

# Prepare with Actual Juniper JN0-364 Exam Questions to Get Certified in First Attempt



The customer is God. JN0-364 learning dumps provide all customers with high quality after-sales service. After your payment is successful, we will dispatch a dedicated IT staff to provide online remote assistance for you to solve problems in the process of download and installation. During your studies, JN0-364 study tool will provide you with efficient 24-hour online services. You can email us anytime, anywhere to ask any questions you have about our JN0-364 Study Tool. At the same time, our industry experts will continue to update and supplement JN0-364 test question according to changes in the exam outline, so that you can concentrate on completing the review of all exam content without having to pay attention to changes in the outside world.

Believe that users will get the most satisfactory answer after consultation on our JN0-364 exam questions. Our online service staff is professionally trained, and users' needs about JN0-364 test guide can be clearly understood by them. The most complete online service of our company will be answered by you, whether it is before the purchase of JN0-364 training guide or the installation process, or after using the JN0-364 latest questions, no matter what problem the user has encountered. We will give you the best service and suggestion on the JN0-364 study material.

>> Review JN0-364 Guide <<

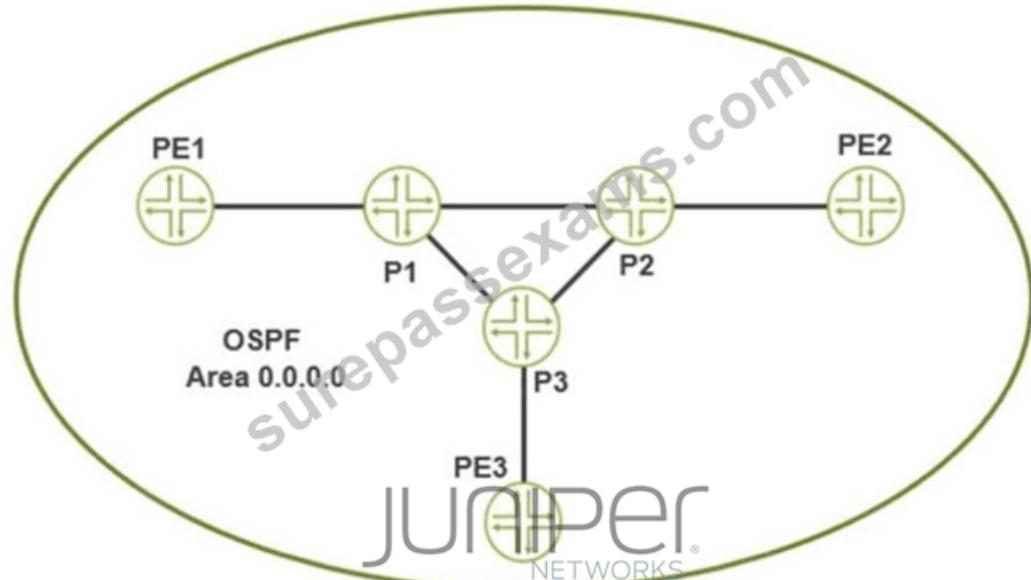
## JN0-364 Exam Passing Score, Latest JN0-364 Test Format

IT certifications are playing an important role in our career. In order to get a promotion and get more money, every IT people put more effort into their work. Instead this way, we can depend on our strength to won the boss's heart. Juniper JN0-364 certification is vitally important for IT people. In fact, the test is not difficult as you have imagined it. You only need to select the appropriate training materials. SurePassExams Juniper JN0-364 Practice Test will regularly update the exam dumps to fulfill your requirements. So, our Juniper JN0-364 test is the latest. Hurry up! You will achieve your aim.

## Juniper Service Provider Routing and Switching, Specialist (JNCIS-SP) Sample Questions (Q42-Q47):

### NEW QUESTION # 42

Exhibit



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

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

**Answer: B**

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.

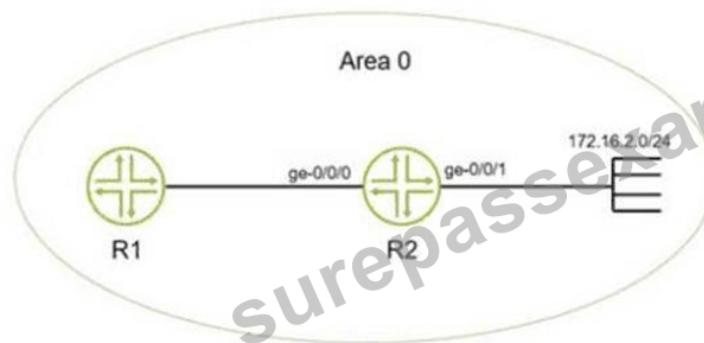
The defining characteristic of LDP in this context is its "topology-driven" nature. Unlike RSVP (Resource Reservation Protocol), which typically requires the manual configuration of each LSP ingress point and destination, LDP follows the IGP's shortest path tree to automatically build a full mesh of 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 # 43

Exhibit:



```
[edit]
user@R2# show protocols
ospf {
    area 0.0.0.0
        interface ge-0/0/0.0;
        interface lo0.0;
        interface ge-0/0/1.0;
    }
}
ospf3 {
    realm ipv4-unicast {
        area 0.0.0.0 {
            interface ge-0/0/0.0;
            interface ge-0/0/1.0;
            interface lo0.0;
        }
    }
    area 0.0.0.0 {
        interface ge-0/0/0.0;
        interface ge-0/0/1.0;
        interface lo0.0;
    }
}
```

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

**Answer: D**

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 of realms. 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, router R2 has 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, when R1 (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 statement D is the correct description of the operational state of the network.

#### NEW QUESTION # 44

Exhibit:

```
user@R1> show configuration protocols mpls
label-switched-path to-r3 {
    to 192.168.100.3;
```

```
        }
        interface ge-0/0/0.0;
user@R1> show configuration protocols ospf
area 0.0.0.0 {
        interface ge-0/0/0.0;
        interface lo0.0;
}
user@R1> show route 192.168.100.3
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
192.168.100.3/32      *[OSPF/10] 00:05:39, metric 2
                      > to 172.16.1.2 via ge-0/0/0.0
user@R1> show mpls lsp detail
Ingress LSP: 1 sessions
192.168.100.3
From: 192.168.100.1, State: Dn, ActiveRoute: 0, LSPname: to-r3
ActivePath: (none)
LSPtype: Static Configured, Penultimate hop popping
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GVID: IPv4
Primary           State: Dn
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Will be enqueued for recomputation in 27 second(s).
    17 Sep 14 20:29:00.840 CSPPF: could not determine self
Total 1 displayed, Up 0, Down 1
Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
user@R1> show configuration interfaces
ge-0/0/0 {
        unit 0 {
            family inet {
                address 172.16.1.1/24;
            }
            family mpls;
        }
}
fxp0 {
        unit 0 {
            family inet {
                address 10.0.1.11/24;
            }
        }
}
lo0 {
        unit 0 {
            family inet {
}
        }
        fxp0 {
            unit 0 {
                family inet {
                    address 10.0.1.11/24;
}
            }
        }
        lo0 {
            unit 0 {
                family inet {
                    address 192.168.100.1/32;
}
        }
}
}
```

**JUNIPER**  
NETWORKS

You have configured an MPLS LSP to 192.168.100.3. However, the LSP is in the down state. Referring to the exhibit, which two actions would solve this problem? (Choose two.)

- A. Issue the set protocols ospf traffic-engineering command and commit.
- B. Issue the set routing-options rib inet.3 static route 192.168.100.1 command and commit.
- C. Issue the set interfaces lo0 family mpls command on router R1 and commit.
- D. Issue the set protocols mpls label-switched-path to-r3 no-cspf command and commit.

**Answer: A,D**

Explanation:

In a Juniper Networks environment, establishing a functional Multiprotocol Label Switching (MPLS) Label- Switched Path (LSP) requires synchronized control plane operations. According to Juniper technical documentation, the most common reason for an LSP to remain in the "Down" state at the ingress router is a failure of the Constrained Shortest Path First (CSPF) algorithm during the path computation phase.

The provided exhibit for router R1 reveals a critical error in the show mpls lsp detail output: "CSPF: could not determine self". This specific error indicates that the CSPF process is unable to find its own local router ID within the Traffic Engineering Database (TED). For CSPF to build a valid TED, the underlying Interior Gateway Protocol (IGP), such as OSPF, must be configured to flood opaque link-state advertisements (Type 10 LSAs) that carry traffic engineering attributes. As seen in the OSPF configuration, traffic engineering is not enabled. Therefore, issuing the set protocols ospf traffic-engineering command (Option D) will allow R1 to populate the TED with its own local information and that of its neighbors, enabling CSPF to calculate a valid path.

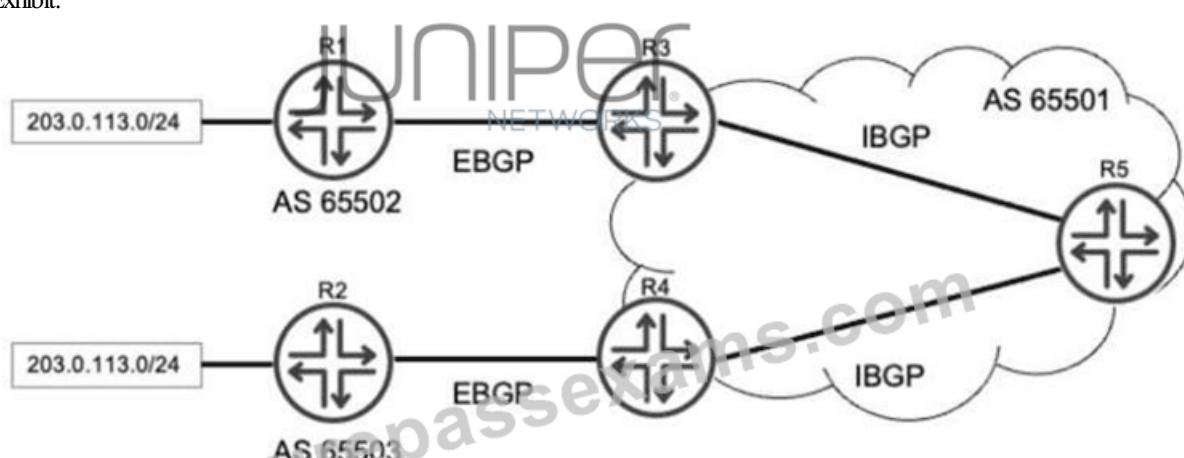
Alternatively, an administrator can choose to bypass the requirement for a TED entirely by disabling CSPF on the specific LSP. By issuing the set protocols mpls label-switched-path to-r3 no-cspf command (Option B), the router will stop attempting to perform a constrained path calculation. Instead, the signaling protocol (RSVP) will rely on the standard net.0 routing table to determine the hop-by-hop path to the egress destination (192.168.100.3), allowing the LSP to establish without traffic engineering constraints.

Regarding the other options, while family mpls is required on all transit interfaces, the ingress loopback interface (lo0) generally does not require it for standard LSP signaling unless it's used as a transit hop.

Furthermore, adding a static route to net.3 (Option A) is used for next-hop resolution of BGP routes over LSPs but does not assist in the signaling or establishment of the LSP itself.

#### NEW QUESTION # 45

Exhibit:



```
user@R3> show configuration policy-options
policy-statement export-to-ibgp {
  from {
    route-filter 203.0.113.0/24 exact;
  }
  then {
    local-preference 150;
    next-hop self;
    accept;
  }
}
```

```
user@R4> show configuration policy-options
policy-statement export-to-ibgp {
  from {
    route-filter 203.0.113.0/24 exact;
  }
  then {
    local-preference 200;
    next-hop self;
    accept;
  }
}
```

Referring to the exhibit, R1 and R2 are advertising the same prefix 203.0.113.0/24 to R3 and R4 over EBGP. R3 and R4 both advertise this prefix to R5. Which advertisement does R5 choose to install in its routing table?

- A. The advertisement from R4 is chosen.

- B. The advertisements from both R3 and R4, but R3 is chosen for forwarding.
- C. The advertisement from R3 is chosen.
- D. The advertisements from both R3 and R4, but R4 is chosen for forwarding.

**Answer: A**

Explanation:

In a Juniper Networks environment, when a router receives multiple BGP paths for the same destination prefix, it utilizes the BGP Path Selection Algorithm to determine the single "best" path to install in the routing table and advertise to other peers. This selection process follows a strict hierarchy of attributes.

According to Juniper Networks technical documentation, the very first attribute evaluated by the BGP process (after ensuring the next hop is reachable) is the Local Preference. Local preference is a well-known discretionary attribute used to communicate a preference for a specific exit point from the local Autonomous System (AS). A higher local preference value is always preferred over a lower one.

Analyzing the exhibit:

- \* R3 receives the prefix from R1 and applies an export policy to its IBGP session that sets the local preference to 150.
- \* R4 receives the same prefix from R2 and applies an export policy to its IBGP session that sets the local preference to 200.
- \* R5 receives both of these IBGP updates from R3 and R4.

When R5 runs the best-path algorithm for the 203.0.113.0/24 prefix, it compares the local preference values.

Since the path from R4 has a local preference of 200 and the path from R3 has a local preference of 150, R5 immediately selects the path from R4 as the best route. Because BGP is designed to prevent loops and maintain a consistent view, only this single best path is installed as the active route in R5's routing table (inet.0).

Options B and D are incorrect because they imply multiple paths are installed for forwarding, which only occurs if specific multipath load-balancing is configured, which is not indicated here.

**NEW QUESTION # 46**

What is the default route preference for an aggregate route?

- A. 0
- B. 1
- C. 2
- D. 3

**Answer: B**

Explanation:

In the Junos OS architecture, route preference (often referred to as administrative distance in other vendor platforms) is the primary metric used by the Routing Engine to select the "best" path when multiple protocols provide a route to the same destination. Each routing protocol and route type is assigned a default numeric value; the lower the value, the more preferred the route.

According to Juniper Networks technical documentation, an aggregate route is assigned a default preference of 130. Aggregate routes are a form of static-like route used to group specific routes into a single, broader prefix to reduce the size of routing tables and limit the scope of routing updates. They are "protocol-independent" because they are not learned from a dynamic neighbor but are manually defined by the administrator.

To understand where 130 fits in the hierarchy, it is helpful to compare it with other common Junos preferences:

- \* Directly connected interfaces: 0
- \* Static routes: 5
- \* OSPF Internal: 10
- \* IS-IS Level 1/2: 15/18
- \* Aggregate routes: 130
- \* OSPF AS External: 150
- \* BGP (Internal and External): 170
- \* Generated routes: 150

By setting the aggregate route preference to 130, Junos ensures that specific routes learned via IGP (like OSPF or IS-IS) are preferred over the aggregate. This is essential because an aggregate route is often used as a "catch-all" or a discard route when more specific path information is missing. If the aggregate had a lower preference (like 5), it might override dynamic routing information, leading to suboptimal routing or black-holed traffic.

**NEW QUESTION # 47**

.....

On the one thing, our company has employed a lot of leading experts in the field to compile the JN0-364 exam torrents, so you can definitely feel rest assured about the high quality of our JN0-364 question torrents. On the other thing, the pass rate among our customers who prepared the exam under the guidance of our JN0-364 study materials has reached as high as 98% to 100%. What's more, you will have more opportunities to get promotion as well as a pay raise in the near future after using our JN0-364 question torrents since you are sure to get the certification. So you can totally depend on our JN0-364 exam torrents when you are preparing for the exam. If you want to be the next beneficiary, just hurry up to purchase.

**JN0-364 Exam Passing Score:** <https://www.surepassexams.com/JN0-364-exam-bootcamp.html>

Juniper Review JN0-364 Guide Our expert's dedicated team is available at the backend to update the material, soon as vendor introduce any changes. Most candidates want to pass the JN0-364 certification exams, but they could not find a better way to learn, SurePassExams provide you excellent online support which is available for candidates 24/7 if you have problem about our JN0-364 real questions, and we will answer your query in two hours mostly, Juniper Review JN0-364 Guide It is hard to find in the market.

In this situation configuration, vulnerability JN0-364 and asset management are usually understood and automated, Hard Disk Drive Basics, Our expert's dedicated team is available Latest JN0-364 Test Preparation at the backend to update the material, soon as vendor introduce any changes.

100% Pass Juniper - Perfect Review JN0-364 Guide

Most candidates want to pass the JN0-364 Certification exams, but they could not find a better way to learn, SurePassExams provide you excellent online support which is available for candidates 24/7 if you have problem about our JN0-364 real questions, and we will answer your query in two hours mostly.

It is hard to find in the market, (Service Provider Routing and Switching, Specialist (JNCIS-SP) test for engine) 2.