

Printable JN0-364 PDF - JN0-364 Exam Demo



JUNIPER JN0-351 STUDY GUIDE PDF

Juniper JNCIS Routing and Switching Certification Questions & Answers

[Details of the Exam-Syllabus-Questions](#)

JN0-351

Juniper Networks Certified Specialist Enterprise Routing and Switching
65 Questions Exam – Variable (60-70% Approx.) Cut Score – Duration of 90 minutes

JN0-364 practice exam takers can even access the results of previous attempts which helps them in knowing and overcoming their mistakes before appearing in the JN0-364 final test. There are thousands of students that bought Lead2PassExam's JN0-364 Practice Exam and got success on their initial tries. We guarantee that if you take our provided Juniper JN0-364 exam dumps you will crack the JN0-364 Exam in a single try.

Our JN0-364 learning question can provide you with a comprehensive service beyond your imagination. JN0-364 exam guide has a first-class service team to provide you with 24-hour efficient online services. Our team includes industry experts & professional personnel and after-sales service personnel, etc. Industry experts hired by JN0-364 Exam Guide helps you to formulate a perfect learning system, and to predict the direction of the exam, and make your learning easy and efficient. Our staff can help you solve the problems that JN0-364 test prep has in the process of installation and download.

>> [Printable JN0-364 PDF](#) <<

JN0-364 Exam Demo - JN0-364 Associate Level Exam

It is our biggest goal to try to get every candidate through the exam. Although the passing rate of our JN0-364 simulating exam is nearly 100%, we can refund money in full if you are still worried that you may not pass the JN0-364 exam. You don't need to worry about the complexity of the refund process at all, we've made it quite simple. And if you really want to pass the exam instead of refund, you can wait for our updates for we will update our JN0-364 Study Guide for sure to make you pass the exam.

Juniper Service Provider Routing and Switching, Specialist (JNCIS-SP) Sample Questions (Q33-Q38):

NEW QUESTION # 33

You are a network architect designing a brand new network. You want to deploy RSVP LSPs in this network. You are currently in the process of choosing whether to run OSPF or IS-IS as your interior gateway protocol. In this scenario, which two statements are correct about IGP traffic engineering extensions in an RSVP network? (Choose two.)

- A. You must explicitly configure IS-IS to carry traffic engineering extensions.
- B. You must explicitly configure OSPF to carry traffic engineering extensions.
- C. In IS-IS, traffic engineering extensions are enabled by default.
- D. In OSPF, traffic engineering extensions are enabled by default.

Answer: B,C

Explanation:

In a Juniper Networks environment, deploying RSVP-signaled LSPs requires a functional Traffic Engineering Database (TED). This database is populated by the Interior Gateway Protocol (IGP) using specific extensions that carry link-state information beyond simple reachability, such as available bandwidth, administrative groups (link coloring), and Maximum Reservable Bandwidth.

The behavior of these extensions differs between OSPF and IS-IS in Junos OS:

* OSPF (Option C): By default, OSPF is a "pure" routing protocol. To support RSVP-TE, it must carry Opaque LSAs (Type 10).

According to Juniper documentation, you must explicitly configure traffic engineering within the OSPF protocol hierarchy using the set protocols ospf traffic-engineering command. Without this command, OSPF will not flood the TE information required by the Constrained Shortest Path First (CSPF) algorithm, and LSPs will fail to establish.

* IS-IS (Option D): IS-IS was designed to be extensible through the use of TLVs (Type, Length, Value).

In Junos OS, IS-IS traffic engineering extensions are enabled by default once the protocol is active.

As soon as you enable IS-IS on an interface, it begins to advertise the wide metrics and TE TLVs (like TLV 22 and 135) necessary for building the TED.

This distinction is a common design consideration for network architects. While IS-IS simplifies the rollout of MPLS by having TE enabled "out of the box," OSPF requires that extra configuration step to transition from a standard IGP to a TE-aware protocol.

NEW QUESTION # 34

Which OSPF packet type is used to initiate and maintain neighbor relationships?

- A. Hello
- B. Link-State Update
- C. Database Description
- D. Link-State Acknowledgment

Answer: A

Explanation:

The Hello packet is the most basic, yet most vital, component of the OSPF protocol. It serves as the primary mechanism for neighbor discovery, parameter negotiation, and "keepalive" functionality. Per Juniper Networks' routing documentation, OSPF routers use the Hello protocol to dynamically discover other OSPF-enabled routers on their directly connected segments.

When OSPF is enabled on a Junos interface, the router begins multicasting Hello packets (typically to the 224.0.0.5 "All OSPF Routers" address). This initiates the neighbor relationship. For two routers to move beyond the Init state and become neighbors, they must agree on several critical parameters contained within the Hello packet:

* Area ID: Routers must be in the same OSPF area.

* Authentication: Passwords or keys must match.

* Timers: The Hello and Dead intervals must be identical.

* Options: Such as Stub area flags.

Beyond the initial "initiation," the Hello packet is used to maintain the relationship. By continuously sending these packets at a fixed interval (the Hello interval), a router signals to its peers that it is still functional. If a router stops receiving Hello packets from a neighbor for a duration exceeding the Dead Interval, it declares the neighbor "down," flushes the associated LSAs from the database, and triggers a new SPF calculation.

Furthermore, on multi-access networks like Ethernet, the Hello packet is the vehicle for the election of the Designated Router (DR) and Backup Designated Router (BDR). By exchanging priority values and Router IDs within the Hello packets, the segment can elect a central point of contact to minimize the number of adjacencies required on the wire.

NEW QUESTION # 35

Which two statements regarding GRE and IP-IP tunnels are correct? (Choose two.)

- A. These tunnels add additional overhead to the packets that traverse them.
- B. These tunnels offer secure encryption mechanisms.
- C. These tunnels do not add any overhead to the packets that traverse them.
- D. These tunnels do not offer encryption mechanisms.

Answer: A,D

Explanation:

In Juniper Networks Junos OS, Generic Routing Encapsulation (GRE) and IP-in-IP (IP-IP) are common tunneling mechanisms used to transport packets across a network by encapsulating them within another protocol. Understanding the header structure and the limitations of these protocols is essential for proper MTU (Maximum Transmission Unit) management and security design.

Overhead (Option A):

Both GRE and IP-IP tunnels operate by adding an additional IP header to the original packet. An IP-IP tunnel (Protocol 4) adds a 20-byte IPv4 header. A GRE tunnel (Protocol 47) adds the same 20-byte delivery IP header plus a minimum 4-byte GRE header (totaling 24 bytes, which can increase if keys or sequencing are used).

Because these headers are added to the payload, the total size of the packet increases. This "overhead" means that if the original packet was already at the MTU limit (e.g., 1500 bytes), the encapsulated packet will exceed it, potentially leading to fragmentation or the need to adjust the TCP MSS (Maximum Segment Size).

Encryption (Option D):

Crucially, according to Juniper Service Provider documentation, neither GRE nor IP-IP provides native encryption or data confidentiality. They are encapsulation protocols, not security protocols. The payload remains in plaintext and is visible to any device along the path. If security and encryption are required for data traversing these tunnels, they must be combined with IPsec (IP Security). While GRE is often used as the "carrier" for IPsec (to allow multicast or dynamic routing protocols which IPsec alone does not support), the GRE protocol itself remains an unencrypted delivery mechanism. Therefore, statements A and D accurately describe the architectural behavior of these tunnel types.

NEW QUESTION # 36

You are the administrator for two Junos routers called R1 and R2. These two routers are directly connected to each other. These two routers run IS-IS and BFD. R1 is configured to send BFD packets every 300 milliseconds. R2 is configured to send BFD packets every 400 milliseconds. In this situation, what is the expected outcome?

- A. Each router will negotiate to send BFD packets at the slowest of the two rates.
- B. Each router will send BFD packets at the rate that has been locally configured.
- C. BFD will fail due to the mismatched timers.
- D. Each router will negotiate to send BFD packets at the fastest of the two rates.

Answer: A

Explanation:

In the context of Juniper Networks High Availability, Bidirectional Forwarding Detection (BFD) is a lightweight protocol designed to provide fast failure detection for the forwarding path. Unlike the slow "hello" mechanisms found in IGPs like OSPF or IS-IS, BFD can detect link or neighbor failures in sub-second intervals.

According to Juniper Networks technical documentation, BFD operates through a negotiation process. When two routers establish a BFD session, they exchange their locally configured Minimum Transmit Interval and Minimum Receive Interval within the BFD control packets. The fundamental rule of BFD negotiation is that the routers must agree on a common timing value that accommodates the slower of the two devices to ensure stability and prevent "false positives" (detecting a failure when none exists simply because one router cannot keep up with the processing speed).

In this scenario, R1 expects to send at 300ms, while R2 is configured for 400ms. During the handshake, R1 informs R2 it is capable of 300ms, but R2 informs R1 it can only support a minimum of 400ms.

Consequently, the routers will negotiate to use the slowest of the two rates (400ms). Specifically, the transmission interval of one router is matched to the receive interval of the other. By choosing the highest common denominator (the slowest rate), the BFD session ensures that both routers have sufficient time to process incoming control packets. This negotiation allows BFD to be highly flexible in heterogeneous environments where different hardware platforms may have varying CPU capabilities for handling rapid heartbeat packets.

NEW QUESTION # 37

Exhibit:

```
user@R1> show route 10.16.2.0/23 exact detail
```

inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
10.16.2.0/23 (1 entry, 1 announced)
*Aggregate Preference: 130
Next hop type: Reject
Address: 0x8f3fd44
Next-hop reference count: 2
State: <Active Int Ext>
Age: 1:39:21
Task: Aggregate
Announcement bits (1): 0-KRT
AS path: I (LocalAgg)
Flags: Depth: 0 Active
AS path list:
AS path: I Refcount: 2
Contributing Routes (2):
10.16.2.0/24 proto Direct
10.16.3.0/24 proto Direct
Which destination IP address will be matched by the aggregate route shown in the exhibit?

- A. packets destined to 10.16.4.183
- **B. packets destined to 10.16.3.79**
- C. packets destined to 10.16.0.4
- D. packets destined to 10.16.1.214

Answer: B

Explanation:

In the Juniper Networks Junos operating system, aggregate routes are used to represent a group of more specific routes with a single, shorter prefix. This technique is essential for reducing the size of routing tables and minimizing the volume of routing updates sent to neighbors. According to Juniper technical documentation, for a destination IP address to "match" a specific route, it must fall within the range defined by the network address and its associated CIDR mask.

The provided exhibit shows a detailed lookup for the aggregate route 10.16.2.0/23. To determine the range of IP addresses covered by a /23 mask, we examine the binary representation of the third octet. A /23 mask means the first 23 bits are fixed. For the address 10.16.2.0:

- * The first two octets (10.16) are fixed.
- * The third octet (2) is 00000010 in binary.
- * The 23rd bit is the second-to-last bit of this octet.
- * The /23 range allows the 24th bit (the last bit of the third octet) and all 8 bits of the fourth octet to vary.

This results in a range where the third octet can be either 2 (00000010) or 3 (00000011). Therefore, the aggregate route 10.16.2.0/23 covers all IP addresses from 10.16.2.0 to 10.16.3.255. The exhibit further confirms this by listing the "Contributing Routes": 10.16.2.0/24 and 10.16.3.0/24.

Analyzing the provided options against this range:

- * 10.16.3.79 (Option A): This address falls squarely within the 10.16.2.0 to 10.16.3.255 range.
- * 10.16.0.4 (Option B): This address falls in the 10.16.0.0/23 range (0.0 to 1.255).
- * 10.16.4.183 (Option C): This address falls in the 10.16.4.0/23 range (4.0 to 5.255).
- * 10.16.1.214 (Option D): This address also falls in the 10.16.0.0/23 range.

Consequently, 10.16.3.79 is the only destination listed that matches the aggregate route shown. It is also important to note the Next hop type: Reject in the exhibit; this means that if a packet matches the aggregate but does not match any of the more specific contributing routes, the router will drop the packet and send an ICMP unreachable message to the source.

NEW QUESTION # 38

.....

All contents of JN0-364 training prep are made by elites in this area rather than being fudged by laymen. Let along the reasonable prices of our JN0-364 exam materials which attracted tens of thousands of exam candidates mesmerized by their efficiency by proficient helpers of our company. Any difficult posers will be solved by our JN0-364 Quiz guide. And we have free demos of our JN0-364 study braindumps for you to try before purchase.

JN0-364 Exam Demo: <https://www.lead2passexam.com/Juniper/valid-JN0-364-exam-dumps.html>

