PassQueen
a.k.a dumpsqueen.com

Service Provider Routing and Switching, Specialist (JNCIS-SP)

Juniper JN0-363

Version Demo

Total Demo Questions: 12

Total Premium Questions: 126

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Routing Policy	10
Topic 2, OSPF	22
Topic 3, IS-IS	10
Topic 4, BGP	25
Topic 5, MPLS	29
Topic 6, Layer 2 VPNs	14
Topic 7, Layer 3 VPNs	3
Topic 8, Mix Questions	13
Total	126

QUESTION NO: 1

Exhibit

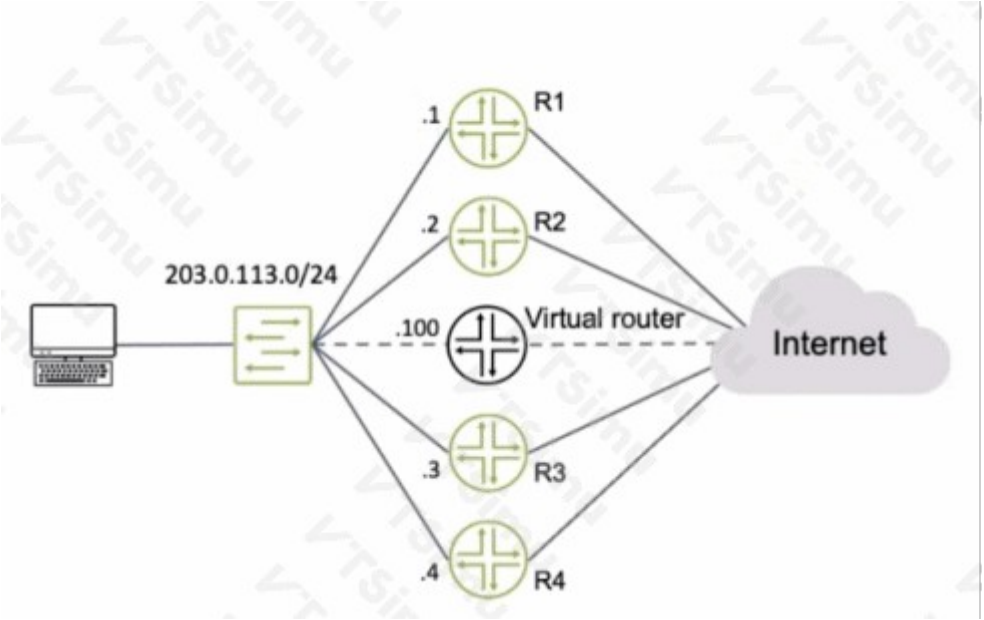

Routers R1 and R4 have a VRRP priority of 90, while R2 and R3 have default VRRP priorities

Referring to the exhibit, which router will be elected as the primary VRRP router?

A. R3

B. R4

C. R2

D. R1

ANSWER: A

Explanation:

R3 is elected as the primary VRRP router. VRRP elects the router advertising the highest priority as the primary router. The default VRRP priority is 100, so the participating routers that retain the default priority are preferred over routers configured with priority 90. When more than one router has the same highest priority, VRRP uses the primary-address tie-breaker: the router with the numerically highest IP address on the VRRP-enabled interface becomes the primary router. In the exhibit, R3 has the higher applicable interface IP address among the routers using the default priority, so it wins the election and advertises itself as the primary for the virtual router. This behavior ensures that the election is deterministic even when administrators have not explicitly assigned distinct priorities. Junos documents that VRRP priority controls primary-router selection, and the VRRP standard defines IP-address comparison as the tie-break mechanism for equal priorities. See Juniper's VRRP configuration documentation and RFC 5798, VRRP primary-router election.

QUESTION NO: 2

Exhibit

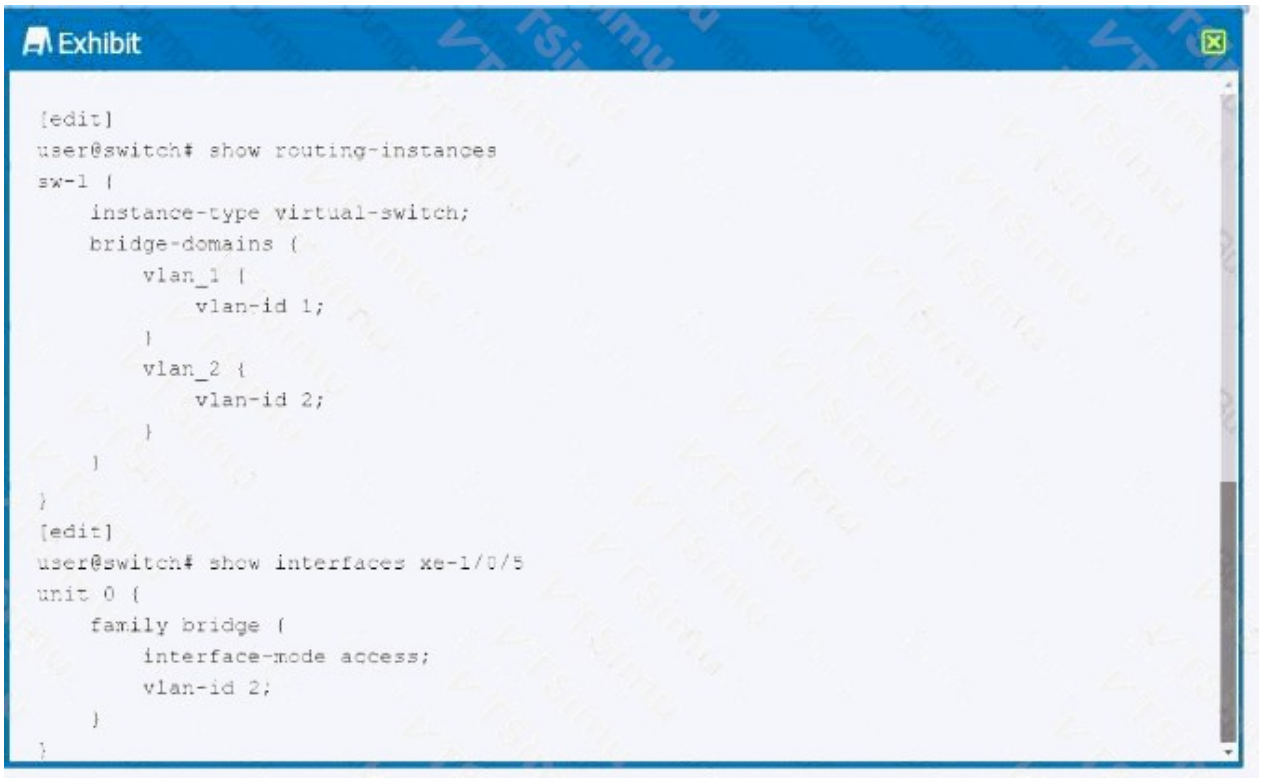

You are asked to assign interface xe-1/0/5 to a virtual switch.

What must be accomplished to complete the configuration?

A. Interface xe-1/0/5 must be added to routing-instance sw-1 vlan_2.

B. Interface xe-1/0'5 must be a trunk port.

C. Interface xe-1/0/5 must be added to routing-instance sw-1.

D. An IRB interface must be configured to routing-instance sw-1 vlan_2.

ANSWER: A

Explanation:

Interface xe-1/0/5 must be added to routing-instance sw-1 vlan_2. A virtual switch routing instance provides Layer 2 switching by defining VLANs within the routing instance and associating member interfaces with the applicable VLAN. The exhibit defines the virtual-switch instance named `sw-1` and includes the VLAN named `vlan_2`, using VLAN ID 2. Because the interface is configured for VLAN ID 2, it must be associated with that specific VLAN under the virtual-switch routing instance for traffic received on the interface to participate in the correct Layer 2 broadcast domain. In Junos configuration hierarchy, this is accomplished under the routing instance's VLAN configuration, typically by assigning the logical interface unit, such as `xe-1/0/5.0`, to `vlan_2`. The virtual switch can then apply its VLAN membership and switching behavior to the interface. Juniper documents that virtual-switch instances use the `vlans` hierarchy to configure VLANs and their interfaces within a routing instance. See the Juniper vlans statement reference for virtual-switch routing instances and the routing instances overview.

QUESTION NO: 3

Which two statements are correct when using LDP? (Choose two.)

A. The Inet.3 table will contain only the paths explicitly defined.

B. The inet.3 table will contain a full mesh of label-switched paths to other LDP-enabled routers.

C. LDP label-switched paths are created by configuring LDP on at least one physical router interface.

D. LDP label-switched paths are created by configuring LDP on the loopbacK Interface.

ANSWER: B C

Explanation:

When LDP is enabled throughout an MPLS domain, each LDP router distributes label bindings for reachable forwarding-equivalence classes and establishes LDP sessions with its directly connected LDP peers. As those bindings are learned and the IGP supplies the underlying next-hop reachability, Junos installs LDP-derived MPLS routes in the `inet.3` routing table. In a normally connected LDP domain, this produces label-switched reachability to the router IDs of the other LDP-enabled routers, effectively forming a full mesh of LDP transport LSPs. These LSPs are then available for use by MPLS services and protocols that resolve through `inet.3`, including RSVP-TE-related transport resolution where applicable.

LDP operation also requires LDP to be enabled on one or more transit-capable physical interfaces. Interface-level LDP configuration causes the router to send and receive LDP link hellos on those links, discover adjacent LDP peers, and establish the sessions through which label mappings are exchanged. The loopback address is commonly used as the LDP transport address and router identifier, but physical interfaces provide the adjacency discovery required to build the LDP signaling topology. Juniper's LDP configuration documentation describes enabling LDP under the protocol hierarchy and associating it with MPLS-enabled interfaces. See Juniper: Configuring LDP and Juniper: MPLS Routing Tables.

QUESTION NO: 4

Which configuration selling prohibits a static route from being redistributed by a dynamic routing protocol?

A. route-filter

B. no-readvertise

C. qualified-next-hop

D. passive

ANSWER: B

Explanation:

`no-readvertise` is the static-route configuration statement that prevents a route from being advertised, or redistributed, into dynamic routing protocols. It is configured under the relevant static route in the `routing-options` hierarchy. The route can still be installed in the local routing table and used for packet forwarding when it is active; the statement specifically controls propagation of that route to routing-protocol peers. This is useful when a static route is needed locally—for example, as a reachability or backup route—but must not be injected into protocols such as OSPF, IS-IS, or BGP through route redistribution policy. Junos documents `no-readvertise` as an option for static routes and describes its purpose as preventing readvertisement of the route by routing protocols. This behavior lets an operator separate local forwarding requirements from externally advertised routing information, helping maintain intended routing-domain boundaries and limiting unnecessary route distribution. See Juniper's static route configuration documentation and the static routing CLI statement reference.

QUESTION NO: 5

You are troubleshooting two OSPF routers that have an adjacency that remains in the ExStart state.

What would cause this problem?

A. mismatched OSPF hello intervals on the OSPF interfaces

B. mismatched authentication settings on the OSPF interfaces

C. mismatched MTU settings on the OSPF interfaces

D. mismatched subnet settings on the OSPF interfaces

ANSWER: C

Explanation:

mismatched MTU settings on the OSPF interfaces is correct because ExStart is the point at which two OSPF neighbors elect a master for database synchronization and begin exchanging Database Description (DBD) packets. A DBD packet includes the interface MTU value. By default, a router checks the MTU advertised by its neighbor before progressing with database exchange. If the received MTU is larger than the receiving interface MTU, the router cannot safely accept the neighbor's advertised packet size and the adjacency can remain stuck in ExStart (or sometimes Exchange) rather than reaching Full. This is a common interoperability and configuration issue on point-to-point links, particularly where one side has a different Ethernet, VLAN, tunnel, or logical-interface MTU. Verify the effective MTU on both OSPF-facing interfaces, including any encapsulation overhead, and configure matching values. Junos also provides an OSPF MTU-ignore capability for situations where it is operationally appropriate, but matching the interface MTUs is the preferred design because it avoids forwarding and fragmentation issues beyond OSPF adjacency formation. Juniper's OSPF documentation describes adjacency establishment and database synchronization behavior: Junos OSPF Overview. The OSPF standard specifies the Interface MTU field in Database Description packets: RFC 2328.

QUESTION NO: 6

You are configuring a Layer 2 VPN using BGP signaling.

Which two statements are correct about route targets? (Choose two.)

A. Route targets are BGP extended communities.

B. An import route target controls which received VPN routes are accepted into a routing instance.

C. A route target is used as the MPLS transport label for a VPN packet.

D. Route targets replace route distinguishers in VPN route advertisements.

ANSWER: A B

Explanation:

Route targets are BGP extended communities used to control VPN route distribution. A VPN routing instance exports its locally learned VPN routes with configured export route targets. It imports received VPN routes when their route target matches a configured import route target. This mechanism allows multiple VPN routing instances to share selected routes without requiring all instances to import the same routes. See RFC 4364.

QUESTION NO: 7

Exhibit

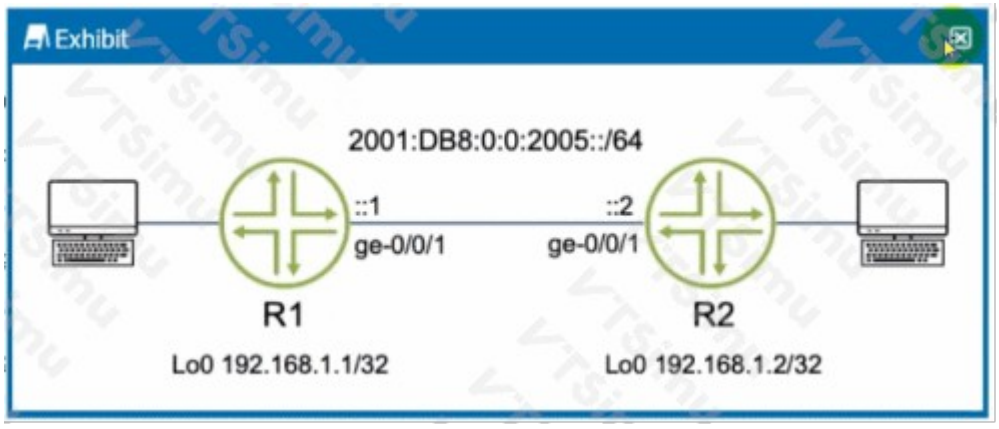

You are asked to configure OSPF between routers R1 and R2 using IPv6 addresses.

Which two tasks will accomplish your objective? (Choose two.)

A. Issue the sec protocols ospf area 0.0.0.0 interface ge-0/0/1.0 command.

B. Under the [edit routing-options] hierarchy, configure a 32-bit router ID.

C. issue the set protocols ospf3 area 0.0.0.0 interface ge-0/0/1.0 command.

D. Under the [edit routing-options] hierarchy, configure a 128-bit router ID.

ANSWER: B C

Explanation:

IPv6 OSPF adjacencies use OSPFv3, which Junos configures under the `protocols ospf3` hierarchy. Therefore, issuing `set protocols ospf3 area 0.0.0.0 interface ge-0/0/1.0` enables OSPFv3 on the specified logical interface and places it in the backbone area. OSPFv3 discovers neighbors and forms adjacencies using IPv6 link-local addressing on the enabled interface, so the participating interface must also have IPv6 addressing enabled as shown by the intended topology. Juniper's OSPFv3 configuration guidance uses this `ospf3` hierarchy for IPv6 unicast routing: Configuring OSPFv3.

Although OSPFv3 carries IPv6 routes, its router ID remains a 32-bit value written in IPv4 dotted-decimal notation. Configuring a 32-bit router ID under `[edit routing-options]` gives the router a stable, explicitly selected OSPF router identifier. This identifier is used in OSPF protocol operation and in the router-LSA; it is not derived from or represented as a 128-bit IPv6 address. Junos documents the routing-options router ID as an IPv4 address, including its use by OSPF: router-id (Routing Options).

QUESTION NO: 8

What are three well-known mandatory BGP attributes? (Choose three.)

A. next hop

B. origin

C. community

D. MED

E. AS path

ANSWER: A B E

Explanation:

The well-known mandatory BGP path attributes are **AS path**, **origin**, and **next hop**. "Well-known" means that every BGP implementation must recognize the attribute, while "mandatory" means that the attribute must be present in every BGP UPDATE message that advertises a reachable destination. The AS path records the autonomous systems a route advertisement has traversed. BGP uses it for loop prevention: a router rejects a route if its own autonomous system number already appears in the received path. The origin attribute indicates how the prefix was introduced into BGP, such as through an IGP, EGP, or an unspecified/incomplete source, and it is also considered during BGP best-path selection. The next hop attribute identifies the IP address that should be used to reach the advertised destination; the receiving router must be able to resolve and reach that next-hop address before the route can be installed and used for forwarding. These classifications are defined in the BGP base specification and are reflected in Junos BGP operational documentation. See RFC 4271, Section 5 and Juniper's BGP Path Attributes documentation.

QUESTION NO: 9

Exhibit

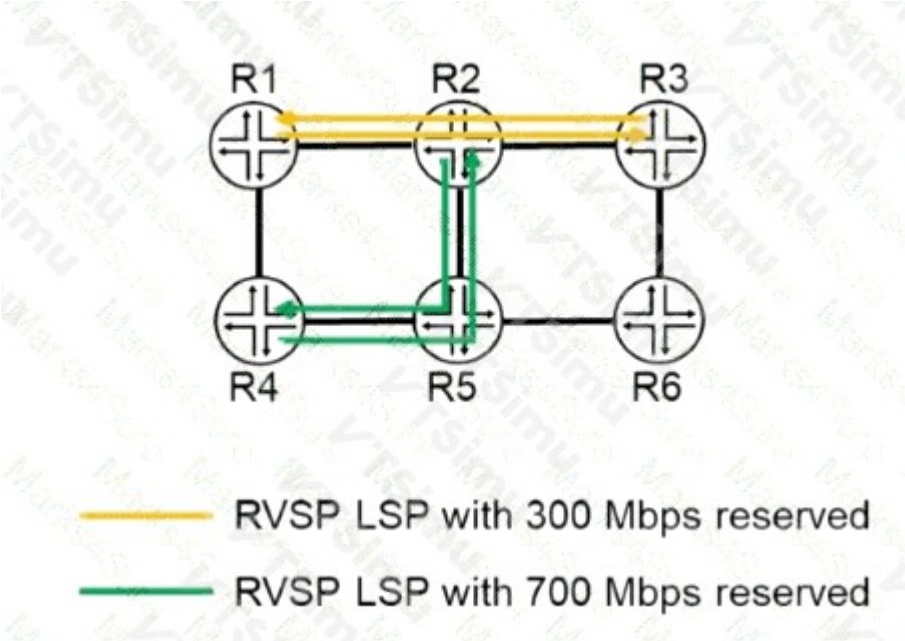

The exhibit shows a topology with 1 Gbps interfaces between routers, and four RSVP LSPs operating with the respective bandwidth reservations.

Which path will be selected for a new LSP from R4 to R6 with a bandwidth reservation of 400 Mbps?

A. R4 - > R1 - > R2 - > R5 - > R6

B. R4 - > R5 - > R6

C. R4 - > R5 - > R2 - > R3 - > R6

D. R4 - > R1 - > R2 - > R3 - > R6

ANSWER: A

Explanation:

R4 - > R1 - > R2 - > R5 - > R6 is selected because RSVP traffic engineering uses a constraint-based shortest-path calculation (CSPF) when establishing a bandwidth-reserved LSP. CSPF evaluates each candidate path using the TE link attributes advertised by the IGP, including the remaining reservable bandwidth. A 400 Mbps request can be signaled only across links that each have at least 400 Mbps of unreserved RSVP bandwidth available at the relevant priority. Based on the reservations shown in the exhibit, the links along R4 through R1, R2, and R5 to R6 satisfy that bandwidth constraint end to end. Consequently, this is the feasible TE path used for the new LSP. RSVP then signals the selected path with PATH and RESV messages, reserving the requested bandwidth on every traversed link so that subsequent CSPF calculations account for the new reservation. Junos RSVP-TE supports this bandwidth-aware path computation and admission-control behavior; see Juniper's RSVP Traffic Engineering documentation and the RSVP-TE signaling specification in RFC 3209.

QUESTION NO: 10

Exhibit

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

A. Traffic ingressing ge-0/0/0 that is tagged with VLAN 101 will egress ge-0/0/1 unchanged.

B. Traffic ingressing ge-0/0/0 that is tagged will VLAN 100 will be dropped.

C. Traffic ingressing ge0/0/0 that Is tagged with VLAN 200 will egress ge-0/0/1 with an outer VLAN tag of 200.

D. Traffic ingressing ge-0/0/0 that is tagged with VLAN 101 will egress ge-0/0M with an outer VLAN tag of 200.

ANSWER: A B

Explanation:

Traffic ingressing ge-0/0/0 that is tagged with VLAN 101 will egress ge-0/0/1 unchanged because VLAN 101 falls within the configured VLAN ID list of 101 through 120 on the ingress logical interface. With flexible VLAN tagging and flexible Ethernet services encapsulation, the router can classify frames using the configured VLAN ID or VLAN ID list while retaining the customer VLAN tag for bridged forwarding when no VLAN rewrite action is configured. Thus, the frame remains associated with its matching service and leaves with VLAN 101 intact. Traffic ingressing ge-0/0/0 that is tagged will VLAN 100 will be dropped because VLAN 100 is outside the accepted VLAN ID list. The vlan-id-list statement defines the VLAN tags that the logical unit accepts and uses for packet classification. A tagged frame whose VLAN ID is not included in that list does not match the logical unit, so it is not admitted to the associated forwarding service. Juniper documents VLAN ID-list matching at vlan-id-list and describes the interface capability enabled by flexible-vlan-tagging.

QUESTION NO: 11

Exhibit button

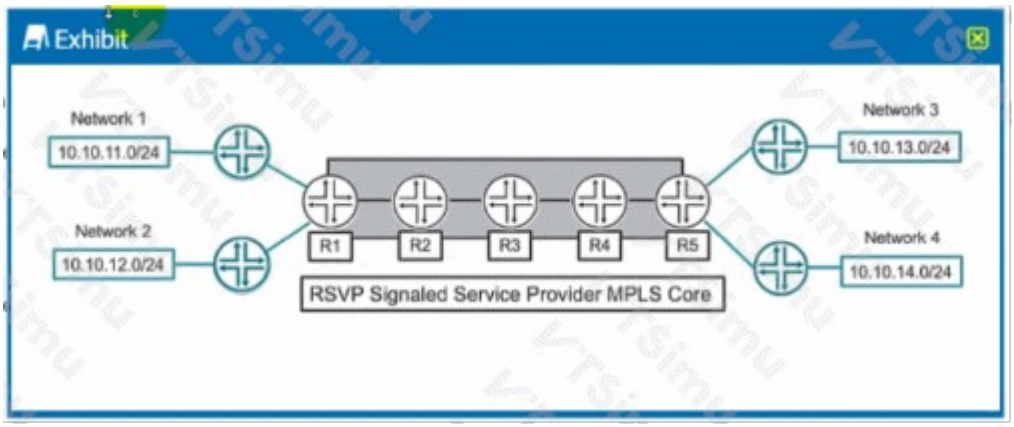

Which two statements are correct about the service provider MPLS network shown in the exhibit? (Choose two.)

A. R3 will perform a label pop operation on the transport MPLS label.

B. Traffic from Network 1 to Network 3 and traffic from Network 1 to Network 4 each need their own unique label-switched path.

C. Traffic from Network 1 to Network 3 and from Network 1 to Network 4 can share the same label-switched path.

D. R3 will perform a label swap operation on the transport MPLS label.

ANSWER: A C

Explanation:

“R3 will perform a label pop operation on the transport MPLS label” is correct because R3 is the penultimate provider router before the egress PE router in the depicted forwarding path. MPLS commonly uses penultimate-hop popping (PHP): the egress router advertises an implicit-null label for its loopback forwarding-equivalence class, causing its upstream neighbor to remove the outer transport label. This lets the egress PE receive the remaining service label, such as the VPN label, and use it directly to select the appropriate customer VRF and route the packet onward.

“Traffic from Network 1 to Network 3 and from Network 1 to Network 4 can share the same label-switched path” is also correct. An MPLS transport LSP is built toward the egress PE router, rather than requiring a separate transport LSP for every customer destination prefix. Where both destination networks are reached through the same egress PE, traffic can use the same outer transport label/LSP. The inner service label distinguishes the relevant VPN forwarding context and destination forwarding behavior after the packet reaches that PE. Juniper’s MPLS documentation describes label stacking and transport forwarding, while its L3VPN documentation describes use of an inner VPN label at the PE: Juniper MPLS Overview and Juniper Layer 3 VPNs.

QUESTION NO: 12

You are bringing a new network online with three MX Series devices enabled for STP. No root bridge priority has been configured. Which statement is true in this scenario?

A. The device with the lowest MAC address will be elected as the root bridge.

B. The device with the highest MAC address will be elected as the root bridge.

C. The device with the lowest numerical lo0 IP address will be elected as the root bridge.

D. The device with the highest numerical lo0 IP address will be elected as The bridge.

ANSWER: A

Explanation:

The device with the lowest MAC address will be elected as the root bridge. STP elects its root bridge by comparing Bridge IDs carried in BPDUs. A Bridge ID is composed of the configured bridge priority and a MAC-address-based system identifier. When no bridge priority is explicitly configured on any of the three devices, they use the same default STP bridge-priority value. Because that priority portion is equal, the MAC-address portion becomes the deciding tie-breaker: the numerically lowest MAC address produces the lowest Bridge ID and therefore wins the root-bridge election.

This behavior is fundamental to predictable STP operation. In a production topology, administrators normally configure a deliberately lower bridge priority on the intended root bridge rather than relying on hardware MAC-address ordering. Doing so ensures that a device replacement or topology expansion does not unexpectedly change the root bridge. Junos documents the STP `bridge-priority` setting and its role in selecting the root bridge in the bridge-priority CLI reference. Juniper's Spanning Tree Protocol overview also describes STP's loop-free topology and root-bridge election behavior.