Get Free Juniper JN0-352 Dumps PDF Questions

Why risk failure? Download updated Juniper Enterprise Routing and Switching, Specialist exam PDF questions today. Practice with real JN0-352 dumps and verified answers designed to help you ace your certification quickly using PrepBolt JN0-352 exam pdf questions and answers.

Thank you for Downloading JN0-352 exam PDF Demo

https://prepbolt.com/JN0-352.html

QUESTIONS & ANSWERS

DEMO VERSION

(LIMITED CONTENT)

Question 1

Question Type: MultipleChoice

You have traffic for a video streaming application traversing a GRE tunnel in your network and users are reporting poor performance. Both ends of the tunnel are using the default settings for a gigabit Ethernet interface, but you observe excessive packet drops due to exceeding the MTU.

Which two steps would you take to improve connectivity for this application? (Choose two.)

Options:

A- Increase the MTU for the member interfaces to 1524 bytes.
B- Route the GRE tunnel through a higher bandwidth link.
C- Configure the clear-dont-fragment option on the member interfaces.
D- Increase the MTU for the member interfaces to 1500 bytes.

Answer:

A, C

Question 2

Question Type: MultipleChoice

Which two statements describe OSPF DR and BDR behavior? (Choose two.)

Options:

A- The BDR becomes the DR if the DR fails.
B- The DR elects the area border router for each area.
C- The DR by default generates Type-2 Network LSAs to describe the multi-access segment.
D- The BDR by default generates Type-2 Network LSAs to describe the multi-access segment.

Answer:

A, C

Explanation:

The Backup Designated Router exists specifically to provide immediate failover redundancy for the Designated Router role on a multi-access broadcast or NBMA segment: the BDR maintains full adjacencies with every other router on the segment concurrently with the DR, precisely so that if the DR ever fails or is withdrawn, the BDR can be promoted directly to DR essentially instantaneously, without needing to wait through a fresh, full DR/BDR election process, and a new BDR election then takes place separately among the remaining DROther routers to fill the now-vacant backup role. This immediate promotion behavior confirms the first statement as correct. The Designated Router's other core responsibility on the segment is originating the Type 2 Network LSA, which describes the multi-access network itself as a pseudonode, listing every router attached to that segment; this LSA type exists specifically to avoid the full-mesh explosion of Router LSA adjacency listings that would otherwise be needed to describe a shared broadcast segment, and only the DR --- never the BDR, and never any DROther --- is responsible for generating and maintaining this particular LSA under normal, stable conditions, which confirms the third statement while directly ruling out the fourth. There is no such thing as the DR 'electing' an area border router; ABR status is instead a role a router acquires organically by having interfaces in more than one OSPF area, entirely independent of any DR/BDR election process on any individual segment. Reference topics: Junos Enterprise Routing -- OSPF, DR/BDR Roles and Type 2 Network LSA Origination.

Question 3

Question Type: MultipleChoice

Which statement describes how Rapid Spanning Tree Protocol (RSTP) identifies an alternate port?

Options:

A- An alternate port is the upstream port that is not receiving BPDUs from the root bridge.
B- An alternate port is the upstream port that provides the least-cost path to the root bridge.
C- An alternate port is the upstream port that provides a backup path to the root bridge and stays in the discarding state.
D- An alternate port is the upstream port that forwards traffic only when the designated port is overloaded.

Answer:

C

Explanation:

RSTP, as defined in IEEE 802.1w, introduces explicit, well-defined port roles that go beyond legacy 802.1D STP's simpler root/designated/blocking model, and the alternate port role is one of RSTP's key refinements. An alternate port is a port that receives BPDUs from a bridge other than the one through

which the local switch's actual root port reaches the root bridge --- in other words, it represents a second, redundant path toward the root bridge learned from a different upstream bridge than the one currently providing the best (root port) path. Because the existing root port already provides the lowest-cost, actively used path to the root, the alternate port is deliberately held in the discarding (non-forwarding) state during normal, stable operation, functioning purely as a pre-computed, ready-to-use backup; critically, RSTP's major performance advantage over legacy STP is that if the current root port fails, the switch can transition its alternate port directly and almost instantly into the forwarding root port role without needing to pass through the lengthy listening and learning timer-based states that classic STP required, since the alternate port's suitability as a backup was already continuously verified through ongoing BPDU reception. It is not simply 'not receiving BPDUs' (an alternate port receives BPDUs continuously, just inferior ones relative to the current root port), it is not itself the least-cost path (that describes the root port), and it does not forward based on designated-port overload, which is not a concept that exists in RSTP's port-role logic at all. Reference topics: Junos Enterprise Switching -- Spanning Tree Protocols, RSTP Port Roles: Alternate and Backup Ports.

Question 4

Question Type: MultipleChoice

[Exhibit]

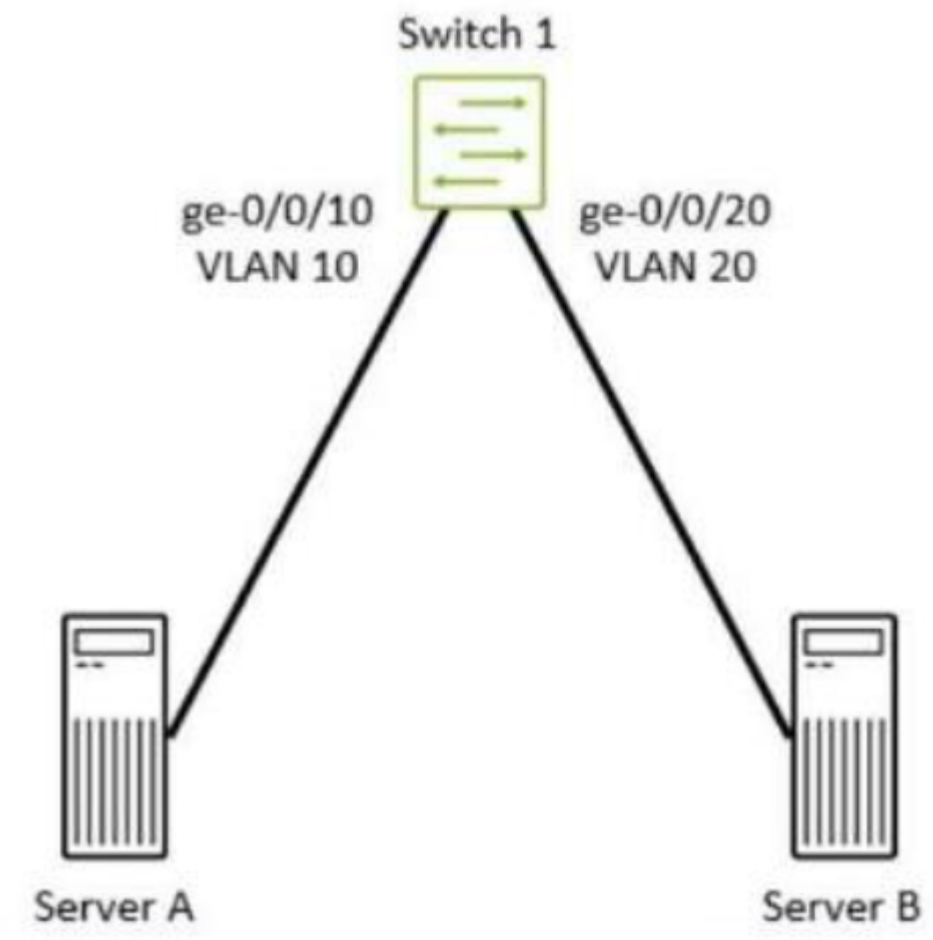

Click the Exhibit button.

Referring to the exhibit, all possible firewall filters are configured along the path from Server A to Server B.

In which order will Switch 1 process transmit (Tx) and receive (Rx) firewall filters?

Options:

A- Rx-Port --> Rx-VLAN --> Rx-Router --> Tx-Router --> Tx-VLAN -> Tx-Port
B- Rx-Port --> Rx-VLAN --> Rx-Router --> Tx-VLAN --> Tx-Port
C- Rx-Port --> Rx-VLAN --> Tx-VLAN --> Tx-Port
D- Rx-Port --> Rx-Router --> Tx-Router --> Tx-Port

Answer:

A

Explanation:

Because Server A resides in VLAN 10 and Server B resides in VLAN 20, traffic between them cannot remain a purely Layer 2 switched flow; it must be routed between the two VLANs, which on Switch 1 means the packet traverses the switch's internal routing engine (via an IRB interface) between the ingress and egress switching stages. Junos processes firewall filters along this path in a strict, layered order that mirrors the packet's actual journey through the device. First, any filter applied at the physical ingress port level (Rx-Port) is evaluated as the frame first arrives on ge-0/0/10. Next, any filter applied to the ingress VLAN (Rx-VLAN, i.e., VLAN 10) is evaluated as the frame is classified into its source broadcast domain. Because the destination subnet requires routing, the packet is then handed to the routing engine, where any input/ingress routed filter (Rx-Router) applied to the receiving IRB interface is evaluated next, followed by any output/egress routed filter (Tx-Router) applied as the now-routed packet is handed from the routing engine back out toward VLAN 20. Finally, the packet passes through any egress VLAN filter (Tx-VLAN, for VLAN 20) and then the egress physical port filter (Tx-Port) on ge-0/0/20 immediately before transmission to Server B. This complete six-stage sequence --- Rx-Port, Rx-VLAN, Rx-Router, Tx-Router, Tx-VLAN, Tx-Port --- reflects every filter-evaluation checkpoint a routed, inter-VLAN packet passes through on an EX Series Layer 3 switch. Reference topics: Junos Enterprise Switching -- Layer 2 Switching and Firewall Filter Processing Order for Routed VLAN Traffic.

Question 5

Question Type: MultipleChoice

You are troubleshooting a Layer 2 topology issue in a campus network. The switches are running RSTP. You need to verify which switch is currently acting as the root bridge and identify the root port on the local device.

Which operational command should the administrator use to view this information?

Options:

A- show spanning-tree interface
B- show spanning-tree bridge
C- show ethernet-switching table
D- show ethernet-switching interface

Answer:

B

Explanation:

The show spanning-tree bridge command is specifically designed to surface exactly the two pieces of information required in this scenario within a single output. Its Root ID field reports the full Bridge Identifier (priority and MAC address) of whichever switch has been elected root bridge for the spanning-tree instance in question, immediately answering which switch in the campus network is currently acting as root, regardless of whether that happens to be the local device or a remote one. In the same output, the command's Root port field explicitly names the specific local interface (for example, ge-0/0/1.0) that the local switch has selected as its path toward that root bridge, directly answering the second part of the requirement without needing to cross-reference a separate command. Together, these two fields give a complete, immediate picture of both root bridge identity and local root port selection from one operational command. show spanning-tree interface instead presents a per-interface breakdown of STP role (root, designated, alternate, backup) and state across all local ports, which can indirectly reveal the root port by looking for the 'Root' role, but it does not directly report the Root ID/root bridge identity in the same structured way that show spanning-tree bridge does. show ethernet-switching table and show ethernet-switching interface pertain entirely to MAC address learning and switching-specific interface configuration respectively, and neither surfaces any spanning-tree topology or root bridge information at all. Reference topics: Junos Enterprise Switching -- Spanning Tree Protocols, Verifying Root Bridge and Root Port with show spanning-tree bridge.

Question 6

Question Type: MultipleChoice

Which two characteristics describe the default behavior of aggregate routes? (Choose two.)

Options:

A- They advertise all contributing routes individually by default.
B- They hide internal routing instability from external peers.
C- They use a default next hop of discard.
D- They use a default next hop of reject.

Answer:

B, D

Explanation:

The entire purpose of route aggregation is summarization --- replacing a set of more-specific contributing routes with a single, broader prefix advertised outward --- and this summarization inherently hides internal routing instability from external peers: if one of the underlying contributing /24 blocks within a summarized /16 flaps up and down due to a local link issue, external peers who only see the stable, unchanging /16 aggregate are completely insulated from that churn, since the aggregate itself remains active and advertised as long as at least one contributing route is present, regardless of how much internal fluctuation occurs among the individual contributors. This stability-hiding characteristic is one of the primary operational benefits that motivates using aggregate routes in the first place, confirming that statement as correct. Regarding next-hop behavior, official Juniper documentation is explicit and unambiguous: when an aggregate route is installed in the routing table, Junos assigns it a reject next hop by default, meaning traffic matching only the aggregate (and no more-specific contributor) is dropped and an ICMP unreachable message is returned to the source; a discard next hop is available only as an optional, explicitly configured alternative when silent dropping without ICMP notification is preferred, but it is not the default. Aggregate routes categorically do not advertise their individual contributing routes; suppressing those specific, more-granular prefixes from onward advertisement in favor of the single summary is the aggregate's defining function, directly contradicting the first statement. Reference topics: Junos Enterprise Routing -- Protocol Independent Routing, Default Reject Next Hop and Route Summarization Benefits.

Question 7

Question Type: MultipleChoice

You implement FBF on router R1 so that traffic from subnet 172.25.0.0/24 uses ISP-A and traffic from subnet 172.25.1.0/24 uses ISP-B. You create forwarding instances for ISP-A and ISP-B. You also configure static default routes inside each instance. However, the static default routes remain inactive.

In this scenario, which action would complete the FBF implementation?

Options:

A- Configure an RIB group that includes the routing instances and inet.0.
B- Apply a firewall filter under the firewall, but not to any interface.
C- Increase the route preference of the static routes.
D- Change the instance type to virtual router.

Answer:

A

Explanation:

A classic and well-documented gotcha in filter-based forwarding deployments is that the static default route configured inside each forwarding instance frequently references a next-hop address that Junos cannot resolve, simply because the interface and other directly connected routes needed to validate that next hop as reachable exist only in the master inet.0 table by default, not automatically inside the newly created forwarding instance's own table. Since a static route's next hop must be resolvable against routes present within its own routing table to become active, the static default routes inside the ISP-A and ISP-B forwarding instances remain inactive precisely because the necessary interface/next-hop routes were never made visible inside those instances. The standard solution taught for this exact FBF scenario is to configure a routing information base (RIB) group that explicitly imports the relevant interface routes between the master inet.0 table and each forwarding instance's table, allowing the static default route's next hop inside each instance to resolve correctly against the leaked routes and become active. Applying a firewall filter without ever attaching it to an interface accomplishes nothing, since an unapplied filter never evaluates any traffic. Increasing the static route's preference value only affects which competing route wins selection and does nothing to resolve an unreachable next hop. Changing the instance type away from forwarding (to virtual-router) would in fact break FBF entirely, since filter-based forwarding specifically depends on the forwarding instance type working together with a routing-instance firewall filter action. Reference topics: Junos Enterprise Routing -- Protocol Independent Routing, Filter-Based Forwarding and RIB Group Route Leaking.

Thank You for trying JN0-352 PDF Demo

To try our JN0-352 practice exam software visit link below

https://prepbolt.com/JN0-352.html

Start Your JN0-352 Preparation

Use Coupon “SAVE50” for extra 50% discount on the purchase of Practice Test Software. Test your JN0-352 preparation with actual exam questions.