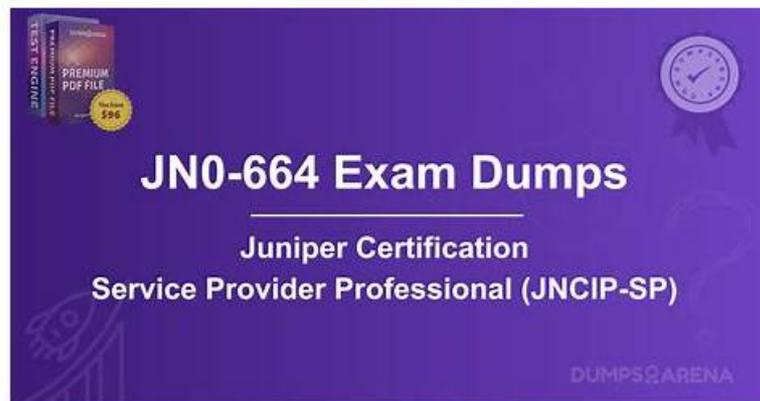# Juniper JN0-664 Exams Dumps, Exam JN0-664 Dumps



2025 Latest Free4Torrent JN0-664 PDF Dumps and JN0-664 Exam Engine Free Share: https://drive.google.com/open?id=1848GN69ZndbVZWyj2VCVeRDeHO_mcnCp

Scenarios of our Service Provider, Professional (JNCIP-SP) (JN0-664) practice tests are similar to the actual JN0-664 exam. You feel like sitting in the real JN0-664 exam while taking these Service Provider, Professional (JNCIP-SP) (JN0-664) practice exams. Practicing under these conditions helps you cope with Juniper JN0-664 Exam anxiety. Moreover, regular attempts of the JN0-664 practice test are also beneficial to enhance your speed of completing the final Service Provider, Professional (JNCIP-SP) (JN0-664) test within the given time.

With JN0-664 study tool, you are not like the students who use other materials. As long as the syllabus has changed, they need to repurchase learning materials. This not only wastes a lot of money, but also wastes a lot of time. Our industry experts are constantly adding new content to JN0-664 exam torrent based on constantly changing syllabus and industry development breakthroughs. We also hire dedicated staff to continuously update our question bank daily, so no matter when you buy JN0-664 Guide Torrent, what you learn is the most advanced. Even if you fail to pass the exam, as long as you are willing to continue to use our JN0-664 study tool, we will still provide you with the benefits of free updates within a year.

>> Juniper JN0-664 Exams Dumps <<

## Top JN0-664 Exams Dumps | High-quality Exam JN0-664 Dumps: Service Provider, Professional (JNCIP-SP) 100% Pass

Through years of efforts and constant improvement, our JN0-664 study materials stand out from numerous study materials and become the top brand in the domestic and international market. Our company controls all the links of JN0-664 study materials which include the research, innovation, survey, production, sales and after-sale service strictly and strives to make every link reach the acme of perfection. Our company pays close attentions to the latest tendency among the industry and the clients' feedback about our JN0-664 Study Materials.

The JN0-664 Exam is one of the most comprehensive exams offered by Juniper Networks. JN0-664 exam covers a wide range of topics, including protocol-independent routing, BGP, OSPF, ISIS, MPLS, Layer 2 VPNs, Layer 3 VPNs, multicast, and QoS. Candidates must have a deep understanding of these topics to pass the exam.

## Juniper Service Provider, Professional (JNCIP-SP) Sample Questions (Q76-Q81):

**NEW QUESTION # 76**
Exhibit
user@R1 show configuration interpolated-profile { interpolate {
fill-level [ 50 75 drop-probability [ > }
class-of-service drop-profiles
];
20 60 ];
Which two statements are correct about the class-of-service configuration shown in the exhibit? (Choose two.)

- A. To use this drop profile, you apply it directly to an interface.
- B. The drop probability gradually increases from 20% to 60% as the queue level increases from 50% full to 75% full
- C. The drop probability jumps immediately from 20% to 60% when the queue level reaches 75% full.
- D. To use this drop profile, you reference it in a scheduler.

**Answer: B,D**

Explanation:
Explanation
class-of-service (CoS) is a feature that allows you to prioritize and manage network traffic based on various criteria, such as application type, user group, or packet loss priority. CoS uses different components to classify, mark, queue, schedule, shape, and drop traffic according to the configured policies.
One of the components of CoS is drop profiles, which define how packets are dropped when a queue is congested. Drop profiles use random early detection (RED) algorithm to drop packets randomly before the queue is full, which helps to avoid global synchronization and improve network performance. Drop profiles can be discrete or interpolated. A discrete drop profile maps a specific fill level of a queue to a specific drop probability. An interpolated drop profile maps a range of fill levels of a queue to a range of drop probabilities and interpolates the values in between.
In the exhibit, we can see that the class-of-service configuration shows an interpolated drop profile with two fill levels (50 and 75) and two drop probabilities (20 and 60). Based on this configuration, we can infer the following statements:
* The drop probability jumps immediately from 20% to 60% when the queue level reaches 75% full. This is not correct because the drop profile is interpolated, not discrete. This means that the drop probability gradually increases from 20% to 60% as the queue level increases from 50% full to 75% full. The drop probability for any fill level between 50% and 75% can be calculated by using linear interpolation formula.
* The drop probability gradually increases from 20% to 60% as the queue level increases from 50% full to
75% full. This is correct because the drop profile is interpolated and uses linear interpolation formula to calculate the drop probability for any fill level between 50% and 75%. For example, if the fill level is
60%, the drop probability is 28%, which is calculated by using the formula: $(60 - 50) / (75 - 50) * (60 - 20) + 20 = 28$.
* To use this drop profile, you reference it in a scheduler. This is correct because a scheduler is a component of CoS that determines how packets are dequeued from different queues and transmitted on an interface. A scheduler can reference a drop profile by using the random-detect statement under the
[edit class-of-service schedulers] hierarchy level. For example: scheduler test { transmit-rate percent 10; buffer-size percent 10; random-detect test-profile; }
* To use this drop profile, you apply it directly to an interface. This is not correct because a drop profile cannot be applied directly to an interface. A drop profile can only be referenced by a scheduler, which can be applied to an interface by using the scheduler-map statement under the [edit class-of-service interfaces] hierarchy level. For example: interfaces ge-0/0/0 { unit 0 { scheduler-map test-map; } }

**NEW QUESTION # 77**
Exhibit

```
[edit routing-instances CE-1]
user@R1# show
protocols {
    bgp {
        group CE-1 {
            type external;
            peer-as 65555;
            neighbor 10.1.1.100;
        }
    }
}
instance-type vrf;
interface ge-0/0/2.0;
route-distinguisher 65512:1;
vrf-target target:65512:100;
[edit routing-instances CE-2]
user@R2# show
protocols {
    bgp {
        group CE-2 {
            type external;
            peer-as 64444;
            neighbor 10.1.5.100;
        }
    }
}
instance-type vrf;
interface ge-0/0/3.0;
route-distinguisher 65512:1;
vrf-target target:65512:100;
```

Referring to the exhibit, which statement is correct?

- A. The route-diatinguisher configuration will stop routes from being shared between CE-1 and CE-2.
- B. The route-distinguisher configuration will allow overlapping routes to be shared between CE-1 and CE-2.
- C. The vrf-target configuration will allow routes to be shared between CE-1 and CE-2.
- D. The vrf-target configuration will stop routes from being shared between CE-1 and CE-2.

**Answer: B**

Explanation:
The route distinguisher (RD) is a BGP attribute that is used to create unique VPN IPv4 prefixes for each VPN in an MPLS network. The RD is a 64-bit value that consists of two parts: an administrator field and an assigned number field. The administrator field can be an AS number or an IP address, and the assigned number field can be any arbitrary value chosen by the administrator. The RD is prepended to the IPv4 prefix to create a VPN IPv4 prefix that can be advertised across the MPLS network without causing any overlap or conflict with other VPNs. In this question, we have two PE routers (PE-1 and PE-2) that are connected to two CE devices (CE-1 and CE-2) respectively. PE-1 and PE-2 are configured with VRFs named Customer-A and Customer-B respectively.

## NEW QUESTION # 78
Exhibit.
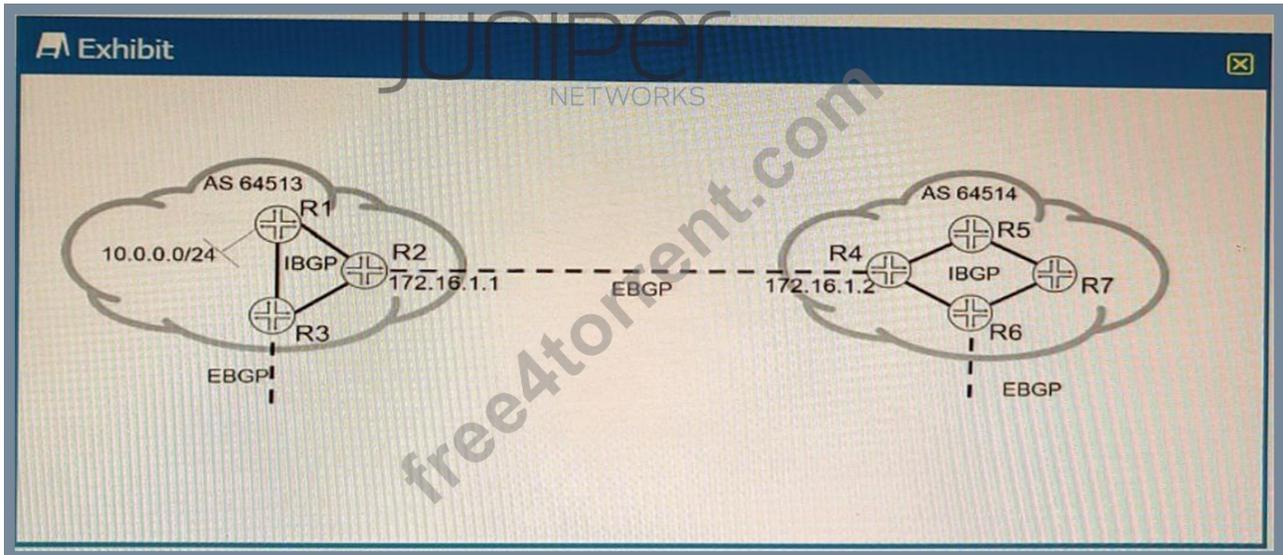


Referring to the exhibit; the 10.0.0.0/24 EBGP route is received on R5; however, the route is being hidden.
What are two solutions that will solve this problem? (Choose two.)

- A. On R4, create a policy to change the BGP next hop to 172.16.1.1 and apply it to IBGP as an export policy
- B. On R4, create a policy to change the BGP next hop to itself and apply it to IBGP as an export policy
- C. Add the external interface prefix to the IGP routing tables
- D. Add the internal interface prefix to the BGP routing tables.

**Answer: B,C**

Explanation:
Explanation
the default behavior for iBGP is to propagate EBGP-learned prefixes without changing the next-hop. This can cause issues if the next-hop is not reachable via the IGP. One solution is to use the next-hop self command on R4, which will change the next-hop attribute to its own loopback address. This way, R5 can reach the next-hop via the IGP and install the route in its routing table.
Another solution is to add the external interface prefix (120.0.4.16/30) to the IGP routing tables of R4 and R5.
This will also make the next-hop reachable via the IGP and allow R5 to use the route. According to 2, this is a possible workaround for a pure IP network, but it may not work well for an MPLS network.

## NEW QUESTION # 79
You have an L2VPN connecting two CEs across a provider network that runs OSPF. You have OSPF configured on both CEs.
Which two statements are correct in this scenario? (Choose two.)

- A. OSPF neighborship is formed between the CEs and PEs.
- B. The CE and PE OSPF areas must match.
- C. The CE and PE OSPF areas can be different.
- D. OSPF neighborship is formed between the two CEs.

**Answer: C,D**

Explanation:

In an L2VPN scenario, the provider network connects two customer edge (CE) devices across a Layer 2 virtual private network. Let's analyze how OSPF operates in this setup.

1. **OSPF Neighborship in L2VPN**:

- An L2VPN provides a Layer 2 connection between two sites, making it transparent to Layer 3 protocols like OSPF. This means the CEs can form OSPF adjacencies directly with each other as if they were on the same local network.

2. **OSPF Configuration on CEs and PEs**:

- **Statement A: OSPF neighborship is formed between the CEs and PEs**:

- Incorrect. In an L2VPN, the provider's network is transparent to the OSPF running on the CEs. OSPF neighborship forms directly between the CEs, not between the CEs and PEs.

- **Statement B: The CE and PE OSPF areas can be different**:

- Correct. Since OSPF adjacencies form directly between the CEs and not between CEs and PEs, the OSPF areas on the CEs and PEs can be different. The provider network acts as a transparent bridge, and OSPF doesn't see the PEs.

- **Statement C: The CE and PE OSPF areas must match**:

- Incorrect. As noted above, because the OSPF neighborship forms directly between the CEs, the OSPF areas on the CEs and PEs do not need to match.

- **Statement D: OSPF neighborship is formed between the two CEs**:

- Correct. The L2VPN makes the connection between the two CEs appear as a direct Layer 2 link, allowing them to form an OSPF adjacency directly.

**Conclusion**:

Given the above analysis, the correct statements are:

**B. The CE and PE OSPF areas can be different.**

**D. OSPF neighborship is formed between the two CEs.**

**References**:

- Juniper Networks Documentation on L2VPNs: [Configuring Layer 2 VPNs](https://www.juniper.net/documentation/en_US/junos/topics/task/configuration/layer-2-vpns-configuring.html)

- OSPF Configuration Guide: [Junos OS OSPF Configuration](https://www.juniper.net/documentation/en_US/junos/topics/concept/ospf-routing-overview.html)

**NEW QUESTION # 80**
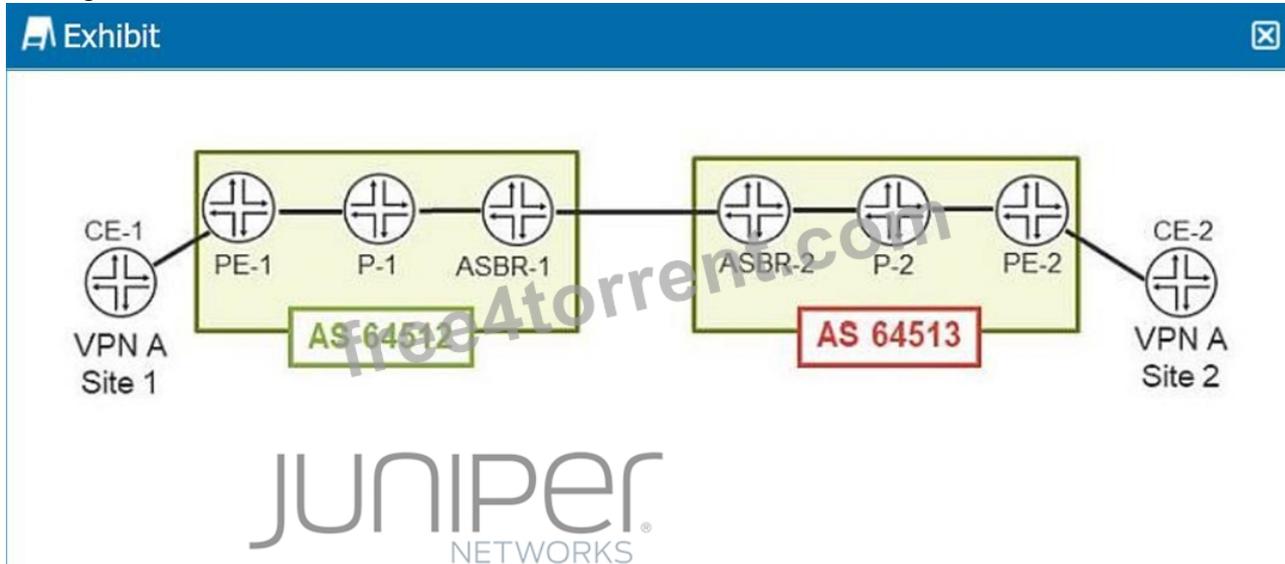


Click the Exhibit hutton.

You are configuring an interprovider Option C Layer 3 VPN to connect two customer sites.

Referring to the exhibit, which three statements are correct? (Choose three.)

- A. ASBR routers maintain the internal routes from its own AS, the loopback address from the other AS PEs, and the L3VPN routes.
- B. ASBR routers maintain the internal routes from its own AS and the loopback addresses from the other AS PEs.
- C. PE routers maintain the internal routes from its own AS, the loopback address from the other AS PEs, and the L3VPN routes.
- D. P routers only maintain the internal routes from their own AS.
- E. P routers maintain the internal routes from its own AS and the loopback address from the other AS PEs.

**Answer: B,C,D**

Explanation:

Interprovider Option C for Layer 3 VPNs involves the use of Autonomous System Boundary Routers (ASBRs) to exchange labeled VPN-IPv4 routes between different Autonomous Systems (AS). This option requires BGP sessions between ASBRs, and the VPN routes are carried end-to-end using MPLS labels. Here's a detailed analysis of the roles of different routers in this scenario:

1. **ASBR Routers**:
- ASBRs are responsible for exchanging VPN-IPv4 routes between different ASes.
- **A. ASBR routers maintain the internal routes from its own AS and the loopback addresses from the other AS PEs.**
- Correct. ASBRs maintain routes to internal destinations within their own AS, and they also need to know the loopback addresses of PEs in the other AS to set up the BGP sessions and MPLS tunnels.

2. **PE Routers**:
- PE routers are responsible for maintaining VPN routes and label information to forward VPN traffic correctly.
- **B. PE routers maintain the internal routes from its own AS, the loopback address from the other AS PEs, and the L3VPN routes.**
- Correct. PE routers need to maintain:
- Internal routes within their AS for routing.
- Loopback addresses of other AS PEs for establishing MPLS LSPs.
- L3VPN routes to provide end-to-end VPN connectivity.

3. **P Routers**:
- P routers are the core routers that do not participate in BGP VPN routing but forward labeled packets based on MPLS labels.
- **C. P routers only maintain the internal routes from their own AS.**
- Correct. P routers maintain the internal routing information to forward packets within the AS and use MPLS labels for forwarding VPN packets. They do not maintain VPN routes or routes from other ASes.

4. **Incorrect Statements**:
- **D. P routers maintain the internal routes from its own AS and the loopback address from the other AS PEs.**
- Incorrect. P routers do not need to maintain the loopback addresses of other AS PEs. They only maintain internal routing and MPLS label information.
- **E. ASBR routers maintain the internal routes from its own AS, the loopback address from the other AS PEs, and the L3VPN routes.**
- Incorrect. ASBR routers do not maintain L3VPN routes. They exchange labeled VPN-IPv4 routes with other ASBRs and forward them to PE routers.

**Conclusion**:

The correct answers are:

**A. ASBR routers maintain the internal routes from its own AS and the loopback addresses from the other AS PEs.**

**B. PE routers maintain the internal routes from its own AS, the loopback address from the other AS PEs, and the L3VPN routes.**

**C. P routers only maintain the internal routes from their own AS.**

**References**:
- Juniper Networks Documentation on Interprovider VPNs: [Interprovider VPN Configuration]
(https://www.juniper.net/documentation/en_US/junos/topics/topic-map/mpls-vpn-interprovider.ht
- MPLS and VPN Architectures, CCIP Edition by Ivan Pepelnjak and Jim Guichard

**NEW QUESTION # 81**

......

The site of Free4Torrent is well-known on a global scale. Because the training materials it provides to the IT industry have no-limited applicability. This is the achievement made by IT experts in Free4Torrent after a long period of time. They used their knowledge and experience as well as the ever-changing IT industry to produce the material. The effect of Free4Torrent's Juniper JN0-664 Exam Training materials is reflected particularly good by the use of the many candidates. If you participate in the IT exam, you should not hesitate to choose Free4Torrent's Juniper JN0-664 exam training materials. After you use, you will know that it is really good.