DUMPSBOSS.

Enterprise Routing and SwitchingProfessional (JNCIP-ENT)

Juniper JN0-650

Version Demo

Total Demo Questions: 9

Total Premium Questions: 92

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Layer 2 Bridging and VLANs	21
Topic 2, Spanning Tree	4
Topic 3, Layer 2 Security	10
Topic 4, OSPF	15
Topic 5, IS-IS	2
Topic 6, BGP	18
Topic 7, IP Multicast	8
Topic 8, Class of Service	8
Topic 9, High Availability	2
Topic 10, Network Services	2
Topic 11, Routing Policy	2
Total	92

QUESTION NO: 1

You must implement EVPN signaling on an EX Series device that is configured with both underlay and overlay networks. Which network protocol accomplishes this task?

A. OSPF underlay network

B. IS-IS underlay network

C. MPLS overlay network

D. EBGP overlay network

ANSWER: D

Explanation:

EBGP overlay network is correct because EVPN uses Multiprotocol BGP (MP-BGP) as its control-plane signaling protocol. BGP advertises and distributes EVPN Network Layer Reachability Information (NLRI), including MAC/IP reachability and Ethernet segment information, between the EVPN virtual tunnel endpoints. On an EX Series EVPN-VXLAN fabric, the overlay BGP sessions carry these EVPN route advertisements, allowing VTEPs to learn remote endpoint locations and establish the forwarding information needed for VXLAN encapsulation.

Using external BGP for the overlay is a common Juniper fabric design because each leaf-to-spine connection can be a separate autonomous-system boundary. This model provides straightforward routing-policy control and naturally limits BGP advertisement scope according to the fabric topology. The IP underlay supplies routable connectivity between VTEP loopback addresses, while the BGP EVPN address family provides the signaling that communicates tenant Layer 2 and Layer 3 endpoint information across that underlay. EVPN's use of BGP for control-plane distribution is defined in RFC 7432; Juniper's EVPN documentation also describes BGP EVPN as the protocol used to exchange EVPN routing information: Juniper EVPN Overview.

QUESTION NO: 2

You are applying well-known BGP communities to control propagation of selected routes from your enterprise edge routers.

Which two statements are correct? (Choose two.)

A. A route tagged with no-advertise is not advertised to any BGP peer.

B. A route tagged with no-export can still be advertised to iBGP peers.

C. A route tagged with no-export is not installed in the local routing table.

D. A route tagged with no-advertise is advertised only to EBGP peers.

ANSWER: A B

Explanation:

The `no-advertise` community prevents a BGP speaker from advertising the route to any BGP peer. The `no-export` community allows advertisement within the local autonomous system, including to iBGP peers, but prevents advertisement to external peers outside the confederation boundary.

Neither community changes the route preference on the local router, and `no-export` does not prevent internal BGP distribution.

QUESTION NO: 3

Which two statements are correct about EVPN/VXLAN deployments? (Choose two.)

A. VNIs are used for encapsulating and de-encapsulating traffic entering and leaving the tunnels

B. VTEPs are used for encapsulating and de-encapsulating traffic entering and leaving the tunnels.

C. VTEPs are used to identify broadcast domains within an EVPN/VXLAN environment.

D. VNIs are used to identify broadcast domains within an EVPN/VXLAN environment.

ANSWER: B D

Explanation:

VTEPs are used for encapsulating and de-encapsulating traffic entering and leaving the tunnels. A VXLAN tunnel endpoint is the device function at the edge of the VXLAN overlay that takes an Ethernet frame received from a local endpoint, adds the outer IP/UDP and VXLAN encapsulation needed to traverse the Layer 3 underlay, and removes that encapsulation when traffic reaches the destination-side endpoint. On Juniper devices, this function enables Layer 2 traffic to be transported across an IP fabric while retaining the original Ethernet-frame context for the attached hosts.

VNIs are used to identify broadcast domains within an EVPN/VXLAN environment. The VXLAN Network Identifier is carried in the VXLAN header and identifies the logical overlay segment to which a frame belongs. In practical EVPN/VXLAN designs, a VNI is associated with a Layer 2 virtual network, commonly mapped to a VLAN or bridge domain, so that endpoints in the same overlay broadcast domain can be extended across the routed underlay. The VNI has a 24-bit value space, supporting a far larger number of isolated Layer 2 segments than traditional VLAN numbering. EVPN supplies the control-plane information used to advertise MAC/IP reachability and VTEP locations, while VXLAN uses the VNI to identify the overlay segment for forwarded traffic. See Juniper's EVPN-VXLAN overview and VTEP overview.

QUESTION NO: 4

You are deploying IP phones in your enterprise networks. When plugged in. the IP phones must automatically be provided with their geographic location details by the EX Series switches.

In this scenario, which protocol should be used to enable this behavior?

A. CDP

B. LLDP-MED

C. LLDP

D. MP-BGP

ANSWER: B

Explanation:

LLDP-MED is the appropriate protocol because it extends standard Link Layer Discovery Protocol capabilities specifically for media endpoint devices, including IP telephones. A switch running LLDP-MED can advertise network-policy information and location-identification information to a connected phone during link discovery. The location data is carried in LLDP-MED location identification TLVs and can represent civic-address, coordinate-based, or emergency-services location details, depending on the deployment requirements. This lets a newly attached phone learn its physical location automatically, which is particularly important for emergency calling and location-aware voice services. On Juniper EX Series switches, LLDP-MED is configured as part of the Ethernet-switching and network-management feature set, and the switch can advertise location information on access interfaces connected to endpoint devices. The endpoint must also support LLDP-MED and be configured to accept the advertised location TLVs. Juniper documents LLDP-MED as the mechanism for advertising endpoint-oriented information, including location identification, in its LLDP-MED documentation. The location formats and TLVs used by LLDP-MED are defined by the Telecommunications Industry Association's TIA-1057 LLDP-MED standard.

QUESTION NO: 5

You are troubleshooting OSPF external route selection. The routing domain contains both Type 1 and Type 2 external LSAs.

Which two statements are correct? (Choose two.)

A. A Type 1 external route includes the internal cost to reach the ASBR in its total cost.

B. A Type 2 external route uses the internal cost to the ASBR as its primary metric.

C. A Type 2 external route considers the internal cost to the ASBR when equal external metrics are compared.

D. A Type 1 external route is always preferred over a Type 2 external route regardless of metrics.

ANSWER: A C

Explanation:

For a Type 1 external route, OSPF adds the cost to reach the advertising ASBR to the external metric advertised in the LSA. This makes the total path cost significant when comparing Type 1 external routes learned through different ASBRs.

For a Type 2 external route, the external metric is the primary comparison value. The internal cost to reach the ASBR is considered only when candidate Type 2 routes have equal external metrics. Therefore, Type 2 routes are appropriate when the external metric should normally dominate the route-selection decision.

QUESTION NO: 6

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.)

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:

Junos OS uses the figure of merit as the accumulated penalty value for BGP route damping. The value represents route instability over time and is increased when a route experiences events that indicate it is flapping. “The route is withdrawn.” is one such event: removal of a previously usable route increases its penalty. “The route is readvertised.” also increases the penalty because a route returning after a withdrawal is part of the same instability pattern. “The route ' s path attributes changed.” identifies the third relevant type of event; Junos specifically treats a change to the AS path as a damping event and adds a penalty to the route’s figure of merit. The accumulated value decays according to the configured half-life. When it rises above the suppress threshold, Junos suppresses the route; once it decays below the reuse threshold, the route may again be used. This behavior prevents persistently unstable BGP reachability information from repeatedly affecting route selection and propagation. Juniper’s BGP route-damping documentation describes the figure-of-merit calculation, penalty events, decay, suppression, and reuse behavior. See Juniper BGP Route Damping documentation and the foundational RFC 2439.

QUESTION NO: 7

Exhibit.

Referring to the exhibit, you see that the 0.0.0.0/0 route is coming from AS 64512.

What is the command to achieve this task?

A. show route receive-protocol bgp < peer ip >

B. show route protocol bgp source-gateway < peer ip >

C. show route protocol bgp next-hop < peer ip >

D. show route receive-protocol bgp < peer ip > hidden

ANSWER: D

Explanation:

`show route receive-protocol bgp < peer ip > hidden` is the appropriate operational command when the objective is to verify a route advertised by a particular BGP peer that is rejected by the configured BGP import policy. Junos evaluates inbound BGP updates against import policy before installing eligible routes into the routing table. A default route such as `0.0.0.0/0` that is explicitly rejected is retained as a hidden route, allowing an operator to inspect it for troubleshooting even though it is not eligible for normal route selection or installation.

The `receive-protocol` form scopes the display to advertisements received from the specified BGP neighbor. Adding `hidden` is essential because it includes routes that are hidden as a result of policy processing or other route-resolution conditions. This lets the operator confirm both that the neighbor in AS 64512 advertised the default route and that the import policy is handling that advertisement as intended. Juniper documents the `show route receive-protocol` command, including its BGP-neighbor filtering and hidden-route display capability, in the Junos OS CLI Reference. Juniper's routing-policy documentation also describes how import policies control whether received routes are accepted into the routing table: Routing Policies Overview.

QUESTION NO: 8

You are asked to configure 802.1X on an EX Series switch port that connects to a downstream switch with several devices attached. You must ensure that the first device is authenticated while all subsequent devices are permitted.

What will accomplish this request?

A. single supplicant

B. MAC RADIUS

C. single-secure supplicant

D. multiple supplicant

ANSWER: A

Explanation:

single supplicant accomplishes this requirement. In single-supplicant mode, the EX Series switch treats the interface as having one supplicant for access-control purposes. The first connected device must complete 802.1X authentication successfully. Once that authentication succeeds, the interface is authorized and the switch permits traffic from the other devices reachable through that interface. This behavior is specifically useful when an EX switch port connects to a downstream device, such as another switch, and it is not practical or desired to authenticate every endpoint individually on the upstream port.

The configuration is applied under the interface's 802.1X authenticator settings using the single supplicant mode. The authenticated endpoint effectively establishes the authorized state for the port, allowing the attached downstream devices to communicate through it. This meets the stated requirement that authentication be performed for the first device while subsequent devices are allowed without requiring separate authentication exchanges on that upstream EX port.

Juniper documents the available 802.1X supplicant modes and their intended port-access behavior in its 802.1X Authentication documentation and the supplicant CLI statement reference.

QUESTION NO: 9

Exhibit

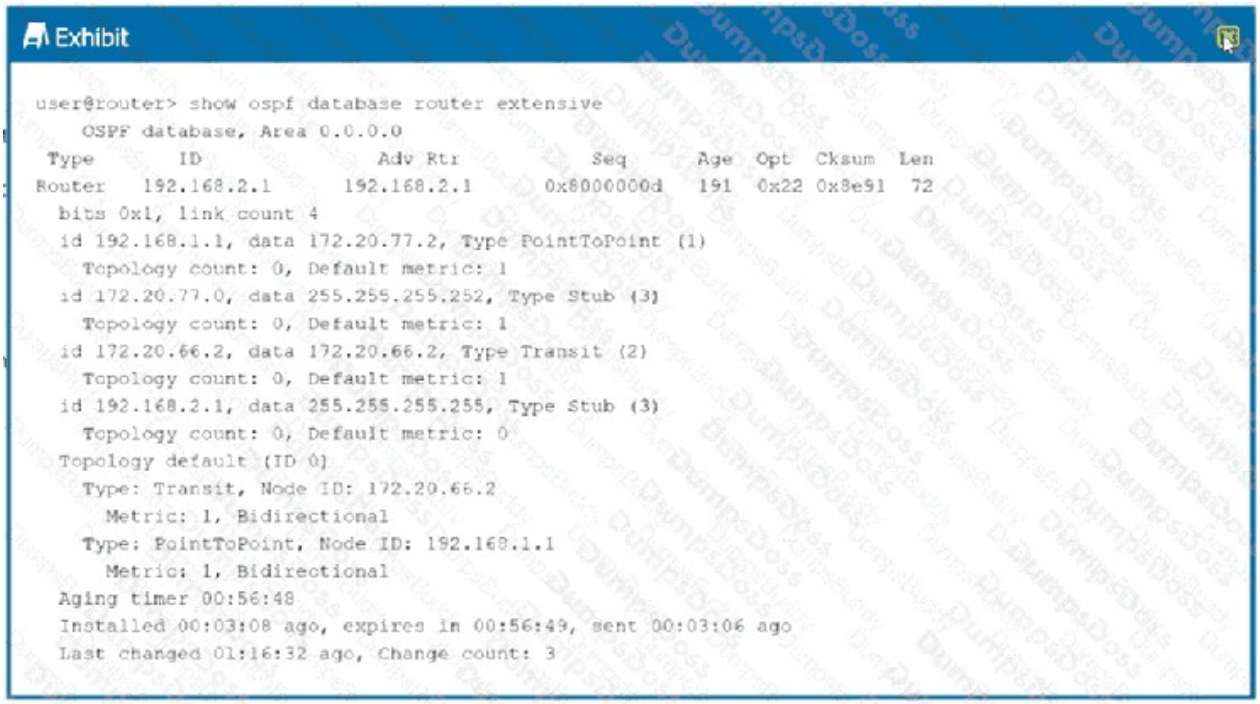

What is the OSPF router type shown in the exhibit for router 192.168.2.1?

A. neither an ABR nor an ASBR

B. an ASBR

C. both an ABR and an ASBR

D. an ABR

ANSWER: D

Explanation:

an ABR is correct because the Router-LSA for router 192.168.2.1 displays `bits 0x1`. The flags in an OSPF Type 1 Router-LSA describe specific capabilities of the originating router. A value of `0x1` sets the B bit (the area-border-router bit),

identifying the router as an Area Border Router. An ABR has interfaces in more than one OSPF area and maintains a separate link-state database for each attached area; it also participates in the backbone area either directly or through an appropriate backbone connection. The Router-LSA is flooded only within its originating area, but its flags allow other routers in that area to identify the originator's OSPF role. In this case, no additional flag is required to establish the ABR role: the set B bit is itself the protocol-defined indication. The OSPF specification defines the Router-LSA flag layout, including B as bit `0x01`, E as `0x02`, and V as `0x04`. Junos operational output exposes these values in the detailed Router-LSA database display. See RFC 2328, Router-LSA format and Juniper show ospf database documentation.