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:
Inlegacy 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
Leaf and spine data centers are used to better accommodate which type of traffic?

A. north-east
B. east-west
C. north-west
D. south-east

Answer: B

Explanation:
In modern data centers, the shift toward leaf-spine architectures is driven by the need to handle increased east- west traffic, which is traffic between servers within the same data center. Unlike traditional hierarchical data center designs, where most traffic was "north-south" (between users and servers), modern applications often involve server-to-server communication (east-west) to enable services like distributed databases, microservices, and virtualized workloads.
Leaf-Spine Architecture:

Leaf Layer: This layer consists of switches that connect directly to servers or end-host devices. These switches serve as the access layer.

Spine Layer: The spine layer comprises high-performance switches that provide interconnectivity between leaf switches. Each leaf switch connects to every spine switch, creating a non-blocking fabric that optimizes traffic flow within the data center.
East-West Traffic Accommodation:
In traditional three-tier architectures (core, aggregation, access), traffic had to traverse multiple layers, leading to bottlenecks when servers communicated with each other. Leaf-spine architectures address this by creating multiple equal-cost paths between leaf switches and the spine. Since each leaf switch connects directly to every spine switch, the architecture facilitates quick, low-latency communication between servers, which is essential for east-west traffic flows.
Juniper's Role:Juniper Networks provides a range of solutions that optimize for east-west traffic in a leaf- spine architecture, notably through:

QFX Series Switches: Juniper??s QFX series switches are designed for the leaf and spine architecture, delivering high throughput, low latency, and scalability to accommodate the traffic demands of modern data centers.

EVPN-VXLAN: Juniper uses EVPN-VXLAN to create a scalable Layer 2 and Layer 3 overlay network across the data center. This overlay helps enhance east-west traffic performance by enabling network segmentation and workload mobility across the entire fabric.
Key Features That Support East-West Traffic:

Equal-Cost Multipath (ECMP): ECMP enables the use of multiple paths between leaf and spine switches, balancing the traffic and preventing any one path from becoming a bottleneck. This is crucial in handling the high volume of east-west traffic.

Low Latency: Spine switches are typically high-performance devices that minimize the delay between leaf switches, which improves the efficiency of server-to-server communications.

Scalability: As the demand for east-west traffic grows, adding more leaf and spine switches is straightforward, maintaining consistent performance without redesigning the entire network.
In summary, the leaf-spine architecture is primarily designed to handle the increase ineast-west trafficwithin data centers, and Juniper provides robust solutions to enable this architecture through its switch platforms and software solutions like EVPN-VXLAN.

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. TheMTU size is checked during theExchangestate.
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
Exhibit:

Exhibit

```
{master:0}[edit switch-options]
user@switch# show
interface ge-0/0/1.0 {
    persistent-learning;
}
```

Referring to the exhibit, which behavior does this configuration enable on the ge-0/0/1.0 interface?

A. This configuration enables a MAC address learned on the interface to be persistently retained in the Ethernet-switching table, even after a reboot.
B. This configuration enables the device to place a MAC address that persistently causes network errors into a special protected VLAN.
C. This configuration enables the device to shut down the interface when a particular MAC address persistently sends broadcast traffic.
D. This configuration enables the interface to learn and remember MAC addresses, until the device is rebooted.

Answer: A

Explanation:
The configuration in the exhibit shows thepersistent-learningfeature enabled on interfacege-0/0/1.0.
Step-by-Step Breakdown:

Persistent Learning:

Persistent-learningensures that the MAC addresses learned on the interface are retained in the Ethernet-switching table, even after a device reboot. This prevents the need to re-learn MAC addresses after the device restarts, improving stability and reducing downtime.

Use Case:

This feature is particularly useful in environments where the re-learning of MAC addresses could cause temporary disruptions or delays in communication, such as in critical Layer 2 network segments.

Command Example:
set switch-options interface ge-0/0/1.0 persistent-learning
Juniper Reference:

Persistent MAC Learning: In Junos, enablingpersistent-learningensures that learned MAC addresses are not lost during reboots, contributing to smoother network operations in environments where stability is crucial.

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
Within your router, you want to verify that you are learning routes from a remote BGP peer at IP address 10.10.100.1. Which command would satisfy the requirement?

A. show route receive-protocol bgp 10.10.100.1
B. show route protocol bgp table inet.0 10.10.100.1
C. show route advertise-protocol bgp 10.10.100.1
D. show route protocol bgp source-gateway 10.10.100.1

Answer: A

Explanation:
To verify that your router islearning routes from a remote BGP peerat a specific IP address (e.g., 10.10.100.1), the correct command to use is show route receive-protocol bgp.
Step-by-Step Breakdown:

BGP Route Learning:The show route receive-protocol bgp command displays the routes that have been received from a specified BGP peer. This helps in confirming that the remote peer is sending routes correctly and that your router is receiving them.

Command Example:
show route receive-protocol bgp 10.10.100.1

This will show all routes that have been received from the BGP peer with IP address 10.10.100.1.
Juniper Reference:

BGP Route Verification: Use this command to troubleshoot and verify that routes from a specific BGP peer are being received.

NEW QUESTION 7
Exhibit:

Exhibit

```
[edit routing-options]
user@Router# show
static {
    route 0.0.0.0/0 {
        next-hop 172.25.11.254;
        qualified-next-hop 172.25.11.200 {
            preference 140;
        }
    }
}
```

Referring to the exhibit, what is the route preference of the 172.25.11.254 next hop?

A. 5
B. 10
C. 130
D. 140

Answer: A

Explanation:
In the exhibit, we see two next-hop addresses for the default static route (0.0.0.0/0):
The first next hop is172.25.11.254, with no specified preference.
The second next hop is172.25.11.200, with a specified preference of140.
Step-by-Step Breakdown:
Default Static Route Preference:If no preference is explicitly set for a next hop in Junos, it defaults to5for static routes.
Determining Preference:In this case, the next hop172.25.11.254does not have an explicit preference defined, so it will use the default value of5. The second next hop has a preference of140, which is higher,meaning it will only be used if the primary next hop is unavailable.
Juniper Reference:
Static Route Preference: In Junos, the default preference for static routes is5, and this value is applied unless overridden by the preference parameter.

NEW QUESTION 8
Which two statements are correct about rules for EBGP and IBGP? (Choose two.)

A. EBGP peers have a TTL of 1, while IBGP peers have a TTL of 255.
B. EBGP peers have a TTL of 255, while IBGP peers have a TTL of 1.
C. EBGP routes are more preferred than IBGP routes.
D. IBGP routes are more preferred than EBGP routes.

Answer: AC

Explanation:
EBGP (External BGP)andIBGP (Internal BGP)operate with different rules due to the nature of their relationships.
Step-by-Step Breakdown:
TTL Differences:
EBGP: By default, EBGP peers have a TTL of 1, meaning they must be directly connected, or the TTL needs to be manually increased for multihop EBGP.
IBGP: IBGP peers within the same AS have a TTL of 255, as they are expected to communicate over multiple hops within the AS.
Preference for EBGP Routes:
Routes learned viaEBGPare typically preferred over IBGP routes. This is because EBGP routes are considered more reliable since they originate outside the AS, while IBGP routes are internal.
Juniper Reference:
BGP Configuration: The different handling of TTL and route preferences between EBGP and IBGP ensures proper route selection and security within Junos-based networks.

NEW QUESTION 9
What is the behavior of the default export policy for OSPF?

A. Accept all routes.

B. Reject all routes.
C. Redistribute all routes.
D. Forward all routes.

Answer: B

Explanation:
In Junos, thedefault export policyforOSPFis toreject all routesfrom being exported.
Step-by-Step Breakdown:
Default Export Policy:By default,OSPFin Junos does not export any routes to other routing protocols or neighbors. This is a safety mechanism to prevent unintended route advertisements.
Custom Export Policies:
If you need to export routes, you must create a customexport policythat explicitly defines which routes to advertise.
Example: You can create an export policy to redistribute static or connected routes into OSPF.
Juniper Reference:
OSPF Export Behavior: In Juniper devices, the default policy for OSPF is to reject route advertisements unless explicitly configured otherwise through custom policies.

NEW QUESTION 10
What are two consequences of having all network devices in a single collision domain? (Choose two.)

A. The amount of network resource consumption does not change.
B. The chance of packet collision is decreased.
C. The chance of packet collision is increased.
D. The amount of network resource consumption is increased.

Answer: CD

Explanation:
Acollision domainis a network segment where data packets can "collide" with one another when being sent on the same network medium.
Step-by-Step Breakdown:
Increased Collision Probability:If all devices are in asingle collision domain, the likelihood of packet collisions increases as more devices attempt to send packets simultaneously, leading to network inefficiencies.
Increased Resource Consumption:More collisions result inincreased network resource consumptionas devices need to retransmit packets, causing higher utilization of bandwidth and slowing down network performance.
Juniper Reference:
Collision Domains: Proper network segmentation using switches reduces collision domains, thereby improving network performance and reducing packet collisions.

NEW QUESTION 10
A generated route is configured under which hierarchy?

A. [edit policy-options]
B. [edit routing-instance]
C. [edit routing-options]
D. [edit protocols]

Answer: C

Explanation:
Agenerated routein Junos OS is configured under the[edit routing-options]hierarchy.
Step-by-Step Breakdown:
Generated Routes:A generated route is created based on the presence of more specific routes in the routing table. It acts as a summary route and is generated when any of its contributing routes are active. This is commonly used to create aggregate routes in OSPF, BGP, or other protocols.
Configuration Hierarchy:The configuration for generated routes is placed under[edit routing-options], where other static and routing policies are also defined.
Command Example:
set routing-options generate route 10.10.0.0/16
Juniper Reference:
Routing Options: Juniper routers use the routing-options hierarchy to configure generated routes and other static routing behaviors.

NEW QUESTION 15
When considering bidirectional forwarding detection, which two statements are correct? (Choose two.)

A. The BFD default minimum interval is 3.
B. You can configure BFD per interface within the protocol stanza.
C. The BFD operation always consists of minimum intervals and multipliers.
D. The BFD default multiplier is 5.

Answer: BC

Explanation:
Bidirectional Forwarding Detection (BFD)is a protocol used to detect faults in the forwarding path between two routers. It provides rapid failure detection, enhancing the performance of routing protocols like OSPF, BGP, and IS-IS.
Step-by-Step Breakdown:
Per Interface Configuration:BFD can be configured on aper-interfacebasis within the protocol stanza (e.g., OSPF, BGP). This allows granular control over where BFD is enabled and the failure detection intervals for specific interfaces.
Minimum Interval and Multiplier:BFD uses aminimum interval(the time between BFD control packets) and amultiplier(the number of missed packets before the path is declared down). The combination of these two defines the detection time for failures.
Juniper Reference:
BFD Configuration: In Juniper, BFD is configurable within routing protocol stanzas, with the failure detection mechanism always based on minimum intervals and

multipliers.

NEW QUESTION 17
Exhibit:

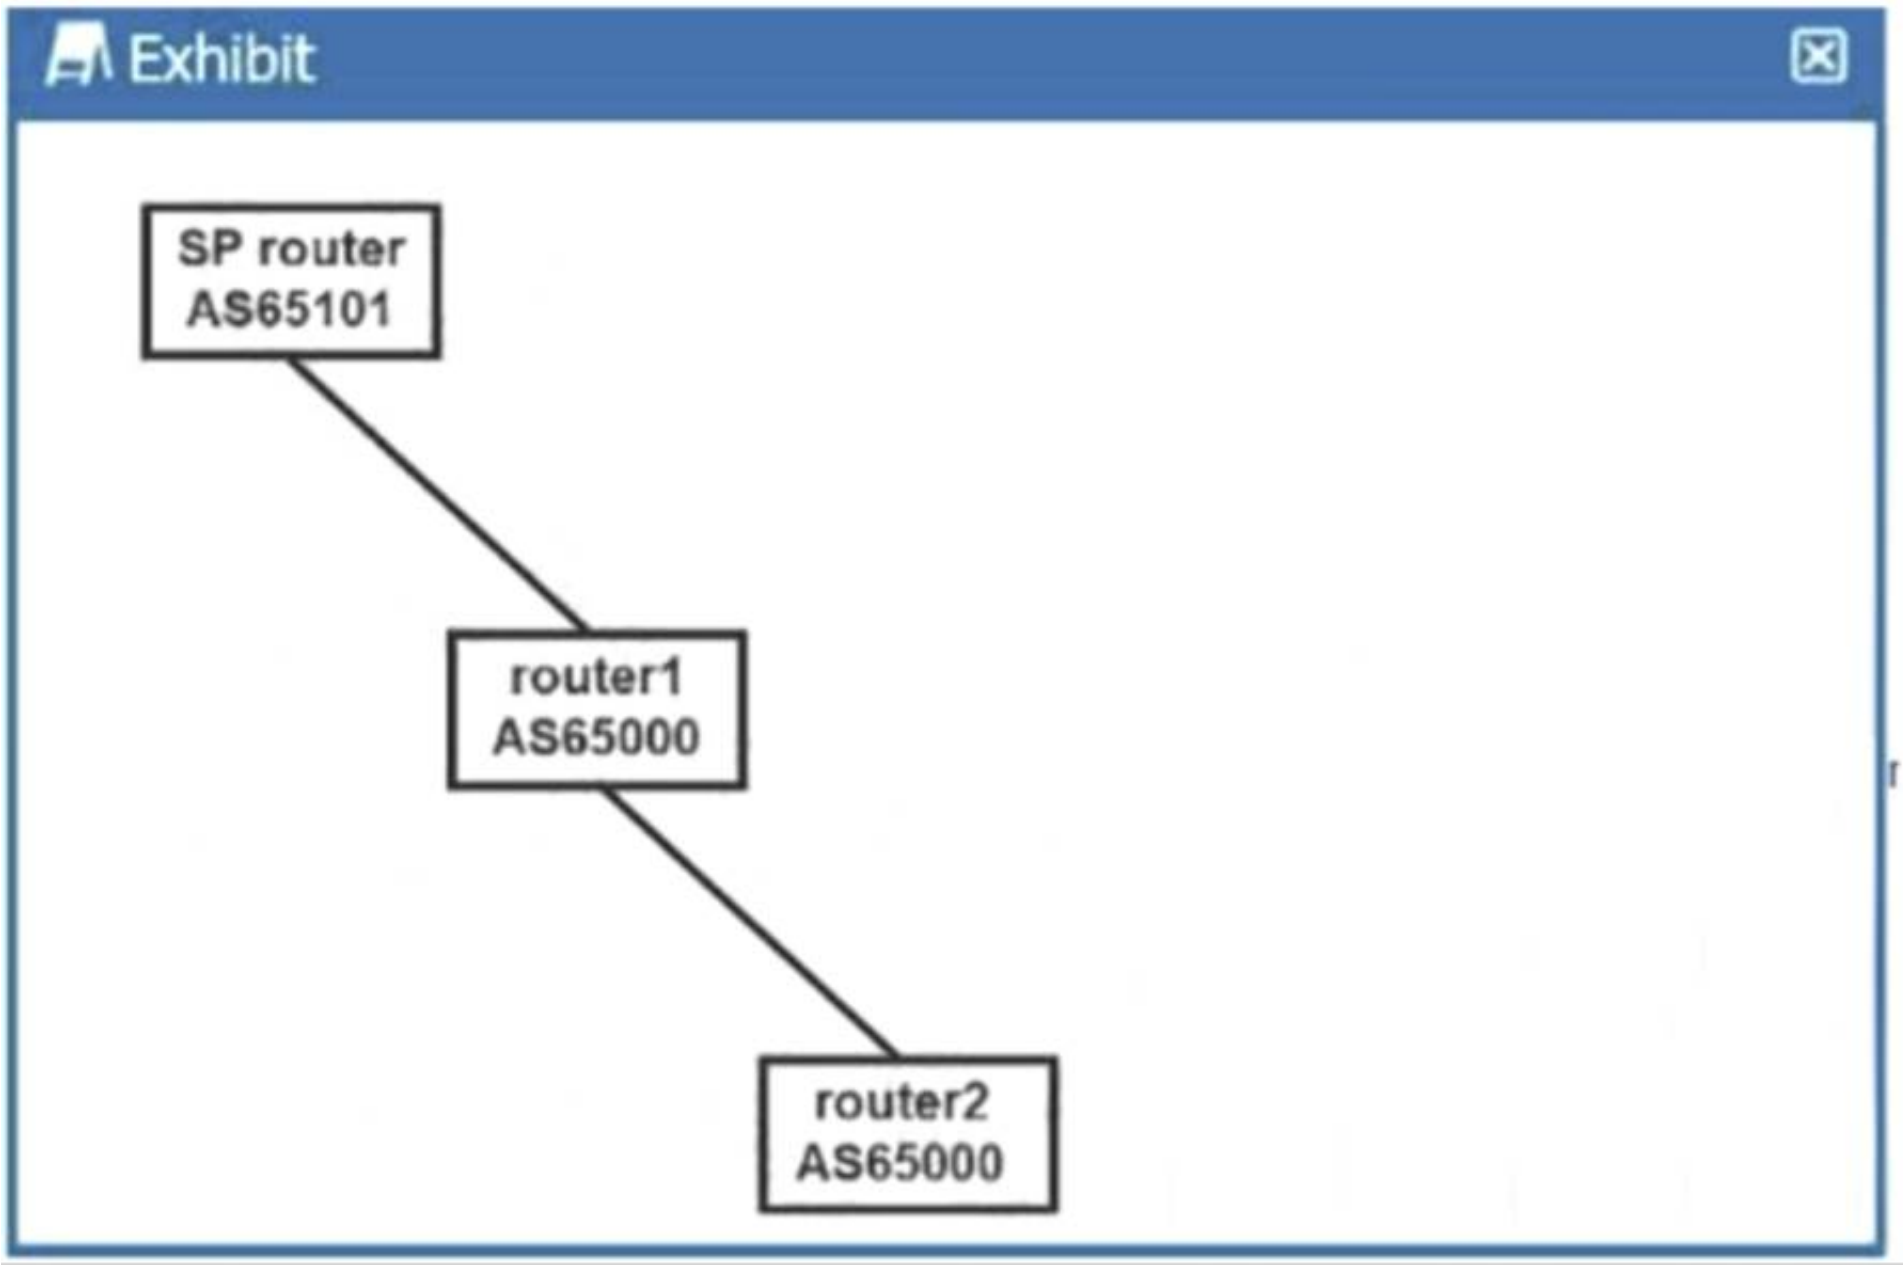

Referring to the exhibit, which two statements are correct about default BGP advertisements? (Choose two.)

A. When routes advertised by router2 are received by the SP router, they will contain the next-hop address of router2.
B. When routes advertised by router2 are received by the SP router, they will contain the next-hop address of router1.
C. When routes advertised by the SP router are received by router2, they will contain the next-hop address of the SP router.
D. When routes advertised by the SP router are received by router2, they will contain the next-hop address of router1.

Answer: BD

Explanation:
The exhibit shows a BGP peering scenario between three routers:router1androuter2are part of the same AS (AS65000), while theSP routeris in a different AS (AS65101). This indicates anEBGP (External BGP)peering between the SP router and router1, andIBGPbetween router1 androuter2.
Step-by-Step Breakdown:
Next-Hop Behavior in BGP:
IBGP: In IBGP, thenext-hop addressis not modified when advertising routes within the same AS. Thus, when router1 advertises routes learned from router2 to the SP router, it will keep thenext-hop addressof router1, not router2.
EBGP: In EBGP, thenext-hop addressis modified. When router1 receives routes from the SP router, it will advertise them to router2 with thenext-hop addressof router1.
Route Propagation:
Routes received byrouter1fromrouter2will be advertised to the SP router with router1 as the next hop.
Similarly, routes advertised by theSP routerwill be passed on to router2, with router1 remaining as the next hop.
Juniper Reference:
BGP Next-Hop: Juniper's BGP implementations follow standard BGP next-hop behavior, where the next-hop is modified in EBGP but not in IBGP, ensuring proper route advertisement across autonomous systems.

NEW QUESTION 19
A switch receives a frame with a MAC address of FF-FF-FF-FF-FF-FF. Which action will the switch take on this frame?

A. It will flood it out of all interfaces, except for the ingress interface.
B. It will flood it out of all interfaces, except for the directly connected VLAN.
C. It will flood it out of all interfaces, except for the next-hop interface.
D. It will flood it out of all interfaces.

Answer: A

Explanation:
AMAC address of FF-FF-FF-FF-FF-FFis the Ethernetbroadcast address. When a switch receives a frame with this destination MAC address, it is required to forward the frame to all interfaces except the one it was received on.
Step-by-Step Breakdown:
Broadcast Frame Handling:When a frame with the broadcast MAC address is received, the switch will flood it out of all active ports that belong to the same VLAN

as the incoming frame. The broadcast frame is not sent back out of the ingress interface (the interface where the frame was originally received).
Purpose of Flooding:Broadcasting is used to ensure that the frame reaches all devices within the broadcast domain (all devices within the same VLAN), which may not have a specific entry for the MAC address in their MAC address table.
Juniper Reference:
Layer 2 Frame Forwarding: Juniper switches flood broadcast frames to all ports in the same VLAN, except the port the frame was received on.

NEW QUESTION 21
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 routesare 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 apreference 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 routesare always preferred over other routes, ensuring that the router forwards traffic through locally connected networks.

NEW QUESTION 26
MACsec provides protection against which two types of threats? (Choose two.)

A. Data decryption
B. Playback attacks
C. Hashing attacks
D. Man-in-the-middle attack

Answer: BD

Explanation:
MACsec (Media Access Control Security)provides data confidentiality, integrity, and origin authenticity at Layer 2, protecting against several types of threats.
Step-by-Step Breakdown:
Man-in-the-Middle Attack Protection:MACsec encrypts traffic at Layer 2, preventingman-in-themiddle attackswhere an attacker intercepts and manipulates traffic between two communicating devices. Since the data is encrypted, any intercepted packets are unreadable.
Protection Against Playback Attacks:MACsec also protects againstplayback attacksby using sequence numbers and timestamps to ensure that old, replayed packets are not accepted by the receiver.
Juniper Reference:
MACsec Configuration: Juniper devices support MACsec for securing Layer 2 communications, ensuring protection against replay and man-in-the-middle attacks in sensitive environments.

NEW QUESTION 30
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 anaggregate 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:discardandreject.
Step-by-Step Breakdown:
Discard:
Thediscardoption will silently drop packets that match the aggregate route. No notification is sent to the sender, and the packet is simply dropped.
Reject:
Therejectoption will drop the packet and also send anICMP Destination Unreachablemessage 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 35
Exhibit:

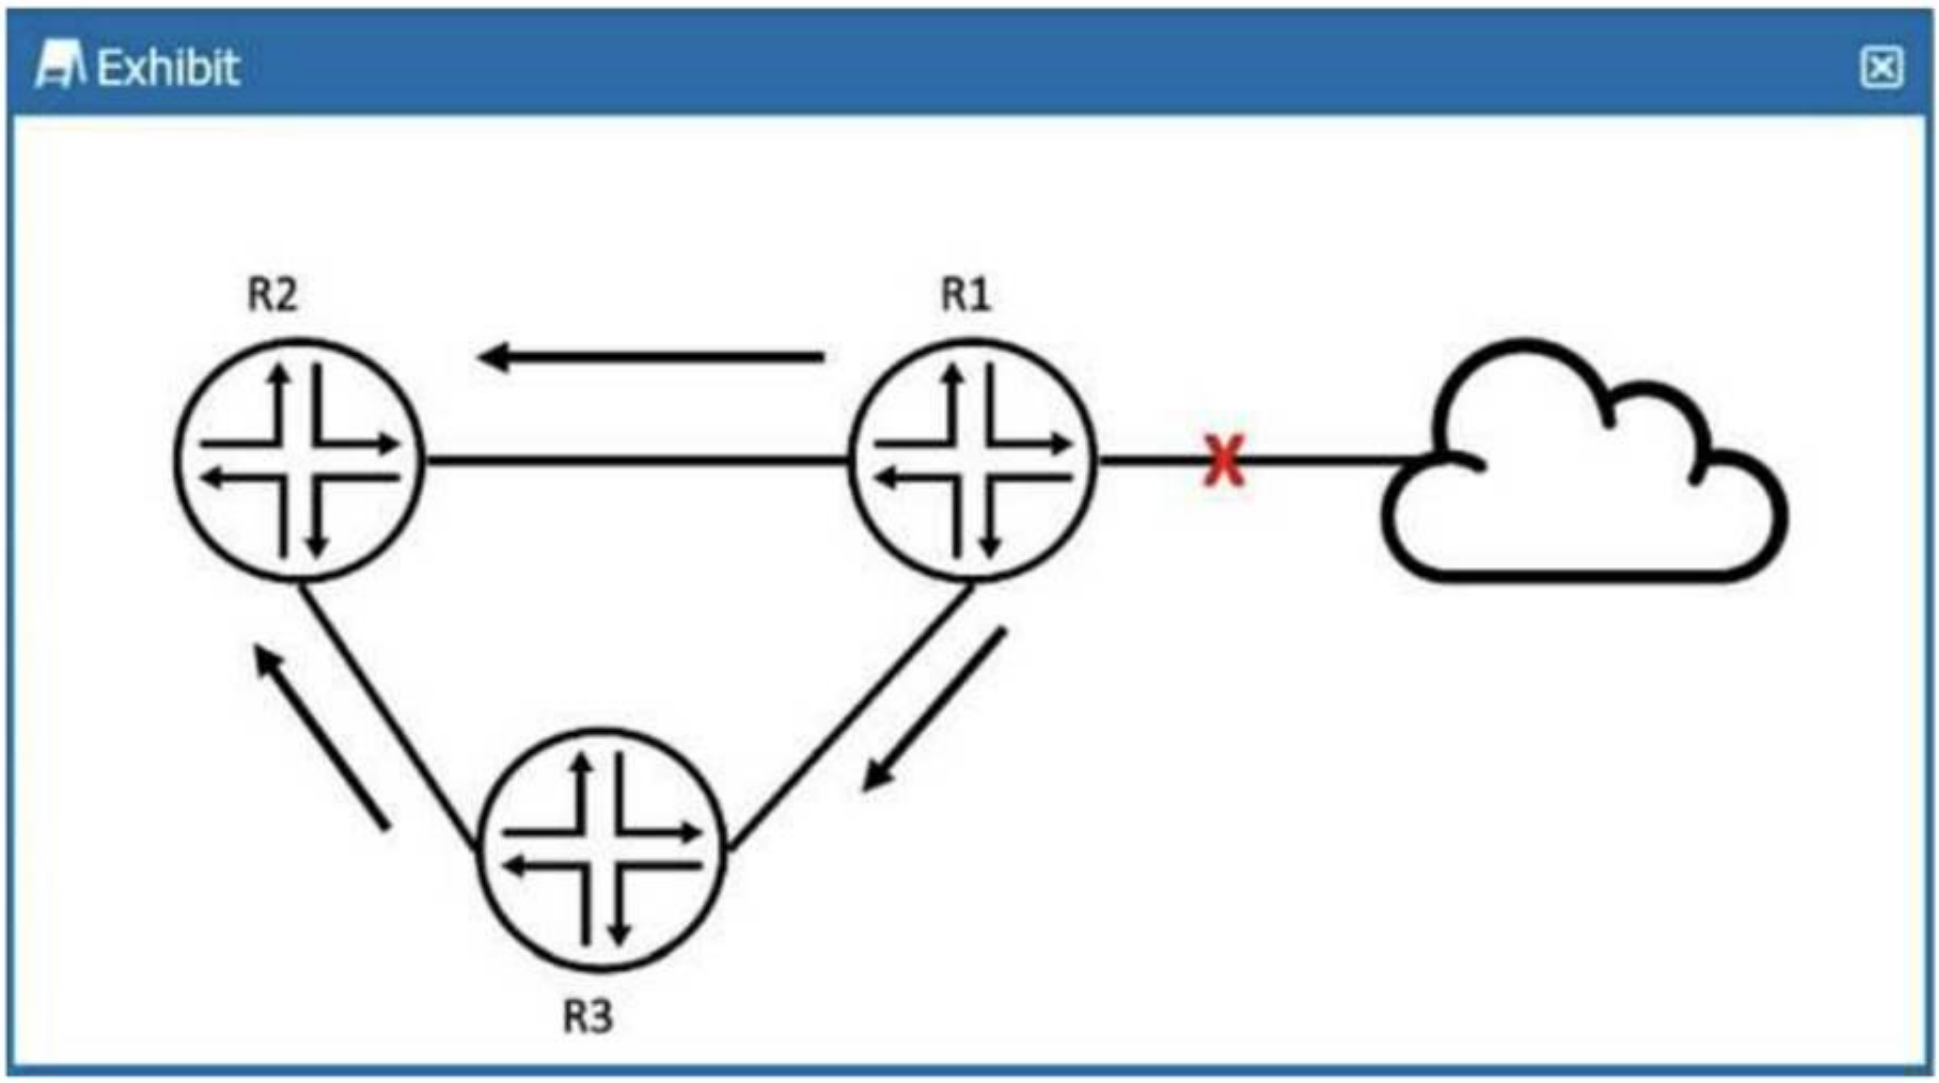

R2 received an OSPF update from R1, and it received the same update from R3.
Referring to the exhibit, what will R2 do?

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 toprevent 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 willignore the update from R3because 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 38
Which statement is correct about the BGP AS path when advertising routes?

A. The order of the AS path is not significant.
B. The local AS number is added to the end of the AS path.
C. The order of the AS path is only significant in IBGP.
D. The local AS number is added to the beginning of the AS path.

Answer: D

Explanation:
The BGPAS (Autonomous System) pathattribute is crucial in path selection and loop prevention. Each BGP router appends itslocal AS numberto the beginning of the AS path when it advertises a route to an external BGP (eBGP) peer.
Step-by-Step Breakdown:
AS Path Attribute:The AS path is a sequence of AS numbers that a route has traversed to reach a destination. Each AS adds its number to the front of the path, allowing BGP to track the route??s history.
Why the Local AS is Added at the Beginning:When advertising a route to an eBGP neighbor, a BGP router adds its own AS number to the beginning of the AS path. This ensures that the AS path reflects the route's journey accurately from the origin to the destination, and prevents loops in BGP. If the route returns to the same AS, the router will detect its AS number in the path and reject the route, preventing routing loops.
Order of the AS Path:The order is significant because BGP uses it to select the best path. A shorter AS path is preferred, as it indicates fewer hops between the source and destination.
Juniper Reference:
AS Path Attribute: Junos devices append the local AS at the start of the AS path before advertising the route to an external peer.

NEW QUESTION 40

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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 FirstTry

* JN0-280 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
Order The JN0-280 Practice Test Here