Get Free Juniper JN0-650 Dumps PDF Questions

Why risk failure? Download updated Juniper Enterprise Routing and Switching, Professional exam PDF questions today. Practice with real JN0-650 dumps and verified answers designed to help you ace your certification quickly using PrepBolt JN0-650 exam pdf questions and answers.

Thank you for Downloading JN0-650 exam PDF Demo

https://prepbolt.com/JN0-650.html

QUESTIONS & ANSWERS

DEMO VERSION

(LIMITED CONTENT)

Question 1

Question Type: MultipleChoice

Click the Exhibit button.

You are troubleshooting suboptimal routing in an enterprise OSPFv2 deployment:

-- R1 is an ABR between Area 0 and Area 1.

-- R2 is an ABR between Area 0 and Area 2.

-- R3 in Area 1 advertises a network that should be reachable by routers in Area 2.

-- The ABRs are connected in the backbone (Area 0).

On router R2, the show ospf database does not display router LSAs or network LSAs from Area 1, but an interarea route to R3's network exists in the routing table. R2 does not have Type-1 or Type-2 LSAs from Area 1 in its OSPF database.

Referring to the exhibit, which statement is correct?

Options:

A- R2 is filtering them using the OSPF import policy.

B- Junos disables Type-1 and Type-2 LSAs by default in Interarea LSDBs.

C- R2 is running OSPFv3, which does not use Type-1 or Type-2 LSAs.

D- Type-1 and Type-2 LSAs are summarized into Type-3 LSAs by ABRs and do not cross area boundaries.

Answer:

D

Explanation:

The behavior described is the fundamental operational principle of OSPF's hierarchical area structure.

LSA Flooding Scope: In OSPFv2, Type-1 (Router) and Type-2 (Network) Link-State Advertisements (LSAs) are restricted to a local area flooding scope. They contain detailed topology information (such as which routers are connected to which links) that is only relevant for calculating the Shortest Path First (SPF) tree within that specific area.

ABR Functionality: Area Border Routers (ABRs), such as R1 in this scenario, act as the gateway between areas. The ABR does not forward Type-1 or Type-2 LSAs from one area to another. Instead, it extracts the prefix information (the network address and mask) from those LSAs and generates a Type-3 (Summary) LSA.

Interarea Reachability: This Type-3 LSA is then flooded into the backbone (Area 0) and subsequently into other non-backbone areas, such as Area 2. This explains why R2 sees an 'interarea route' to R3's network in its routing table and has the corresponding Type-3 LSA in its database, but cannot see the original Type-1 or Type-2 LSAs from Area 1.

Why other options are incorrect: Option A is incorrect because this is standard protocol behavior, not the result of a manual filter. Option B is incorrect as there is no such 'disable' feature; the protocol design inherently prevents these LSAs from crossing boundaries. Option C is incorrect because OSPFv3 still uses Type-1 and Type-2 LSAs (renamed but functionally similar for topology), and the prompt specifically identifies the deployment as OSPFv2.

Question 2

Question Type: MultipleChoice

When configuring Q-in-Q tunneling, which type of tunneling involves the swapping of S-VLANs with C-VLANs?

Options:

A- L2PT

B- many-to-many

C- VLAN rewrites

D- all-in-one

Answer:

C

Explanation:

In a Juniper Q-in-Q (Layer 2 tunneling) environment, VLAN rewrites (specifically the swap operation) provide the most granular control over how customer traffic (C-VLANs) is mapped to service provider traffic (S-VLANs).

VLAN Rewrites (The Swap Operation): This method involves replacing the incoming customer VLAN tag with a service provider tag as the frame enters the tunnel. This is technically a 'swap' because the original C-VLAN tag is removed and the S-VLAN tag is written in its place. At the egress of the tunnel, the S-VLAN tag is swapped back for the original C-VLAN tag. This is often used when different customers use the same C-VLAN IDs and the provider needs to keep them unique within their core.

Many-to-Many: This is a mapping style where multiple customer VLANs are mapped to multiple service provider VLANs, but it typically relies on the 'push' (stacking) operation rather than a literal 'swap' of the tag itself.

All-in-One: This is the simplest form of Q-in-Q where all traffic entering an interface is 'pushed' into a single S-VLAN tag, regardless of any existing C-VLAN tags. No swapping occurs; the original tags are simply buried under the new provider tag.

L2PT (Layer 2 Protocol Tunneling): This is a feature used to tunnel Layer 2 control protocols (like STP, CDP, or LLDP) across a provider network by encapsulating them or changing their destination MAC addresses. It does not involve the swapping of VLAN tags.

Question 3

Question Type: MultipleChoice

The network you support currently has a mixture of MC-LAG and Virtual Chassis being used to provide redundant connectivity from various IDFs. A project to modernize the architecture and move to EVPN-VXLAN using ESI-LAG will be starting soon. You want to avoid IDFs losing connectivity as the core devices are migrated to EVPN-VXLAN. Which action will accomplish this task?

Options:

A- Enable EVPN-VXLAN on the leaf devices before starting the migration.
B- Remove the links between the MC-LAG leaf devices and disable Virtual Chassis.
C- Enable no-core-isolation on the core devices.
D- Enable network-isolation-profile on the leaf devices.

Answer:

C

Explanation:

In an EVPN-VXLAN environment using ESI-LAG (Ethernet Segment Identifier Link Aggregation), the Core Isolation feature is a safety mechanism designed to prevent traffic blackholing. When a leaf switch (acting as a VTEP) loses its BGP peering or its link to the IP fabric core, it assumes it is 'isolated' from the rest of the network. To protect the network, the switch automatically shuts down its local member links of any multi-homed ESI-LAG to force traffic to the peer switch that still has core connectivity.

However, during a migration or in specific transitional topologies where the core might be temporarily unreachable or not yet fully established, this feature can cause the leaf switches to shut down all downstream IDF (Intermediate Distribution Frame) connections, leading to a total loss of connectivity.

The Solution (Option C): By enabling the no-core-isolation statement under the [edit protocols evpn] hierarchy, you instruct the switch to disable this automatic shutdown behavior. This ensures that even if the BGP session or core links are not yet stable during the migration process, the ESI-LAG interfaces remain Up, allowing the IDFs to maintain connectivity to their local default gateways or other local resources.

Why others are incorrect: Enabling EVPN-VXLAN before migration (Option A) does not address the isolation logic. Removing MC-LAG/Virtual Chassis links prematurely (Option B) would cause an immediate outage. The network-isolation-profile (Option D) is typically used for different loop prevention scenarios and does not override the specific core-isolation check that affects multi-homed ESIs.

Question 4

Question Type: MultipleChoice

Which three conditions does the Junos OS use to create a figure-of-merit value for each BGP route in the routing table? (Choose three.)

Options:

A- The route is withdrawn.
B- The route has a configured local preference.
C- The route is readvertised.
D- The route is from IBGP or EBGP.
E- The route's path attributes changed.

Answer:

A, C, E

Explanation:

In Junos OS, the figure-of-merit is an internal numerical value used by the BGP Route Damping mechanism to quantify the stability of a specific BGP route. When a route 'flaps' (becomes unstable), the Junos OS increases this value by assigning a specific penalty based on the type of event that occurred:

The route is withdrawn (Option A): When a BGP neighbor sends a withdrawal message for a prefix, Junos assigns a significant penalty---typically 1,000 by default---to the route's figure-of-merit.

The route is readvertised (Option C): When a previously withdrawn route is announced again by the neighbor, the Junos OS adds another penalty to the figure-of-merit value to track the continued instability.

The route's path attributes changed (Option E): If a route remains reachable but its BGP attributes (such as the AS Path or Community) change, it is considered a flap. Junos assigns a smaller penalty---typically 500 by default---for this attribute change.

As the figure-of-merit accumulates and exceeds a predefined suppress threshold, the route is suppressed and no longer used for forwarding. Over time, the value decreases according to a configured half-life until it falls below a reuse threshold, at which point the route becomes active again. Factors like local preference or the origin (IBGP/EBGP) do not trigger these internal penalty increments for the figure-of-merit.

Question 5

Question Type: MultipleChoice

Exhibit

You need to configure the non-RP routers in your network to dynamically learn the location of the RP using standard protocols only. You have already applied the configuration shown in the exhibit, but the routers have not learned the location of the RP.

Which additional step must you take to fulfill these requirements?

Options:

A- Configure MSDP.
B- Configure BSR.
C- Configure Anycast RP.
D- Configure Auto-RP

Answer:

B

Explanation:

The exhibit shows a multicast network running PIM Sparse Mode where R2 is configured as a local Rendezvous Point (RP). However, the non-RP routers (like R5) show no learned RP in their show pim rps output. To resolve this using standard protocols for dynamic RP discovery, you must implement a mechanism that automates the distribution of RP information across the PIM domain.

Bootstrap Router (BSR) (Option B): BSR is the industry-standard (RFC 5059) mechanism for dynamically electing an RP and distributing that information throughout a PIM-SM domain.

It uses two roles: Candidate-RPs (C-RPs), which announce their desire to be an RP, and Candidate-BSRs (C-BSRs), which collect these announcements and flood them to all other routers in the network via PIM bootstrap messages.

BSR is a 'standard protocol' as it is part of the PIMv2 specification. Once configured on R2 (as C-RP) and potentially R1 or R2 (as C-BSR), all other routers will dynamically learn that R2 is the RP for the specified group ranges.

Auto-RP (Option D): While this also provides dynamic RP discovery, it is a proprietary Cisco protocol. Although Junos supports it for interoperability, the question specifically asks for 'standard protocols only,' making BSR the preferred choice.

MSDP (Option A): Multicast Source Discovery Protocol (MSDP) is used to share information about active sources between different PIM domains or for Anycast-RP. It does not handle the initial discovery of an RP within a single domain for standard PIM-SM operations.

Anycast RP (Option C): This is a technique used for RP redundancy and load balancing. While it can use BSR or MSDP for synchronization, it is a design architecture rather than a standalone 'standard protocol' for the basic dynamic discovery of an RP in the manner requested here.

Configuration Example for BSR in Junos OS 24.4: To enable this, you would add the following to the RP (R2):

set protocols pim rp bootstrap-discovery

set protocols pim rp candidate-rp 10.222.1.2 group-ranges 224.0.0.0/4

set protocols pim rp candidate-bsr 10.222.1.2

Question 6

Question Type: MultipleChoice

You run a multivendor switching environment where you have configured VSTP. You have 450 VLANs and notice that some of your VLANs do not function properly. How should you change the configuration to get all 450 VLANs working?

Options:

A- Include the force-version stp; statement in your configuration

B- Enable RSTP to handle additional VLANs.

C- Set the VLAN max-age to 3600 or more.

D- Increase the bridge priority on all VLANs to at least 16k.

Answer:

B

Explanation:

VSTP (VLAN Spanning Tree Protocol) is Juniper's implementation that provides a separate spanning tree instance for each VLAN, ensuring compatibility with Cisco's PVST+. However, it has significant scaling limitations:

Instance Limits: On many Juniper EX and QFX series switches, VSTP is restricted to a maximum of 253 or 510 VLAN instances depending on the software version (ELS vs. non-ELS). In your scenario, having 450 VLANs exceeds the standard 253-instance limit found on many platforms.

The Solution (Option B): When the number of VLANs exceeds the VSTP capacity, the recommended best practice is to enable RSTP (Rapid Spanning Tree Protocol). Unlike VSTP, RSTP runs a single spanning tree instance for the entire switch, regardless of how many VLANs are configured. This ensures that all 450 VLANs are protected from loops without hitting the hardware's instance-count limit.

Other Options: Option A (force-version) only affects the BPDU format for compatibility but doesn't solve the instance limit. Option C and Option D are parameter tuning actions that do not address the architectural limitation of the number of running instances.

Thank You for trying JN0-650 PDF Demo

To try our JN0-650 practice exam software visit link below

https://prepbolt.com/JN0-650.html

Start Your JN0-650 Preparation

Use Coupon “SAVE50” for extra 50% discount on the purchase of Practice Test Software. Test your JN0-650 preparation with actual exam questions.