

JN0-364 Test Assessment & JN0-364 Valid Exam Cost

JNCIA JN0-104 Practice Test Questions and Answers 2023 with complete solution

Three statements that describe Transmission Control Protocol?

- Uses a best effort delivery approach
- had delivery notification and error-checking mechanisms
- applications include HTTP and SMTP
- TCP is faster than User Datagram Protocol (UDP)
- TCP is a transport layer protocol - has delivery notification and error-checking, applications include HTTP and SMTP
- TCP is transport layer protocol

Two benefits of the disaggregated JUNOS OS?

- A-The Junos VM becomes hardware-independent and can be deployed on any hypervisor without modifications.
- B-Platform drivers and forwarding engine are removed from the control plane to increase performance.
- C-Increased flexibility to use different control plane versions
- D-The architecture facilitates programmability through provisioning the control plane, the data path and the platform API - B and D

Describes the connection between control plane and forwarding plane

- No rate limiter is configured by default
 - control traffic is preferred over exception traffic during congestion
 - A rate limiter is configured by default
 - Exception traffic is preferred over control traffic during congestion - Control traffic is preferred, a rate-limiter is configured by default
- The JUNOS OS sends all exception traffic destined for the RE over the internal link that connects the control and forwarding plane. The JUNOS OS rate limits the exception traffic traversing the internal link to protect the RE from DOS attacks. During congestion preference is given to local and control traffic destined for the RE

A packet enters a JUNOS device. No matching destination entry exists in the forwarding table. How will the device respond? - The PFE responds to the source with a destination unreachable message. The PFE, not the RE

Two examples of transit traffic:

- A-SCP traffic destined for the router's loopback interface
- B-SFTP traffic that enters and exits the same interface on a local router
- C-SCP traffic that enters one interface and exits another on a local router

If you want to get a higher salary or a promotion on your position, you need to work harder! Purchase our JN0-364 learning materials and stick with it. Then your strength will protect you. For as long as you study with our JN0-364 exam questions, then you will find that the content of our JN0-364 preparation braindumps is all the hot hit of the newest knowledge and keypoints of the subject, you will learn so much to master the skills which will help you solve your problems in your work. And besides, you can achieve the certification for sure with our JN0-364 study guide.

If you buy our JN0-364 study materials, then you can enjoy free updates for one year. After you start learning, I hope you can set a fixed time to check emails. If the content of the JN0-364 practice guide or system is updated, we will send updated information to your e-mail address. Of course, you can also consult our e-mail on the status of the product updates. I hope we can work together to make you better use our JN0-364 simulating exam.

>> JN0-364 Test Assessment <<

JN0-364 Valid Exam Cost | JN0-364 Practice Tests

By propagating all necessary points of knowledge available for you, our JN0-364 practice materials helped over 98 percent of former exam candidates gained successful outcomes as a result. Our JN0-364 practice materials have accuracy rate in proximity to 98 and over percent for your reference. Up to now we classify them as three versions. They are pdf, software and the most convenient one app. Each of them has their respective feature and advantage including new information that you need to know to pass the test.

Juniper Service Provider Routing and Switching, Specialist (JNCIS-SP) Sample Questions (Q41-Q46):

NEW QUESTION # 41

You are monitoring OSPF on a router and notice frequent state changes between Full and Down. Which condition would cause this behavior?

- A. route preference mismatch
- B. MTU mismatch
- C. physical interface flapping
- D. area ID mismatch

Answer: C

Explanation:

When troubleshooting OSPF in a service provider environment, distinguishing between "stuck" adjacencies and "flapping" adjacencies is the first step. A session that transitions frequently between Full and Down indicates that the relationship can be established successfully (meaning parameters match), but it cannot be maintained.

According to Juniper Networks documentation, the most common cause for a session to drop from Full to Down is the expiration of the Dead Interval. If a router does not receive a Hello packet within the Dead Interval (usually 40 seconds), it tears down the adjacency. A physical interface flapping (Option A) is the primary trigger for this. If the physical link or the underlying transport (like a leased line or a microwave link) goes down even momentarily, the OSPF process immediately detects the interface failure, flushes the neighbors, and moves the state to Down. As soon as the interface comes back up, the routers perform the Hello exchange and reach the Full state again, creating the flapping cycle.

Analysis of other options:

* MTU Mismatch (Option D): This typically causes the adjacency to get "stuck" in the Exchange or ExStart state. The routers can exchange small Hello packets, but when they try to send larger Database Description (DBD) packets that exceed the MTU, the packets are dropped, preventing the session from ever reaching "Full."

* Area ID Mismatch (Option C): This prevents the adjacency from even reaching the Init state; the routers will never form a neighbor relationship.

* Route Preference (Option B): This affects which route is chosen for the forwarding table but has no impact on the OSPF neighbor state machine itself.

NEW QUESTION # 42

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 10.16.1.214
- D. packets destined to 10.16.0.4

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 # 43

Exhibit:

Referring to the exhibit, R1 and R2 are configured to run IS-IS. The IS-IS adjacency between R1 and R2 is up. What does the output of the show isis interface command tell you about R1?

- A. R1 sends Level 1 hello PDUs to R2.
- B. R1 advertises a Level 1 metric of 100 and a Level 2 metric of 100 toward R2 in its link-state PDU.
- C. R1 is not configured to use wide metrics.
- **D. R1 only forms a Level 2 adjacency with R2.**

Answer: D

Explanation:

In the IS-IS (Intermediate System to Intermediate System) protocol as implemented in Junos OS, routers can operate at two hierarchical levels: Level 1 (L1) for intra-area routing and Level 2 (L2) for inter-area backbone routing. By default, a Juniper router and its interfaces are configured to act as Level 1/2, meaning they will attempt to form adjacencies at both levels simultaneously.

According to Juniper Networks technical documentation, the show isis interface command provides a granular view of how the protocol is interacting with specific local links. In the provided exhibit, we must examine the L (Level) column and the DR (Designated Router) status columns to understand R1's operational state.

* Level Configuration: Under the L column for both the physical interface $ge-0/0/0.0$ and the loopback $lo0.0$, the value is strictly 2.

This indicates that these interfaces have been explicitly configured to operate only at Level 2.

* Adjacency Capabilities: For the interface $ge-0/0/0.0$, the Level 1 DR field is marked as Disabled. This confirms that R1 is not participating in Level 1 operations on this link; it will not transmit Level 1 Hello PDUs, nor will it listen for them. Consequently, R1 is incapable of forming a Level 1 adjacency with R2 on this segment.

* Metric Implications: The exhibit shows an L1/L2 Metric of $100/100$. In Junos, "narrow" metrics (the default) are limited to a maximum value of 63 per interface. A metric of 100 indicates that wide metrics (wide-metrics-only) have been enabled. Therefore, option A is incorrect because the router is using wide metrics.

Since the prompt states the adjacency is "up," and the interface is restricted to Level 2, we can conclude that R1 only forms a Level 2 adjacency with R2 (Option B). Even though an L1 metric of 100 is displayed in the table as a configured value, it is not actually "advertised" in a Link-State PDU because the Level 1 protocol is disabled on that interface.

NEW QUESTION # 44

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

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

Answer: C

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 # 45

You are designing a high availability solution for a Juniper router with dual Routing Engines (RE). You want to ensure that the routing protocol state is preserved during an RE switchover. You have already enabled graceful Routing Engine switchover (GRES) and you want to avoid relying on helper routers to maintain the routing protocol state. In this scenario, which feature would accomplish this behavior?

- A. bidirectional forwarding detection
- B. non-stop active routing
- C. graceful restart
- D. non-stop active bridging

Answer: B

Explanation:

When designing High Availability (HA) for Juniper Service Provider routers, understanding the interaction between the control plane and data plane is vital. The user has already enabled Graceful Routing Engine Switchover (GRES), which synchronizes the interface and kernel state between the primary and backup Routing Engines (REs). However, GRES by itself does not preserve the routing protocol state (like OSPF adjacencies or BGP sessions).

To achieve the preservation of the routing protocol state without relying on external "helper" routers, you must implement Non-Stop Active Routing (NSR). According to Juniper Networks documentation, NSR uses the infrastructure provided by GRES to also synchronize the routing protocol process (rpd) information.

Under NSR, the backup RE maintains a "hot" standby state of all routing protocols. If the primary RE fails, the backup RE takes over immediately. Because it already possesses the full routing table and peer session states, the peering neighbors are unaware that a switchover occurred. No protocol adjacency resets occur, and traffic continues to flow uninterrupted.

It is crucial to differentiate NSR from Graceful Restart (Option C). While Graceful Restart also aims to maintain traffic flow during a switchover, it does require help from neighboring routers (known as "helper mode"). If the neighbors do not support or are not configured for Graceful Restart, the sessions will drop.

Since the user explicitly stated they want to "avoid relying on helper routers," Graceful Restart is not the correct solution.

Non-stop Active Bridging (Option A) provides a similar "hitless" failover but specifically for Layer 2 environments (STP/VLANs) rather than Layer 3 routing protocols. BFD (Option B) is a failure detection protocol used to speed up convergence but does not preserve state during an RE failover; in fact, without NSR, BFD would likely trigger a faster teardown of the session during a switchover. Therefore, NSR is the only feature that meets the requirement for independent control-plane preservation.

