

Juniper

Exam Questions JN0-280

Data Center Associate (JNCIA-DC)



NEW QUESTION 1

What are three correct layer names used in legacy hierarchical network design? (Choose three.)

- A. Access layer
- B. Modular layer
- C. Aggregation layer
- D. Core layer
- E. Function layer

Answer: ACD

Explanation:

In legacy hierarchical network design, three key layers are used to create a scalable and structured network:

Step-by-Step Breakdown:



Access Layer:



The access layer is where end devices, such as computers and IP phones, connect to the network.

It typically involves switches that provide connectivity for devices at the edge of the network.



Aggregation Layer (Distribution Layer):



The aggregation layer (also called the distribution layer) aggregates traffic from multiple access layer devices and applies policies such as filtering and QoS.

It also provides redundancy and load balancing.



Core Layer:



The core layer provides high-speed connectivity between aggregation layer devices and facilitates traffic within the data center or between different network segments.

Juniper Reference:



Legacy Hierarchical Design: Juniper networks often follow the traditional three-layer design (Access, Aggregation, and Core) to ensure scalability and high performance.

NEW QUESTION 2

When a MAC limiting violation occurs, the switch performs which two actions by default? (Choose two.)

- A. No logging takes place.
- B. It causes Layer 2 loops.
- C. The port is disabled.
- D. It drops the packet.

Answer: CD

Explanation:

When a MAC limiting violation occurs on a Juniper switch, the switch will perform the following actions by default:

Step-by-Step Breakdown:



Port Disabled: When the number of MAC addresses on an interface exceeds the configured limit, the port is automatically disabled to prevent further violations.

This is a protective mechanism to prevent MAC address flooding.



Packet Dropped: Additionally, packets from the violating MAC address are dropped to prevent any further communication from that address. This ensures that only valid MAC addresses are allowed to communicate through the interface.



Example Configuration:

```
set ethernet-switching-options secure-access-port interface <interface-name> mac-limit 5
```



If more than five MAC addresses are learned, the port is disabled, and excess packets are dropped.

Juniper Reference:



MAC Limiting: When the switch detects a MAC limiting violation, it disables the port and drops further packets from the violating MAC addresses to maintain network security.

NEW QUESTION 3

Which state in the adjacency process do OSPF routers check the MTU size?

- A. Init
- B. Exchange
- C. Done
- D. ExStart

Answer: B

Explanation:

In OSPF, routers exchange link-state information in different stages to establish full adjacency. The MTU size is checked during the Exchange state.

Step-by-Step Breakdown:



OSPF Adjacency Process:

- OSPF routers go through multiple stages when forming an adjacency:Down,Init,2-Way,ExStart,Exchange,Loading, andFull.
 - Exchange State:
 - During theExchangestate, OSPF routers exchangeDatabase Description (DBD)packets to describe their link-state databases. TheMTU sizeis checked at this stage to ensure both routers can successfully exchange these packets without fragmentation.
 - If there is anMTU mismatch, the routers may fail to proceed past the Exchange state.
- Juniper Reference:
- MTU Checking in OSPF: Junos uses the Exchange state to check for MTU mismatches, ensuring that routers can properly exchange database information without packet fragmentation issues.

NEW QUESTION 4

Which statement is correct about IBGP?

- A. It requires a physical full mesh.
- B. It requires a logical full mesh.
- C. It ensures that the local and remote peers use different AS numbers.
- D. It ensures that duplicate AS numbers are not present in the AS path.

Answer: B

Explanation:

InIBGP (Internal Border Gateway Protocol), all routers within the same AS (Autonomous System) must have a logical full-mesh topology. This means that every IBGP router must be able to communicate with every other IBGP router directly or indirectly to ensure proper route propagation.

Step-by-Step Breakdown:

- Logical Full Mesh:
 - In an IBGP setup, routers do not re-advertise routes learned from one IBGP peer to another IBGP peer. This rule is in place to prevent routing loops within the AS.
 - To ensure full route propagation, alogical full meshis required, meaning every IBGP router must peer with every other IBGP router in the AS. This can be done either directly or via route reflection or confederation.
 - Physical Full Mesh Not Required:The physical topology does not need to be a full mesh, but the BGP peering relationships must form a logical full mesh. Techniques like route reflectors or BGP confederations can reduce the need for manual full-mesh peering.
- Juniper Reference:
- IBGP Configuration: IBGP logical full mesh requirements can be simplified usingroute reflectorsto avoid the complexity of manually configuring many IBGP peers.

NEW QUESTION 5

Which two statements are true about how switches handle Layer 2 traffic? (Choose two.)

- A. The MAC address is learned based on the destination MAC address.
- B. The MAC address is learned based on the source MAC address.
- C. Traffic is forwarded based on the source MAC address.
- D. Traffic is forwarded based on the destination MAC address.

Answer: BD

Explanation:

In Layer 2 switching, switches learn MAC addresses based on thesource MAC addressof incoming frames and forward frames based on thedestination MAC address.

Step-by-Step Breakdown:

- MAC Learning:When a switch receives a frame, it records thesource MAC addressand the port on which it arrived. This allows the switch to know where to send traffic destined for that MAC address.
 - Forwarding Based on Destination:The switch then looks at thedestination MAC addressand forwards the frame out of the port associated with that MAC address. If the MAC is unknown, the switch floods the frame to all ports.
- Juniper Reference:
- Layer 2 Switching: Juniper switches use source MAC addresses to build MAC tables and forward traffic based on the destination MAC address.

NEW QUESTION 6

Which operation mode command will display the mapping between the VLAN ID and ports on a switch?

- A. show route
- B. show ethernet-switching table
- C. show interfaces terse
- D. show vlans

Answer: D

Explanation:

To display the mapping between VLAN IDs and ports on a Juniper switch, the `show vlans` command is used.

Step-by-Step Breakdown:

➤ **VLAN Information:** The `show vlans` command displays detailed information about VLAN configurations, including the VLAN ID, associated interfaces (ports), and VLAN membership.

➤ **Command Example:** `show vlans`

➤ This command will provide an output listing each VLAN, its ID, and the interfaces associated with the VLAN, enabling network engineers to quickly verify VLAN to port mappings.

Juniper Reference:

➤ **VLAN Verification:** Use the `show vlans` command to verify which VLANs are configured on the switch and the ports that are members of those VLANs.

NEW QUESTION 7

Which three technologies improve high availability and convergence in a data center network? (Choose three.)

- A. graceful restart (GR)
- B. Bidirectional Forwarding Detection (BFD)
- C. link loss adjacency
- D. Failover Group (FG)
- E. link aggregation group (LAG)

Answer: ABE

Explanation:

High availability and fast convergence are critical in data center networks to minimize downtime and maintain optimal performance. The following technologies contribute to achieving these goals:

➤ **Graceful Restart (GR):**

➤ GR allows routers to maintain forwarding state during control plane restarts, ensuring continuous packet forwarding while minimizing network disruptions.

➤ **Bidirectional Forwarding Detection (BFD):**

➤ BFD provides fast detection of path failures, allowing routing protocols to converge quickly by detecting link failures much faster than traditional timers.

➤ **Link Aggregation Group (LAG):**


➤ LAG increases both redundancy and bandwidth by combining multiple physical links into one logical link, providing load balancing and fault tolerance.

Juniper Reference:

➤ **High Availability Techniques:** These technologies are fundamental in ensuring rapid recovery and failover within Juniper-based data center environments.

NEW QUESTION 8

How many stages are shown in the exhibit?



- A. 2
- B. 5
- C. 6
- D. 3

Answer: D

Explanation:

The exhibit shows a Folded IP Clos Architecture, which is also referred to as a 3-stage Clos network design. This architecture typically consists of two layers of switches:

Spine Layer: The top row of switches.

Leaf Layer: The bottom row of switches.

Step-by-Step Breakdown:

Clos Architecture: A 3-stage Clos network has two types of devices: spine and leaf. In this design, each leaf switch connects to every spine switch, providing a high level of redundancy and load balancing.

Stage Explanation:

Stage 1: The first set of leaf switches.

Stage 2: The spine switches.

Stage 3: The second set of leaf switches.

The Folded Clos architecture shown here effectively "folds" the 3-stage design by combining the ingress and egress leaf layers into one, reducing it to two visible layers, but still maintaining the overall 3-stage architecture.

Juniper Reference:

IP Clos Architecture: The 3-stage Clos design is commonly used in modern data centers for high availability, redundancy, and scalability.

NEW QUESTION 9

Which statement is correct about per-flow load balancing?

- A. Packets associated with the same flow are sent through different egress ports.
- B. The packets are guaranteed to arrive at their destination in a different order in which they were sent.
- C. Packets associated with the same flow are sent through the same egress port.
- D. The packets are guaranteed to arrive at their destination in the same order in which they were sent.

Answer: C

Explanation:

Per-flow load balancing ensures that packets within the same flow are always forwarded over the same path, ensuring that packet order is preserved.

Step-by-Step Breakdown:

Flow Definition: A flow is typically defined by a combination of packet attributes like source/destination IP, source/destination port, and protocol type. Packets that belong to the same flow are routed over the same path to avoid reordering.

Per-Flow Behavior: In per-flow load balancing, the hashing algorithm ensures that all packets in a particular flow use the same egress port, maintaining order across the network.

Juniper Reference:

Load Balancing in Juniper: This method ensures that flows are balanced across multiple paths while preventing packet reordering within a single flow.

NEW QUESTION 10

A routing policy has been created to advertise OSPF routes in BGP. Which statement is correct in this scenario?

- A. Apply the policy as an export policy within BGP.
- B. Apply the policy as an export policy within OSPF.
- C. Apply the policy as an import policy within BGP.
- D. Apply the policy as an import policy within OSPF.

Answer: A

Explanation:

When advertising OSPF routes into BGP, the appropriate routing policy should be applied as an export policy in BGP.

Step-by-Step Breakdown:

OSPF to BGP Route Advertisement: Routes learned via OSPF (a dynamic IGP) need to be exported into BGP to be advertised to external BGP peers. In Junos OS, this is done using export policies.

Export Policies in BGP: An export policy controls which routes are advertised out of a BGP session. In this scenario, the routing policy must be applied to BGP as an export policy to export the OSPF-learned routes to external BGP peers.

Policy Configuration: Example configuration:

```
set policy-options policy-statement EXPORT_OSPF term 1 from protocol ospf
```

```
set policy-options policy-statement EXPORT_OSPF term 1 then accept
```

```
set protocols bgp group export EXPORT_OSPF
```

This policy ensures that only OSPF routes are exported into BGP.

Juniper Reference:

Routing Policy: Export policies are used in BGP to control route advertisements to peers, including those learned via OSPF.

NEW QUESTION 10

How does OSPF calculate the best path to a particular prefix?

- A. It finds the path with the numerically lowest cost.
- B. It finds the path with the shortest autonomous system path.
- C. It finds the path with the least number of hops.
- D. It finds the path with the numerically lowest route preference.

Answer: A

Explanation:

OSPF (Open Shortest Path First) calculates the best path based on the cost of the route, which is derived from the bandwidth of the interfaces along the path.

Step-by-Step Breakdown:

OSPF Path Selection:

OSPF assigns a cost to each link, typically based on the link's bandwidth (higher bandwidth equals lower cost).

The OSPF algorithm computes the shortest path to a destination by adding the costs of all links in the path. The path with the numerically lowest total cost is chosen as the best path.

Cost Calculation: The OSPF cost can be manually adjusted or automatically calculated using the default formula:

$$\text{Cost} = \frac{\text{Reference Bandwidth}}{\text{Link Bandwidth}}$$

$\text{Cost} = \text{Link Bandwidth} / \text{Reference Bandwidth}$

Juniper Reference:

OSPF Best Path Selection: OSPF selects the path with the lowest cumulative cost, ensuring efficient use of higher-bandwidth links in Junos networks.

NEW QUESTION 13

You want to enable a Junos device to support aggregated Ethernet interfaces. In this scenario, which configuration hierarchy would you use?

- A. [edit switch-options]
- B. [edit system]
- C. [edit interfaces]
- D. [edit chassis]

Answer: D

Explanation:

To configure aggregated Ethernet (AE) interfaces on a Junos device, the configuration is done under the [edit chassis] hierarchy.

Step-by-Step Breakdown:

Chassis Configuration: The chassis configuration is responsible for enabling the hardware to support Link Aggregation Groups (LAGs), allowing multiple physical interfaces to be bundled into a single logical interface for load balancing and redundancy.

Command Example:

```
set chassis aggregated-devices ethernet device-count
```

This command enables a specific number of aggregated Ethernet interfaces on the device.

Juniper Reference:

LAG Configuration in Junos: The chassis hierarchy is used to allocate and manage hardware resources for aggregated Ethernet interfaces in Juniper devices.

NEW QUESTION 15

What is the definition of a trunk interface on a switch?

- A. An interface that carries multiple VLANs.
- B. An interface that carries high bandwidth.
- C. An interface that connects directly to powerful servers.
- D. An interface that carries excess traffic.

Answer: A

Explanation:

A trunk interface on a switch is used to carry traffic for multiple VLANs between switches or between a switch and another network device, like a router. Trunk

interfaces use 802.1Q tagging to identify which VLAN the traffic belongs to.

Step-by-Step Breakdown:

Trunk Ports:

Trunk ports are typically used for inter-switch links or switch-to-router links where multiple VLANs need to be carried over the same physical connection.

VLAN traffic is tagged with a VLAN ID to ensure that it is properly identified as it crosses the trunk link.

* 802.1Q VLAN Tagging:

Trunk ports use 802.1Q to tag Ethernet frames with the VLAN ID. This ensures that frames are correctly forwarded to the appropriate VLANs on the other side of the trunk.

Juniper Reference:

Trunk Interface Configuration: In Juniper switches, trunk ports are configured to carry tagged traffic for multiple VLANs, which is essential for interconnecting multiple network segments.

NEW QUESTION 16

In the Junos OS, which feature is used to create an alternate next hop with a unique preference for a static route?

- A. Preference
- B. Resolve
- C. Next-hop
- D. Qualified-next-hop

Answer: D

Explanation:

In Junos OS, the qualified-next-hop feature is used to specify an alternate next hop for a static route, along with a unique preference value.

Step-by-Step Breakdown:

Qualified-Next-Hop: A qualified-next-hop allows you to define multiple next hops for a static route, each with its own preference. This provides flexibility by allowing the router to choose the best available next hop based on reachability and preference.

Use Case: If the primary next hop becomes unreachable, the router can automatically switch to the alternate next hop defined by the qualified-next-hop with a higher preference value.

Command Example:

```
set routing-options static route 10.10.10.0/24 qualified-next-hop 192.168.1.1 preference 5
```

```
set routing-options static route 10.10.10.0/24 qualified-next-hop 192.168.1.2 preference 10
```

Preference: The next hop with the lowest preference is chosen first. If it becomes unavailable, the router will use the higher preference next hop.

Juniper Reference:

Qualified-Next-Hop: This feature is used to configure backup or alternate next hops for static routes in Juniper devices.

NEW QUESTION 19

What are two reasons why you would deploy an IP fabric instead of a traditional Layer 2 network in a data center? (Choose two.)

- A. Layer 2 networks only support a single broadcast domain.
- B. IP fabrics are better suited to smaller networks where scale is less important.
- C. Layer 3 networks support load balancing.
- D. Layer 2 networks are susceptible to loops.

Answer: CD

Explanation:

IP fabrics are Layer 3-centric network designs often used in data centers due to their scalability, efficient routing, and loop-free architecture.

Step-by-Step Breakdown:

Layer 3 Load Balancing: IP fabrics use Equal-Cost Multipath (ECMP) to distribute traffic across multiple paths, providing effective load balancing and improving bandwidth utilization. This capability is absent in traditional Layer 2 networks, which do not support ECMP for routing decisions.

Layer 2 Loops: Layer 2 networks are prone to loops because of the lack of TTL (Time-to-Live) mechanisms. Spanning Tree Protocol (STP) is required to prevent loops, but it can introduce inefficiencies by blocking links. In contrast, IP fabrics based on Layer 3 protocols are loop-free and do not need STP.

Juniper Reference:

IP Fabric: Juniper's IP fabric solutions offer efficient Layer 3 routing with built-in load balancing and loop prevention, making them ideal for modern data center architectures.

NEW QUESTION 23

Exhibit:

Exhibit

```
[edit protocols ospf]
user@router# show
area 0.0.0.0 {
    interface xe-0/0/4.0 {
        bfd-liveness-detection {
            minimum-interval 400;
            multiplier 5;
        }
    }
}
```

Referring to the exhibit, at which interval will the interface be considered down if no hello packets are received?

- A. 2000 seconds
- B. 400 milliseconds
- C. 400 seconds
- D. 2000 milliseconds

Answer: D

Explanation:

The exhibit shows the configuration of Bidirectional Forwarding Detection (BFD) for OSPF on interface xe-0/0/4.0, with the following parameters:

minimum-interval: 400 milliseconds

multiplier: 5

Step-by-Step Breakdown:

BFD Liveness Detection: BFD is used to detect link failures at sub-second intervals, providing faster convergence times for routing protocols like OSPF.

The minimum-interval is the time between BFD control packets (in milliseconds), and the multiplier indicates how many missed BFD packets trigger a failure.

Calculating Failure Detection Time: The failure detection interval is calculated as:

Failure Interval = minimum-interval * multiplier
 $\text{Failure Interval} = \text{minimum-interval} \times \text{multiplier}$

In this case:

$400 \text{ milliseconds} \times 5 = 2000 \text{ milliseconds (2 seconds)}$

$400 \text{ milliseconds} \times 5 = 2000 \text{ milliseconds (2 seconds)}$

Conclusion: If no BFD control packets are received within 2000 milliseconds (2 seconds), the interface will be considered down, triggering OSPF to recalculate routes.

Juniper Reference:

BFD Configuration: BFD parameters such as minimum-interval and multiplier are used to fine-tune the failure detection time for faster convergence.

NEW QUESTION 28

You want to minimize topology disruptions in your network when the rpd process restarts on a device. Which service would accomplish this task?

- A. Bidirectional Forwarding Detection (BFD)
- B. link aggregation groups
- C. graceful restart (GR)
- D. Virtual Chassis

Answer: C

Explanation:

Graceful Restart (GR) is a feature that allows a router to maintain forwarding even when the routing process (e.g., the rpd process in Junos) is restarting, minimizing disruption to the network.

Step-by-Step Breakdown:

Graceful Restart Function: During a GR event, the forwarding plane continues to forward packets based on existing routes, while the control plane (rpd process) is restarting. This prevents traffic loss and maintains routing stability.

Minimizing Disruptions: GR is particularly useful in ensuring continuous packet forwarding during software upgrades or routing protocol process restarts.

Juniper Reference:

Graceful Restart in Junos: GR ensures high availability by maintaining forwarding continuity during control plane restarts, enhancing network reliability.

NEW QUESTION 32

Which route is preferred by the Junos OS software routing tables?

- A. Static
- B. Aggregate
- C. Direct
- D. BGP

Answer: C

Explanation:

In Junos OS, direct routes are the most preferred routes in the routing table, having the highest priority.

Step-by-Step Breakdown:

Direct Routes:

Direct routes represent networks that are directly connected to the router's interfaces. Since these routes are directly accessible, they are assigned the highest priority and always take precedence over other types of routes.

Preference Values:

Direct routes have a preference of 0, which is the most preferred in Junos. Static routes, OSPF routes, and BGP routes have higher preference values and will only be used if there are no direct routes to the destination.

Juniper Reference:

Direct Route Preference: In Junos, direct routes are always preferred over other routes, ensuring that the router forwards traffic through locally connected networks.

NEW QUESTION 33

When using spine and leaf fabric architectures, what is the role of each device? (Choose two.)

- A. Spine nodes are used for host connectivity.
- B. Spine nodes are used for transit to other leaf nodes.
- C. Leaf nodes are used for traffic to other leafs.
- D. Leaf nodes are used for host connectivity.

Answer: BD

Explanation:

In a spine-leaf fabric architecture, which is commonly used in data center designs, each device has a distinct role to ensure efficient and scalable network traffic flow.

Step-by-Step Breakdown:

Spine Nodes:

The spine nodes form the backbone of the fabric and are responsible for transit traffic between leaf nodes. They connect to every leaf switch and provide multiple paths for traffic between leaf nodes, ensuring redundancy and load balancing.

Leaf Nodes:

The leaf nodes are used for host connectivity. These switches connect to servers, storage, or edge routers. They also connect to the spine switches to reach other leaf switches.

Juniper Reference:

Spine-Leaf Architecture: In Juniper's IP fabric designs, spine switches handle inter-leaf communication, while leaf switches manage host and endpoint connectivity.

NEW QUESTION 34

Which Junos OS routing table stores IPv6 addresses?

- A. inet.0
- B. inet0.6
- C. inet.6
- D. inet6.0

Answer: D

Explanation:

In Junos OS, routing information is stored in different routing tables depending on the protocol and address family. For IPv6 addresses, the routing table used is inet6.0.

Step-by-Step Explanation:

Routing Tables in Junos:

inet.0: This is the primary routing table for IPv4 unicast routes.

inet6.0: This is the primary routing table for IPv6 unicast routes.

inet.3: This routing table is used for MPLS-related routing.

Other routing tables, like inet.1, inet.2, are used for multicast and other specific purposes.

inet6.0 Routing Table: When IPv6 is enabled on a Juniper router, all the IPv6 routes are stored in the inet6.0 table. This includes both direct routes (connected networks) and learned routes (from dynamic routing protocols like OSPFv3, BGP, etc.).

Verification: To view IPv6 routes, the command `show route table inet6.0` is used. This will display the contents of the IPv6 routing table, showing the network prefixes, next-hop addresses, and protocol information for each route.

Juniper Reference:

Junos Command: Use `show route table inet6.0` to check IPv6 routing entries.

IPv6 Routing: Ensure that the IPv6 protocol is enabled on interfaces and that routing protocols like OSPFv3 or BGP are properly configured for IPv6 traffic handling.

NEW QUESTION 36

When troubleshooting an OSPF neighborhood, you notice that the router stopped at the ExStart state. What is the cause of this result?

- A. The priority is set to 255.
- B. There is an interval timing mismatch.
- C. There is an area ID mismatch.
- D. There is an MTU mismatch.

Answer: D

Explanation:

When an OSPF (Open Shortest Path First) neighborship is stuck in the ExStart state, it usually points to a mismatch in Maximum Transmission Unit (MTU) settings between two routers trying to establish the adjacency. The ExStart state is where OSPF routers negotiate the master-slave relationship and exchange DBD (Database Description) packets.

Step-by-Step Breakdown:

OSPF Neighbor States: OSPF goes through several states to establish an adjacency with a neighbor:

Down: No hello packets have been received.

Init: Hello packets are received, but bidirectional communication isn't confirmed.

2-Way: Bidirectional communication is established.

ExStart: The routers are negotiating who will be the master and who will be the slave, and begin to exchange DBD packets.

Exchange: The routers start exchanging the database information.

Loading: The routers process the Link-State Advertisements (LSAs).

Full: The adjacency is fully established.

MTU Mismatch Issue:

During the ExStart state, both OSPF routers must agree on their MTU values. If there is an MTU mismatch between the two routers, OSPF neighbors will fail to move from the ExStart to the Exchange state. The router with the larger MTU setting will not accept DBD packets from the router with a smaller MTU because the packets may exceed the smaller MTU size.

In Juniper devices, this behavior can be identified by examining the MTU settings using the show interfaces command and ensuring both routers have matching MTU configurations. To resolve this issue, either match the MTU settings on both routers or configure OSPF to ignore MTU mismatches using the command set protocols ospf ignore-mtu.

NEW QUESTION 40

What is the primary purpose of an IRB Layer 3 interface?

- A. to provide load balancing
- B. to provide a default VLAN ID
- C. to provide inter-VLAN routing
- D. to provide port security

Answer: C

Explanation:

The primary purpose of an IRB (Integrated Routing and Bridging) interface is to enable inter-VLAN routing in a Layer 3 environment. An IRB interface in Junos combines the functionality of both Layer 2 bridging (switching) and Layer 3 routing, allowing devices in different VLANs to communicate with each other.

Step-by-Step Breakdown:

VLANs and Layer 2 Switching:

Devices within the same VLAN can communicate directly through Layer 2 switching. However, communication between devices in different VLANs requires Layer 3 routing.

IRB Interface for Inter-VLAN Routing:

Without an IRB interface, devices in different VLANs would not be able to communicate.

Configuration:

In Juniper devices, the IRB interface is configured by assigning Layer 3 IP addresses to it. These IP addresses serve as the default gateway for devices in different VLANs.

Example configuration:

```
set interfaces irb unit 0 family inet address 192.168.1.1/24
```

```
set vlans vlan-10 l3-interface irb.0
```

This allows VLAN 10 to use the IRB interface for routing.

Juniper Reference:

IRB Use Case: Inter-VLAN routing is essential in data centers where multiple VLANs are deployed, and Juniper's EX and QFX series switches support IRB configurations for this purpose.

NEW QUESTION 43

You are configuring an aggregate route. In this scenario, which two statements are correct? (Choose two.)

- A. Reject will silently drop the traffic.
- B. Discard will silently drop the traffic.
- C. Reject will send an ICMP Destination Unreachable message back to the sender.
- D. Discard will send an ICMP Destination Unreachable message back to the sender.

Answer: BC

Explanation:

When configuring an aggregate route, you have options for how to handle traffic that matches the route but does not match any more specific route in the routing table. Two actions can be taken: discard and reject.

Step-by-Step Breakdown:

Discard:

The discard option will silently drop packets that match the aggregate route. No notification is sent to the sender, and the packet is simply dropped.

Reject:

The reject option will drop the packet and also send an ICMP Destination Unreachable message back to the sender. This informs the sender that the packet could not be delivered because there is no specific route available.

Juniper Reference:

Aggregate Routes: The reject and discard next-hop options provide different levels of feedback when packets cannot be routed, and they can be used to control how unreachable destinations are handled.

NEW QUESTION 44

What is the main purpose of Bidirectional Forwarding Detection (BFD)?

- A. to detect network path failures
- B. to determine if the forwarding routes are correct
- C. to detect the forwarding protocol
- D. to determine packet round-trip latency

Answer: A

Explanation:

Bidirectional Forwarding Detection (BFD) is a network protocol used to detect failures in the network path between two devices quickly.

Step-by-Step Breakdown:

Path Failure Detection: BFD provides a low-overhead mechanism for detecting failures in forwarding paths across Layer 3 networks. It is much faster than traditional routing protocol timers and can detect failures within milliseconds.

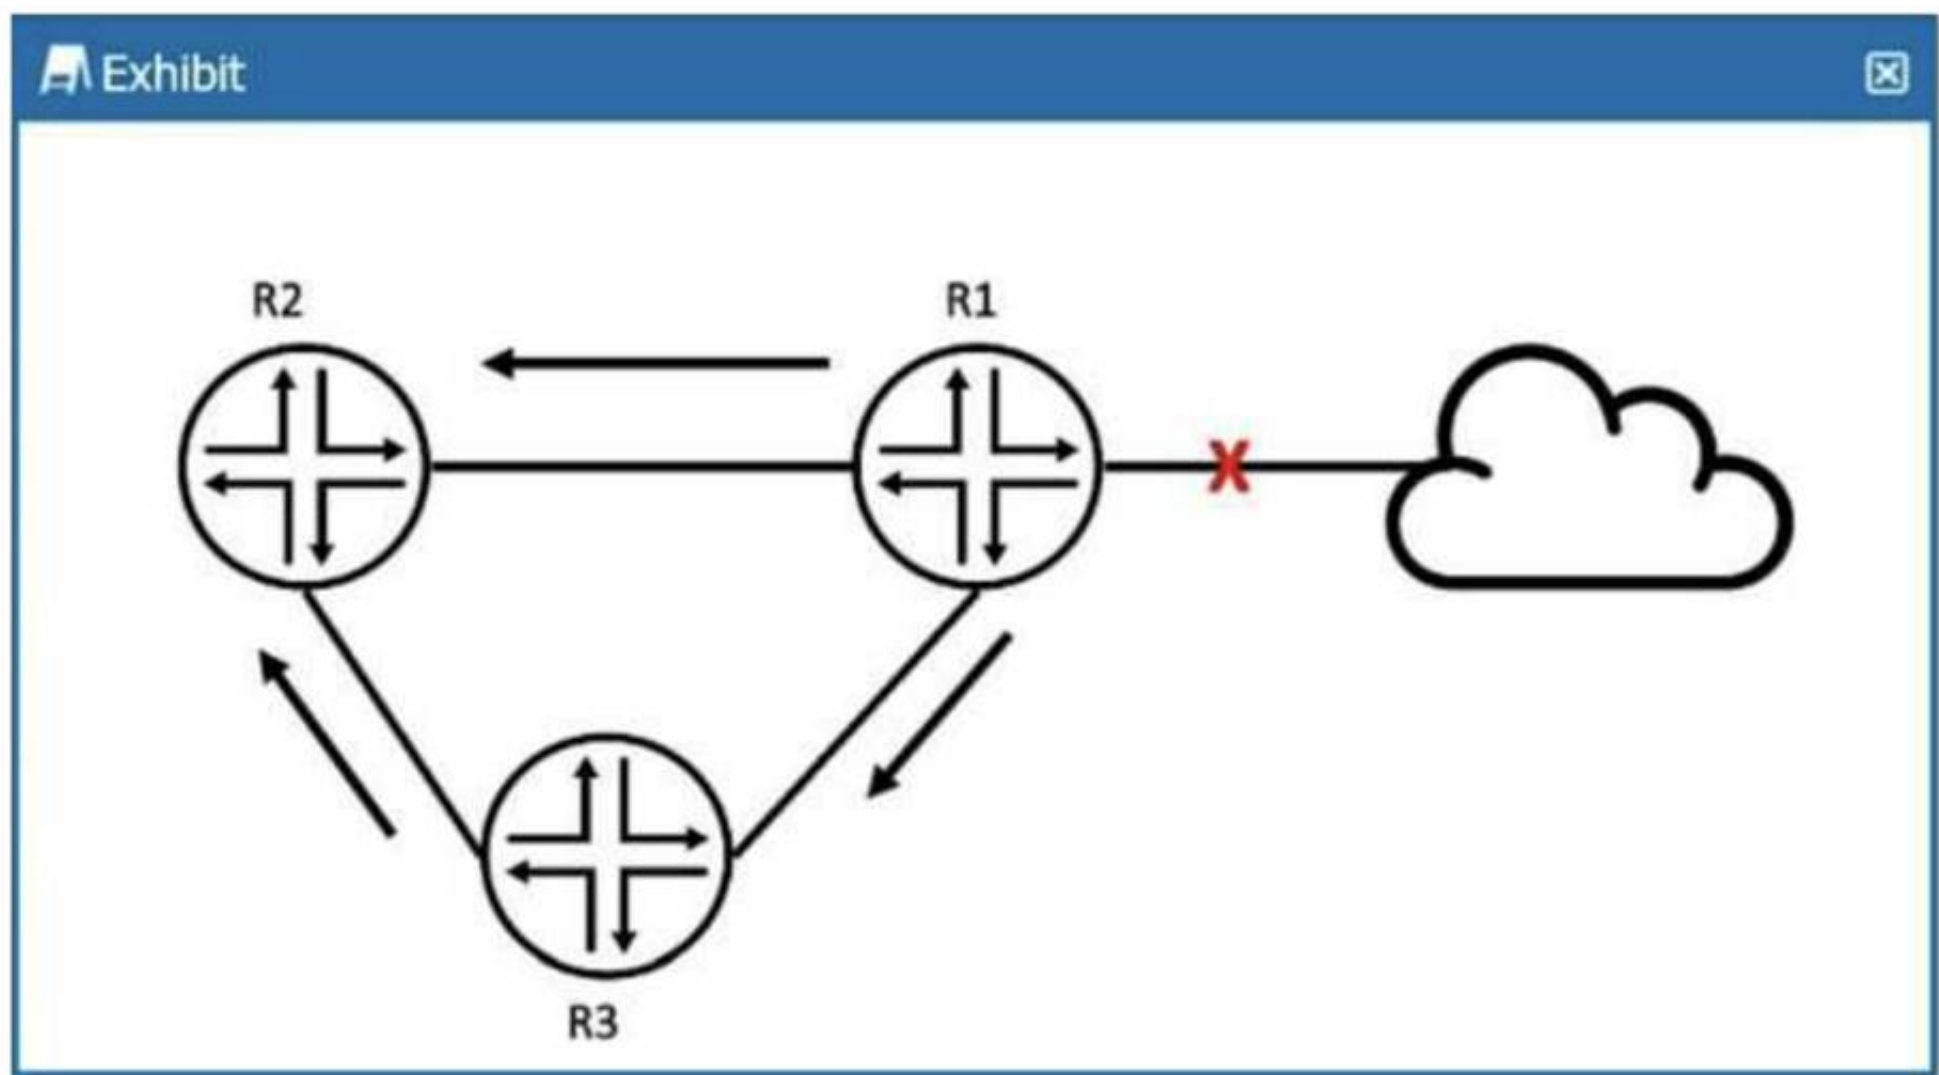
BFD in Routing: BFD can be integrated with routing protocols like OSPF, BGP, or IS-IS to trigger a faster convergence when a network path goes down.

Juniper Reference:

BFD Configuration: Juniper devices use BFD to monitor network paths and ensure fast failure detection, enhancing network resilience.

NEW QUESTION 47

Exhibit:



R2 received an OSPF update from R1, and it received the same update from R3. Referring to the exhibit, what will R2 do?

- A. R2 ignores the update from R1.
- B. R2 does nothing with R3's update.
- C. R2 ignores the update from R3.
- D. R2 acknowledges R3 and discards it.

Answer: C

Explanation:

In the exhibit, R2 receives the same OSPF update from both R1 and R3. OSPF has mechanisms to prevent unnecessary processing of duplicate LSAs (Link-State Advertisements).

Step-by-Step Breakdown:

OSPF LSA Processing:

OSPF uses LSAs to exchange link-state information between routers. When a router receives an LSA, it checks if it already has a copy of the LSA in its Link-State Database (LSDB).

Duplicate LSAs: If R2 has already received and processed the update from R1, it will ignore the update from R3 because it already has the same LSA in its database. OSPF uses the concept of flooding, but it does not reprocess LSAs that it already knows about.

R2 Behavior: R2 will keep the update from R1 (the first one it received) and will ignore the same LSA from R3, as it is already in the LSDB.

Juniper Reference:

OSPF LSA Processing: Junos adheres to OSPF standards, ensuring that duplicate LSAs are not

processed multiple times to avoid unnecessary recalculations.

NEW QUESTION 50

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

JN0-280 Practice Exam Features:

- * JN0-280 Questions and Answers Updated Frequently
- * JN0-280 Practice Questions Verified by Expert Senior Certified Staff
- * JN0-280 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * JN0-280 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The JN0-280 Practice Test Here](#)