically when transitioning from the ExStart state to the Exchange state—routers must determine what information they are missing from their neighbors without sending the entire database at once, which would be highly inefficient.

The Database Description (DBD) packet, also known as a DDP, is the mechanism used for this summary exchange. According to Juniper Networks technical documentation, the DBD packet does not contain full Link-State Advertisements (LSAs). Instead, it contains only the LSA headers, which include the LSA type, the ID of the advertising router, and the sequence number.

By exchanging these headers, a Juniper router can compare the neighbor's database summary against its own local LSDB. If the router identifies a header in the DBD packet that represents a newer or missing entry, it records that LSA in its "Link-State Request List." This collaborative "handshake" ensures that only the necessary, updated information is requested in the subsequent Link-State Request (LSR) phase. It is important to distinguish this from the Link-State PDU (LSP) mentioned in Option D, which is actually the term used in the IS-IS protocol, not OSPF. In OSPF, the functional unit is the LSA, and the transport vehicle for the initial summary is the DBD packet. This methodical synchronization is what allows OSPF to scale effectively in large service provider environments.

NEW QUESTION # 23

□ In the exhibit, Site A is sending traffic to Site B. R1 adds MPLS label 7166 to direct the traffic to R5. Which two criteria did R1 use

to determine which label number to add to the traffic? (Choose two.)

- A. the destination address of the traffic
- B. a label number received from R5
- C. a label number advertisement received from R2
- D. the source address of the traffic

Answer: A,C

Explanation:

In a Juniper Networks MPLS environment, the process by which a router determines how to forward traffic involves both the control plane and the data plane. When R1 (acting as an Ingress Label Edge Router, or LER) receives an IP packet from Site A destined for Site B, it must perform a lookup to decide whether to forward the packet via standard IP routing or via an MPLS Label Switched Path (LSP).

The first criterion R1 uses is the destination address of the traffic (Option C). Upon receiving the native IP packet, R1 looks up the destination IP in its routing table (typically inet.0). If the destination matches a prefix that is associated with an LSP—such as the loopback address of R5 or a prefix reachable via R5—the router identifies the appropriate Forwarding Equivalence Class (FEC). The FEC essentially groups packets that should be forwarded in the same manner over the same path. Without identifying the destination, the router cannot map the traffic to the correct MPLS tunnel.

The second criterion is the label number advertisement received from R2 (Option D). MPLS relies on downstream label allocation. In this topology, R2 is the immediate downstream "next hop" for R1 on the path to Site B. For the LSP to be established, R2 must signal a label to R1 using a protocol like LDP (Label Distribution Protocol) or RSVP (Resource Reservation Protocol). This label (in this case, 7166) tells R1: "If you want to send traffic to the destination associated with this LSP, wrap it in this specific label so I know how to process it." R1 does not use the source address (Option A) for standard label mapping, nor does it receive the label directly from R5 (Option B) in a hop-by-hop signaling model; it must use the label provided by its direct neighbor, R2. Therefore, by combining the destination IP (to find the path) and the label provided by the next hop (to encapsulate the packet), R1 successfully directs the traffic through the MPLS core.

NEW QUESTION # 24

Exhibit:

```
user@R2> show route 198.51.100.1
inet.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
198.51.100.1/32 *[Static/5] 5d 21:02:26
> to 203.0.113.65 via ge-0/0/3.0
user@R2> show route 172.20.110.0/24
inet.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active,
* = Both
172.20.110.0/24 *[Static/5] 10:43:01
> via gr-0/0/0.0
```

Referring to the exhibit, traffic destined to which network will be sent through the tunnel?

- A. 0.0.0.0/0
- B. 172.20.110.0/24
- C. 203.0.113.65
- D. 198.51.100.1/32

Answer: B

Explanation:

To determine which traffic is being sent through a tunnel in a Junos OS environment, an administrator must analyze the routing table output for the exit interface associated with each destination prefix. The provided exhibit shows the results of the show route command on router R2 for two specific destination networks.

In the first output, the destination 198.51.100.1/32 is an active static route. The next-hop information specifies that traffic for this address is sent to the gateway 203.0.113.65 via the interface ge-0/0/3.0. According to Juniper Networks interface naming conventions, the prefix ge- denotes a Gigabit Ethernet interface, which represents a standard physical connection. Therefore, this traffic does not traverse a tunnel.

In the second output, the destination 172.20.110.0/24 is also an active static route. However, the next-hop for this network is listed as via gr-0/0/0. In the Junos operating system, the gr- prefix explicitly identifies a Generic Routing Encapsulation (GRE) tunnel interface. GRE is a widely used protocol in service provider networks to encapsulate various network layer protocols over an IP backbone, effectively creating a virtual point-to-point link. Because the routing table has installed the route for 172.20.110.0/24 specifically via the gr- interface, all traffic destined for this network will be encapsulated and sent through the tunnel.

The other choices are incorrect for the following reasons:

- * 203.0.113.65 (Option B): This is the next-hop IP address for the physical Gigabit Ethernet path; it is not a destination network directed to a tunnel.
- * 0.0.0.0/0 (Option C): There is no information in the exhibit regarding a default route.
- * 198.51.100.1/32 (Option D): As identified by the ge-interface prefix in the exhibit, traffic for this destination is sent via a physical Ethernet link.

NEW QUESTION # 25

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

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

Answer: C

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 IS-IS. 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 # 26

.....

ExamsLabs JN0-364 valid test will assist you to pass your JN0-364 actual test with ease. You will never regret to choose our JN0-364 exam engine test. Here are some outstanding properties which can benefit all of you. The detailed explanations are offered where available to ensure you fully understand why to choose the correct answers. All the questions cover the main points which the JN0-364 Actual Exam required. The answers of each question are correct and verified by our IT experts which can ensure you 100% pass.

Exam JN0-364 Tips: <https://www.examslabs.com/Juniper/JNCIS-SP/best-JN0-364-exam-dumps.html>

The ExamsLabs Service Provider Routing and Switching, Specialist (JNCIS-SP) (JN0-364) PDF dumps file work with all devices and operating system, Get Command on Every Bit of Juniper JN0-364 Exam Questions, If you still worry about further development in IT industry you are doing the right thing now to scan our website about JN0-364 certification exam prep and our good JN0-364 passing rate, At the same time, many people pass the exam for the first time under the guidance of our JN0-364 practice exam.

If the Start/Stop Jobs button has a stop sign JN0-364 on it, that means click me to stop sending data to the printer, Attracting

