Pass Guaranteed Quiz Professional Juniper - JN0-364 Training Courses


Pass Guaranteed 2023 Juniper JN0-104: Reliable Junos, Associate (JNCIA-Junos) Latest Test Format

Try Approved Juniper JN0-104 Exam Questions To Pass JN0-104 Exam

Compared with other education platform on the market, Actual4test is more reliable and highly efficiently. It provide candidates who want to pass the JN0-364 exam with high pass rate study materials, all customers have passed the exam in their first attempt. They all need 20-30 hours to learn on our website can pass the exam. JN0-364 Exam Dump is really a high efficiently exam tool that can help you save much time and energy to do other things.

Now you can trust Actual4test JN0-364 exam questions as these Service Provider Routing and Switching, Specialist (JNCIS-SP) (JN0-364) exam questions have already helped countless candidates in their JN0-364 exam preparation. They easily got success in their challenging and dream Juniper JN0-364 Certification Exam. Now they have become certified Juniper professionals and offer their services to top world brands.

>> JN0-364 Training Courses <<

Pass Guaranteed Quiz Professional Juniper - JN0-364 - Service Provider Routing and Switching, Specialist (JNCIS-SP) Training Courses

Great concentrative progress has been made by our company, who aims at further cooperation with our candidates in the way of using our JN0-364 exam engine as their study tool. Owing to the devotion of our professional research team and responsible working staff, our training materials have received wide recognition and now, with more people joining in the JN0-364 Exam army, we has become the top-raking JN0-364 training materials provider in the international market. Believe in our JN0-364 study guide, you will succeed in your exam!

Juniper Service Provider Routing and Switching, Specialist (JNCIS-SP) Sample Questions (Q35-Q40):

NEW QUESTION # 35
In IS-IS, what would you use to control which external routes are installed in the routing table?

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

Answer: D

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.
Animport 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), animport policycan 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 anexport policy (Option A). In Junos, an export policy is used to take routesalready inthe routing table and push themoutto 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 theinterface 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, theimport policyis the correct tool.

NEW QUESTION # 36
Exhibit:
You have configured IPv4 and IPv6 in your network and all OSPF neighbors are established. You apply the configuration shown in the exhibit. Which statement is true in this scenario?

- A. There will only be an OSPFv3 entry in R1 for network 172.16.2.0/24.
- B. There will be an OSPFv2 and OSPFv3 entry in R1 for network 172.16.2.0/24.
- C. There will only be an OSPFv2 entry in R1 for network 172.16.2.0/24.
- D. There will not be a route in R1 for network 172.16.2.0/24.

Answer: B

Explanation:
In a Juniper Networks environment running Junos OS, understanding the interaction between different versions of OSPF is essential for multi-protocol environments.OSPFv2(defined in RFC 2328) is the standard protocol used for routing IPv4 unicast traffic.OSPFv3(defined in RFC 5340) was originally developed to support IPv6 routing. However, OSPFv3 was later extended via RFC 5838 to support multiple address families (AF), allowing it to carry IPv4 unicast, IPv4 multicast, and other address types within a single OSPF instance.
According to Juniper technical documentation, Junos OS implements this multi-AF support in OSPFv3 through the use ofrealms. When the realm ipv4-unicast statement is configured under the [edit protocols ospf3] hierarchy, the OSPFv3 process becomes capable of calculating and advertising IPv4 routes.
In the provided exhibit, routerR2has a dual-protocol configuration. First, it is running standard OSPFv2, with the ge-0/0/1.0 interface (which is directly connected to the 172.16.2.0/24 network) participating in Area 0.
This ensures that the prefix is advertised as a standard IPv4 LSA to its neighbor,R1. Second, R2 is running OSPFv3 with the realm ipv4-unicast specifically enabled on that same ge-0/0/1.0 interface. Because of this realm, OSPFv3 also treats the 172.16.2.0/24 prefix as a reachable IPv4 destination and advertises it to R1 as an OSPFv3 IPv4-unicast LSA.
As a result, whenR1(which is also running both protocols) receives these routing updates, it will see the same destination prefix advertised by two different protocols. Its routing table (inet.0) will contain one entry learned from the OSPFv2 process and a second, separate entry learned from the OSPFv3 process. While the Junos Routing Engine will ultimately select one as the "active"

route based on route preference (both protocols have a default preference of 10), both entries will technically exist within the Routing Information Base (RIB). This confirms that statementBis the correct description of the operational state of the network.

NEW QUESTION # 37
Exhibit:
Referring to the exhibit, why is the ge-0/0/0.0 interface shown as belonging to Level 3?

- A. This interface is configured as a broadcast interface, that uses Level 3 as shorthand for both Level 1 and Level 2.
- B. This interface is configured as a point-to-point interface, that uses Level 3 as shorthand for both Level 1 and Level 2.
- C. This interface connects to a super spine.
- D. This interface is configured as a broadcast interface that has three adjacencies with other routers on the shared LAN.

Answer: B

Explanation:
In theIS-IS (Intermediate System to Intermediate System)protocol as implemented in Junos OS, the output of operational commands uses specific numerical representations to denote the hierarchy levels of a neighbor adjacency. Understanding these values is crucial for troubleshooting peering relationships in a multi-level IS- IS network.
According to Juniper Networks technical documentation, the show isis adjacency command displays the status of the neighbors. The "L" column indicates the level of the adjacency:
* Level 1:Indicates the adjacency is strictly for intra-area routing.
* Level 2:Indicates the adjacency is strictly for backbone/inter-area routing.
* Level 3:This is ashorthand representationused by Junos to indicate that a single adjacency has been established forboth Level 1 and Level 2 simultaneously.
The critical distinction in this question lies in the interface type. On abroadcast interface(such as standard Ethernet), IS-IS typically establishes and maintains separate adjacencies for Level 1 and Level 2. In the CLI output for a broadcast link, you would generally see two separate lines for the same neighbor-one for Level
1 and one for Level 2.
However, on apoint-to-point (P2P)interface, IS-IS can negotiate both levels within a single adjacency. When this occurs, Junos consolidates the output into a single entry and usesLevel 3to signify that the adjacency is functional for both levels. Since the exhibit shows ge-0/0/0.0 as Level 3, it confirms that the link is configured with a point-to-point encapsulation (either natively or via the interface-type p2p command) and is acting as a Level 1/2 adjacency.
Option B is incorrect as the number "3" refers to protocol levels, not the count of neighbors. Option C is a reference to data center architectures that does not influence IS-IS level nomenclature. Option D is incorrect because, as noted, broadcast interfaces display these levels separately rather than using the Level 3 shorthand.

NEW QUESTION # 38
What are three extension headers supported by IPv6? (Choose three.)

- A. header checksum
- B. fragment
- C. hop-by-hop options
- D. protocol
- E. destination options

Answer: B,C,E

Explanation:
One of the most significant architectural improvements in IPv6 is the move from a complex, variable-length header (as seen in IPv4) to a streamlined, fixed-length base header of 40 bytes. Additional functionality that was previously handled by "Options" in IPv4 is now moved toExtension Headers, which are inserted between the IPv6 base header and the upper-layer protocol (TCP/UDP).
According to Juniper Networks technical documentation and RFC 8200, the following are valid IPv6 Extension Headers:
* Hop-by-Hop Options (Option B):This header carries optional information that must be examined by every node along the delivery path. It is used for features like the Router Alert and Jumbo Payload options.
* Fragment (Option E):Unlike IPv4, where any router can fragment a packet, in IPv6, fragmentation is performed only by the source node. The Fragment header contains the information necessary for the destination to reassemble the packet (Offset, Identification, and More Fragments flag).

* Destination Options (Option A):This header carries information intended only for the destination node. It can appear twice: once before a routing header and once after.
Why other options are incorrect:
* Protocol (Option C):In IPv4, this was a field in the header. In IPv6, this is replaced by theNext Header field, which identifies the type of the following header (whether it's an extension header or the upper- layer protocol).
* Header Checksum (Option D):This field was entirely removed in IPv6. IPv6 relies on the data link layer (Ethernet) and the transport layer (TCP/UDP) to perform error detection, significantly reducing the processing overhead for routers in the core of a service provider network.

NEW QUESTION # 39
Exhibit:
user@Router-1> show route 172.24/16
inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
...
172.24.0.0/24 *[OSPF/150] 01:31:31, metric 0, tag 0
> to 172.20.0.2 via ge-0/0/2.0
to 172.20.1.2 via ge-0/0/3.0
user@Router-1> show route forwarding-table
Routing table: default.inet
Internet:
Destination Type RtRef Next hop Type Index NhRef Netif
...
172.24.0.0/24 user 0
172.20.0.2 ucst 551 2 ge-0/0/2.0
172.20.1.2 ucst 552 2 ge-0/0/3.0
Referring to the exhibit, which two statements are true? (Choose two.)

- A. The default route load-balancing behavior of this router has been modified.
- B. The router is performing default route load-balancing behavior.
- C. This router will choose both next hops in the routing table.
- D. This router will only choose the next hop with a > next to it in the routing table.

Answer: A,C

Explanation:
In Junos OS, understanding the distinction between theRouting Information Base (RIB)and theForwarding Information Base (FIB)is fundamental to analyzing traffic patterns and load-balancing behavior. The RIB (show route) contains all prefixes learned via various protocols, while the FIB (show route forwarding-table) contains only the active next-hops that are actually programmed into the Packet Forwarding Engine (PFE).
According to Juniper Networks technical documentation, the default behavior for Junos OS when encounteringEqual-Cost Multipath (ECMP)routes is to select only a single next-hop from the available candidates in the RIB and install that single path into the FIB. In a default state, even if the show route output displays multiple next-hops for a destination like 172.24.0.0/24, only one would have the active route symbol (
>) and only that one would appear in the forwarding table.
In the provided exhibit, the show route output shows two next-hops for 172.24.0.0/24, but only the first one (172.20.0.2) is marked with the>symbol as the active selection. However, the subsequent show route forwarding-table output reveals thatboth next-hops(172.20.0.2 and 172.20.1.2) are currently present in the forwarding table for that same destination. This discrepancy indicates that thedefault load-balancing behavior has been modified (Option B). This modification is typically achieved by creating a routing policy with the action then load-balance per-packet (which actually results in flow-based load balancing) and applying it to the forwarding table via the export statement under [edit routing-options forwarding-table].
Because the forwarding table now contains both next-hops, the router is no longer restricted to a single path.
Therefore, therouter will choose both next-hops in the routing table (Option D)for packet forwarding, distributing flows across the two available Gigabit Ethernet interfaces (ge-0/0/2.0 and ge-0/0/3.0). This ensures higher utilized bandwidth and provides redundancy at the data plane level.

NEW QUESTION # 40
......

Actual4test are supposed to help you pass the JN0-364 exam smoothly. Don't worry about channels to the best JN0-364 study materials so many exam candidates admire our generosity of offering help for them. Up to now, no one has ever challenged our leading position of this area. The existence of our JN0-364 learning guide is regarded as in favor of your efficiency of passing the exam. Over time, our company is becoming increasingly obvious degree of helping the exam candidates with passing rate up to 98 to 100 percent. All our behaviors are aiming squarely at improving your chance of success on JN0-364 Exam.

Valid JN0-364 Test Sims: https://www.actual4test.com/JN0-364_examcollection.html

Juniper JN0-364 Training Courses The certificate is of great value in the job market, JN0-364 Practice Exam Questions with 100% Guaranteed Success, Would you like to obtain JN0-364 certificate, Juniper JN0-364 Training Courses The strength of our the IT elite team will make you feel incredible, Our JN0-364 exam prep and practice test own the best solution to work out your problem, you can download immediately JN0-364 guide torrent material the moment you finish your payment, with no one minute delay.

Our experienced experts and certified trainers used their rich-experience and professional knowledge to do the study of JN0-364 Actual Test file for many years and finally has developed the best training materials about the real exam.

Pass Guaranteed Quiz JN0-364 - Updated Service Provider Routing and Switching, Specialist (JNCIS-SP) Training Courses

Delete the network object, The certificate is of great value in the job market, JN0-364 Practice Exam Questions with 100% Guaranteed Success, Would you like to obtain JN0-364 certificate?

The strength of our the IT elite team will make you feel incredible, Our JN0-364 exam prep and practice test own the best solution to work out your problem, you can download immediately JN0-364 guide torrent material the moment you finish your payment, with no one minute delay.