


JN0-364 Practice Exam - Exam JN0-364 Dump



No doubt the Juniper JN0-364 certification is a valuable credential that offers countless advantages to JN0-364 exam holders. Beginners and experienced professionals can validate their skills and knowledge level with the Service Provider Routing and Switching, Specialist (JNCIS-SP) JN0-364 Exam and earn solid proof of their proven skills.

As a market leader, our company is able to attract quality staffs on our JN0-364 exam materials , it actively seeks out those who are energetic, persistent, and professional to various JN0-364 certificate and good communicator. And we believe that the key of our company's success is its people, skills, and experience on JN0-364 Study Guide. Over 50% of the account executives and directors have been with the Group for more than ten years. We have strong strength to lead you to success!

>> JN0-364 Practice Exam <<

JN0-364 Exam Pdf Vce & JN0-364 Exam Training Materials & JN0-364 Study Questions Free

The up-to-date Juniper JN0-364 exam answers will save you from wasting much time and energy in the exam preparation. The content of our Juniper JN0-364 Dumps Torrent covers the key points of exam, which will improve your ability to solve the difficulties of Juniper JN0-364 real questions.

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

NEW QUESTION # 41

Which BGP attribute is optional, transitive, and is passed unchanged to other BGP peers if not recognized?

- A. AS Path
- B. MED
- C. Community
- D. Origin

Answer: C

Explanation:

BGP attributes are categorized into four distinct types based on how they are handled by a BGP speaker: Well-known mandatory, Well-known discretionary, Optional transitive, and Optional non-transitive.

Understanding these categories is essential for traffic engineering and ensuring consistent policy across an Autonomous System. According to Juniper Networks technical documentation, the **Community** attribute is classified as an optional transitive attribute. The term "optional" implies that a BGP implementation is not required to support or recognize the attribute. However, because it is "transitive," if a Juniper router receives an update containing a community tag that it does not recognize or has no specific policy for, it must accept the attribute and pass it along to other BGP peers unchanged. This ensures that community-based policies can be

signaled across intermediate ASes that may not be configured to act upon those specific tags.

In contrast:

* Origin (Option A) and AS Path (Option B) are well-known mandatory attributes. Every BGP update must include these, and every BGP-compliant router must recognize them.

* MED (Option D) (Multi-Exit Discriminator) is an optional non-transitive attribute. If a router receives a MED and advertises that route to an EBGP peer, the MED is typically stripped away (unless specific configurations like path-selection cisco-non-deterministic are used), as it is intended only to influence the immediate neighboring AS.

The Community attribute (defined in RFC 1997) is a powerful tool in Junos OS, often used for tagging routes to trigger specific routing policies, such as setting local preference or identifying the geographic origin of a prefix. By being transitive, it allows for sophisticated administrative control across complex multi-provider environments.

NEW QUESTION # 42

How are routing loops prevented in external BGP networks?

- A. Routing policies must be used to drop looped routes.
- B. Routing policies must be used to accept valid routes.
- **C. By default, a router receiving a route with its own AS in the AS Path attribute will not use the route.**
- D. By default, a router receiving a route with its own AS in the AS Path attribute will use the route.

Answer: C

Explanation:

BGP is a path-vector protocol, and its primary mechanism for ensuring a loop-free topology across the global internet is the AS_PATH attribute. This attribute is a "well-known mandatory" attribute that records every Autonomous System (AS) a prefix has passed through.

According to Juniper Networks Service Provider documentation, the loop prevention rule for External BGP (EBGP) is straightforward: when a router receives a BGP Update from an EBGP peer, it examines the AS_PATH list. If the router's own local AS number is already present in the list, it indicates that the advertisement has already traversed the local AS and has returned. To prevent a routing loop, the router will not use the route and will implicitly discard the update (Option D).

This behavior is a default, hard-coded function of the BGP protocol and does not require the administrator to write manual routing policies (Options B and C) to achieve basic loop prevention. While there are advanced features like as-path-expand or allow-as-in that can modify this behavior for specific design requirements (such as in certain Hub-and-Spoke MPLS VPN topologies), the standard operational default is to reject any route where the local AS is detected in the path. This ensures that traffic does not circulate infinitely between Autonomous Systems.

NEW QUESTION # 43

You are asked to configure a new network environment that will be based on IPv6 and use OSPF. In this scenario, which two statements correctly identify configuration task considerations? (Choose two.)

- **A. Participating interfaces are only required to be configured with the IPv6 protocol family and address.**
- B. The router ID used must be based on a 128-bit identifier value.
- **C. The router ID used must be based on a 32-bit identifier value.**
- D. Participating interfaces must be configured with both IPv4 and IPv6 protocol families and addresses.

Answer: A,C

Explanation:

When transitioning to an IPv6 environment using OSPFv3 (the version of OSPF designed for IPv6), there are significant architectural differences compared to OSPFv2 (IPv4). According to Juniper Networks technical documentation, OSPFv3 was redesigned to be more protocol-agnostic.

Router ID (Option C):

Despite OSPFv3 routing IPv6 (which uses 128-bit addresses), the OSPF Router ID remains a 32-bit value formatted like an IPv4 address (e.g., 1.1.1.1). This is a common point of confusion. In a pure IPv6 environment where no IPv4 addresses are configured on any interfaces, a Juniper router cannot automatically derive a Router ID. Therefore, the administrator must manually configure a 32-bit Router ID under [edit routing-options] for the OSPFv3 process to initialize.

Interface Configuration (Option D):

OSPFv3 runs directly over the IPv6 link-local scope. Unlike OSPFv2, it does not require an IPv4 address to function. Therefore, interfaces are only required to be configured with family inet6 (Option D). You do not need "dual-stack" (both IPv4 and IPv6) functionality just to run OSPFv3. The protocol uses the link-local address (fe80::/10) of the interface for neighbor adjacencies and as

the next hop for routing updates. This separation allows OSPFv3 to carry multiple "address families" (both IPv4 and IPv6 unicast) if needed, but the base requirement for an IPv6-only network is simply the family inet6 configuration.

NEW QUESTION # 44

A service provider is onboarding a new enterprise customer that operates multiple branch offices, each with its own set of VLANs. The customer requires transparent Layer 2 connectivity between sites while maintaining separation of internal VLANs. The provider must also ensure that customer VLAN identifiers do not conflict with other customers on the shared infrastructure. Which solution would provide the desired results?

- A. Extend customer VLANs using Q-in-Q tunneling.
- B. Provide Internet access with NAT and firewall services.
- C. Deliver Layer 3 VPN services using MPLS.
- D. Aggregate customer traffic using GRE tunnels.

Answer: A

Explanation:

In a service provider environment, Q-in-Q tunneling (also known as 802.1ad or double-tagging) is the standard solution for transporting multiple customer VLANs over a shared provider backbone while maintaining total separation.

According to Juniper Networks documentation, Q-in-Q works by adding a second 802.1Q tag (the Service Provider tag or S-tag) to the customer's already tagged frames (the Customer tag or C-tag). This creates a

"tunnel" at Layer 2. This solution specifically addresses all the customer's requirements:

* Transparent Layer 2 Connectivity: Because the provider simply encapsulates the customer's frames, the customer's internal BPDU traffic (like Spanning Tree) and VLAN tags are preserved and delivered transparently to the remote site.

* Separation of Internal VLANs: The customer can run their own internal VLAN IDs (1-4094) without the provider needing to know or manage them.

* Conflict Avoidance: Different customers on the same provider infrastructure are assigned unique S-tags. Even if two different customers both use "VLAN 10" internally, they remain isolated because their traffic is encapsulated in different provider S-tags.

Why other options are incorrect:

* Layer 3 VPN (Option B): While MPLS L3VPNs are common, they provide Layer 3 (IP) connectivity, not the "transparent Layer 2" connectivity requested.

* GRE Tunnels (Option C): GRE is a Layer 3 encapsulation and does not natively provide the transparent VLAN bridging required for a multi-site Layer 2 service.

* NAT/Firewall (Option D): These are security and address-translation services for internet access and do not facilitate site-to-site Layer 2 bridging.

NEW QUESTION # 45

How are routing loops prevented in internal BGP networks?

- A. Internal BGP routes are never readvertised to other external BGP neighbors.
- B. Internal BGP routes are never readvertised to other internal BGP neighbors.
- C. External BGP routes are never readvertised to other external BGP neighbors.
- D. External BGP routes are never readvertised to other internal BGP neighbors.

Answer: B

Explanation:

The prevention of routing loops within an Autonomous System (AS) is handled differently than loop prevention between ASes. While External BGP (EBGP) uses the AS_PATH attribute to detect loops, Internal BGP (IBGP) does not modify the AS_PATH. Therefore, a different mechanism is required to ensure that a route does not circulate infinitely inside the network.

This mechanism is known as the IBGP Split Horizon rule. According to Juniper Networks documentation and the BGP standard (RFC 4271), a BGP speaker must not advertise a route learned via an IBGP peer to any other IBGP peer. In simpler terms, "what is learned internally, stays local." This rule ensures that a route only travels one "hop" inside the AS—from the router that learned it from an external source to all other internal routers.

Because of this rule, IBGP routers do not naturally propagate routes through each other. This creates a requirement for a full mesh of IBGP sessions, where every BGP-speaking router in the AS must have a direct peering session with every other BGP-speaking router. To mitigate the scaling issues of a full mesh in large service provider networks, architects use Route Reflectors or Confederations, which are authorized exceptions to the Split Horizon rule.

Option B is incorrect because EBGP peers do not advertise EBGP routes to other EBGP peers (this is how the internet works). Option

C is incorrect because EBGP-learned routes must be sent to IBGP peers so the internal network knows how to reach the outside world. Option D is incorrect because internal routes must be sent to external peers to advertise your network to the internet.

NEW QUESTION # 46

To make sure your possibility of passing the certificate, we hired first-rank experts to make our JN0-364 practice materials. So the proficiency of our team is unquestionable. They help you to review and stay on track without wasting your precious time on useless things. By handpicking what the JN0-364 practice exam usually tested in exam and compile them into our JN0-364 practice materials, they win wide acceptance with first-rank praise. To go with the changing neighborhood, we need to improve our efficiency of solving problems as well as the new contents accordingly, so all points are highly fresh about in compliance with the syllabus of the exam.

Exam JN0-364 Dump: https://www.passcollection.com/JN0-364_real-exams.html

First, users can have a free trial of JN0-364 learning materials, to help users better understand the JN0-364 study materials, You can download the demo of JN0-364 study guide here to see if you really need it or not, Juniper JN0-364 Practice Exam Then after confirming, we will refund you, Juniper JN0-364 Practice Exam While each of them still hopes to pass the exam and get the certification.

It is very perfect for the project managers JN0-364 and project management aspirants. Examining the 'ImageView' control, First, users can have a free trial of JN0-364 Learning Materials, to help users better understand the JN0-364 study materials.

100% Pass Quiz Accurate Juniper - JN0-364 - Service Provider Routing and Switching, Specialist (JNCIS-SP) Practice Exam

You can download the demo of JN0-364 study guide here to see if you really need it or not. Then after confirming, we will refund you. While each of them still hopes to pass the exam and get the certification.

After downloading it also support offline operate.

