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
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 aMAC limiting violationoccurs 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 isautomatically disabledto prevent further violations. This is a protective mechanism to prevent MAC address flooding.

Packet Dropped:Additionally, packets from the violating MAC address aredroppedto 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
You are creating an IP fabric underlay and want to use OSPF as your routing protocol. In this scenario, which statement is correct?

A. All leaf devices must be configured in separate OSPF areas.
B. All leaf and spine devices must be the same model to ensure the proper load-balancing behavior.
C. Interface speeds should be the same throughout the fabric to ensure that all links are utilized.
D. All spine devices must use the same router ID.

Answer: C

Explanation:
When creating anIP fabric underlayusingOSPFas the routing protocol, consistentinterface speedsare important to ensure optimal traffic distribution and utilization of all links.
Step-by-Step Breakdown:

OSPF and Interface Speeds:OSPF calculates the cost of a link based on its bandwidth. The default cost calculation in OSPF is:

A black text on a white background Description automatically generated

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

If interface speeds vary significantly, OSPF may choose paths with lower cost (higher bandwidth), resulting in some links being underutilized.

Equal Utilization:To ensure that all links are equally utilized in an IP fabric, it is recommended to maintain uniform interface speeds across the fabric. This ensures balanced load sharing across all available paths.
Juniper Reference:

IP Fabric with OSPF: Juniper recommends consistent interface speeds to maintain even traffic distribution and optimal link utilization in IP fabric underlay designs.

NEW QUESTION 4
Which statement is correct about aggregate routes?

A. The default next hop is discard.
B. The default next hop is readvertise.
C. The default next hop is resolve.
D. The default next hop is reject.

Answer: D

Explanation:
An aggregate route is a summarized route that is created by combining multiple specific routes into a single, broader route. In Junos OS, when an aggregate route is configured, its default next hop is set toreject.
Step-by-Step Explanation::

Aggregate Route:Aggregate routes are used to reduce the size of routing tables by representing a collection of more specific routes with a single summary route. They help improve routing efficiency and scalability, especially in large networks.

Default Next Hop Behavior:

When you configure an aggregate route in Junos OS, it has arejectnext hop by default.

Therejectnext hop means that if a packet matches the aggregate route but there is no more specific route in the routing table for that destination, the packet will be discarded, and an ICMP "destination unreachable" message is sent to the source.

This behavior helps to prevent routing loops and ensures that traffic isn't forwarded to destinations for which there is no valid route.

Modifying Next Hop:If needed, the next hop behavior of an aggregate route can be changed todiscard (which silently drops the packet) or to another specific next hop. However, by default, the next hop is set toreject.
Juniper Reference:

Junos Command: set routing-options aggregate route <route> reject to configure an aggregate route with a reject next hop.

Verification: Use show route to verify the presence and behavior of aggregate routes.

NEW QUESTION 5
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 6
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 7
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 8
What is the default route preference of a static route in the Junos OS?

A. 10
B. 1
C. 5

Answer: D

Explanation:
In Junos OS, the default route preference for a static route is 5. Route preference values are used to determine which route should be installed in the routing table when multiple routes to the same destination are available.
Step-by-Step Breakdown: Static Route Preference:
A static route, by default, has a preference of 5, making it a highly preferred route. Lower preference values are more preferred in Junos, meaning static routes

take precedence over most dynamic routing protocol routes, such as OSPF (preference 10) or BGP (preference 170).
Route Preference:
Route preference is a key factor in the Junos routing decision process. Routes with lower preference values are preferred and installed in the forwarding table.
Juniper Reference:
Static Routes: In Junos, the default preference for static routes is 5, making them more preferred than most dynamic routes.

NEW QUESTION 9
Which two statements about IBGP are correct? (Choose two.)

A. By default, IBGP has a TTL of 1.
B. IBGP uses AS path for loop prevention.
C. By default, IBGP has a TTL of 255.
D. IBGP uses full mesh for loop prevention.

Answer: CD

Explanation:
:
IBGP (Internal Border Gateway Protocol)is used to exchange routing information betweenrouters within the same AS (Autonomous System).
Step-by-Step Breakdown:
TTL of 255:
By default, IBGP sessions are established with aTTL (Time to Live)value of255. This allows IBGP neighbors to communicate over multiple hops within the AS without requiring any additional configuration.
Full Mesh Requirement:
IBGP requires alogical full meshbetween all IBGP routers to ensure that routing information is fully distributed within the AS. Since IBGP does not propagate routes learned from one IBGP peer to another by default, a full mesh topology is needed unlessroute reflectorsorBGP confederationsare used.
Juniper Reference:
IBGP Full Mesh: Juniper recommends using route reflectors in large networks to simplify IBGP full-mesh requirements.

NEW QUESTION 10
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 10
Which three actions are required to implement filter-based forwarding? (Choose three.)

A. You must create an instance-type forwarding routing instance.
B. You must create an instance-type vrf routing instance.
C. You must create a match filter.
D. You must create a security policy.
E. You must create a RIB group.

Answer: ACE

Explanation:
Filter-Based Forwarding (FBF) in Junos OS allows traffic to be routed based on specific criteria such as source address, rather than just the destination address. This is useful in scenarios like policy routing or providing multiple paths for different types of traffic.
Step-by-Step Breakdown:

- Instance-Type Forwarding:You must create aninstance-type forwardingrouting instance. This routing instance allows for different routing tables based on the incoming packet filter.
- Command:

set routing-instances FBF-instance instance-type forwarding

- Match Filter:You need to create afilterto match the traffic that will be forwarded according to your custom routing policy. This filter is applied to an interface to determine which traffic will use the custom forwarding instance.
- Command Example:

set firewall family inet filter FBF-filter term 1 from source-address <address>

set firewall family inet filter FBF-filter term 1 then routing-instance FBF-instance

RIB Group:ARIB (Routing Information Base) groupis necessary to share routes between the primary routing table and the custom routing instance. This allows FBF traffic to use the routing information from other routing tables.

Command Example:
set routing-options rib-groups FBF-group import-rib inet.0
set routing-instances FBF-instance routing-options rib-group FBF-group
Juniper Reference:

FBF Configuration: Filter-based forwarding requires these specific steps to redirect traffic to a custom routing table based on filter criteria.

NEW QUESTION 11
Which statement is correct about a three-stage IP fabric underlay?

A. Every ingress interface into the fabric is only two hops away from the egress interface.
B. Every spine device can communicate directly with other spine devices.
C. Every leaf device can communicate directly with other leaf devices.
D. Every server that connects to a three-stage IP fabric must be multihomed.

Answer: A

Explanation:
In athree-stage IP fabric(also known as a Clos fabric), traffic between any two points (ingress to egress) in the fabric isonly two hops away.
Step-by-Step Breakdown:

Three-Stage IP Fabric:

Leaf Layer: Leaf switches connect directly to servers and edge devices.

Spine Layer: Spine switches provide connectivity between leaf switches but do not connect to each other directly.

Two-Hop Communication:In this architecture, every leaf switch is connected to every spine switch.
Therefore, when a packet enters the fabric via an ingress leaf switch, it is forwarded to a spine switch, which then directs the packet to the correct egress leaf switch. This path always involves exactly two hops:

Ingress leaf # Spine # Egress leaf.

Benefits:This consistent two-hop path ensures predictable latency and makes the network highly scalable while maintaining low complexity.
Juniper Reference:

IP Fabric Architecture: This two-hop property of Clos fabrics is a hallmark of spine-leaf designs, as supported by Juniper's QFX and EX switches in data centers.

NEW QUESTION 16
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 17
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 22
Which two statements are correct about EVPN-VXLAN overlay networking? (Choose two.)

A. It is the only option to provide reachability between servers that reside in the same network segment in a data center.
B. BGP provides the control plane within the overlay network.
C. An encapsulation of the original packet is required to transport the packet across the network.
D. OSPF provides the control plane within the overlay network.

Answer: BC

Explanation:
EVPN-VXLANis an overlay technology used in data center networks to extend Layer 2 services over a Layer 3 network.
Step-by-Step Breakdown:
BGP Control Plane:BGP (Border Gateway Protocol)is used as the control plane for EVPN-VXLAN. BGP advertises MAC addresses and IP address reachability information across the VXLAN network, enabling efficient multi-tenant Layer 2 connectivity over a Layer 3 infrastructure.
Encapsulation:VXLAN (Virtual Extensible LAN)encapsulates Layer 2 frames into Layer 3 packets. This encapsulation allows Layer 2 traffic to be transported across a Layer 3 network, effectively creating a tunnel for Ethernet frames.
Juniper Reference:
EVPN-VXLAN Configuration: Juniper supports EVPN-VXLAN with BGP as the control plane, allowing scalable Layer 2 connectivity over a routed infrastructure in modern data centers.

NEW QUESTION 26
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 advertisingOSPF routesintoBGP, the appropriaterouting policyshould be applied as anexport policyin 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 usingexport 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 toBGP as an export policyto 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 29
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 configureaggregated Ethernet (AE) interfaceson a Junos device, the configuration is done under the[edit chassis]hierarchy.
Step-by-Step Breakdown:
Chassis Configuration:Thechassisconfiguration is responsible for enabling the hardware to supportLink 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: Thechassishierarchy is used to allocate and manage hardware resources for aggregated Ethernet interfaces in Juniper devices.

NEW QUESTION 30
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 32
Which statement is correct about areas in OSPF?

A. An OSPF area is used to segment Layer 2 broadcast domains.
B. OSPF areas are used to isolate the effects of a broadcast storm.
C. OSPF areas are used to reduce the size of the link-state database.
D. An OSPF area is used to signify the autonomous system to which each device belongs.

Answer: C

Explanation:
InOSPF (Open Shortest Path First), areas are used to segment a network into smaller, more manageable pieces to improve scalability. By dividing a network into areas, OSPF can reduce the size of thelink-state database (LSDB), which helps routers process updates more efficiently.
Step-by-Step Breakdown:
Purpose of OSPF Areas:OSPF areas allow for hierarchical routing within the OSPF domain. Routers in the same area have identical LSDBs, but routers in different areas do not exchange full link-state information. Instead, they exchange summarized routes, which reduces the LSDB size and CPU/memory usage.
Benefits:Reducing the LSDB size improves scalability and ensures faster convergence in larger networks. Area 0 is the backbone area, and all other areas must connect to it, forming a hierarchical structure.
Juniper Reference:
OSPF Configuration: Areas in OSPF are configured to optimize network performance by limiting the scope of link-state advertisements (LSAs) to within an area.

NEW QUESTION 36
Which two statements describe an IP fabric? (Choose two.)

A. An IP fabric allows devices to always be one hop away.
B. An IP fabric depends on Layer 2 switching.
C. An IP fabric uses spine and leaf devices.
D. An IP fabric provides traffic load sharing.

Answer: CD

Explanation:
AnIP fabricis a network topology designed to provide a scalable, low-latency architecture that is typically implemented in modern data centers. It usesspine and leafswitches and enables efficienttraffic load sharingacross the network.
Step-by-Step Breakdown:
Spine-Leaf Architecture:
Leaf Devices: These switches connect to servers and edge devices within the data center. Each leaf switch connects to every spine switch.
Spine Devices: These high-performance switches interconnect all the leaf switches. There are no direct connections between leaf switches or spine switches. This architecture ensures that any two endpoints within the fabric are only one hop away from each other, minimizing latency.
Traffic Load Sharing:
An IP fabric leveragesEqual-Cost Multipath (ECMP)to distribute traffic evenly across all available paths between leaf and spine switches, providing effective load balancing. This ensures that no single link becomes a bottleneck and that traffic is distributed efficiently across the network.
Juniper Reference:
Juniper providesQFX Seriesswitches optimized for IP fabric topologies, allowing for scalable deployments in modern data centers.
EVPN-VXLAN: Often used in IP fabrics to extend Layer 2 services across the fabric with Layer 3 underlay, enabling both efficient routing and bridging.

NEW QUESTION 37
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, thequalified-next-hopfeature 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:Aqualified-next-hopallows 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 thequalified-next-hopwith 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 41
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 fabricsare 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 useEqual-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 networksare 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: Juniperâ€™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 43
Which two statements are correct about aggregate routes and generated routes? (Choose two.)

A. An aggregate route does not have a forwarding next hop.
B. An aggregate route has a forwarding next hop.
C. A generated route has a forwarding next hop.
D. A generated route does not have a forwarding next hop.

Answer: AC

Explanation:
Aggregate routesandgenerated routesare used to create summarized routes in Junos, but they behave differently in terms of forwarding.
Step-by-Step Breakdown:
Aggregate Routes:
Anaggregate routesummarizes a set of more specific routes, but it does not have a direct forwarding next hop. Instead, it points to the more specific routes for actual packet forwarding.
Generated Routes:
Agenerated routealso summarizes specific routes, but it has aforwarding next hopthat is determined based on the availability of contributing routes. The generated

route can be used to directly forward traffic.
Juniper Reference:
Aggregate and Generated Routes: In Junos, aggregate routes rely on more specific routes for forwarding, while generated routes can forward traffic directly based on their next-hop information.

NEW QUESTION 48
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 52
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 ofBidirectional 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.
Theminimum-intervalis the time between BFD control packets (in milliseconds), and themultiplierindicates 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}Failure Interval=minimum-interval??multiplier
In this case:

400?milliseconds??5=2000?milliseconds(2seconds)400 \, \text{milliseconds} \times 5 = 2000 \, \text{milliseconds} (2 seconds)400milliseconds??5=2000milliseconds(2seconds)
Conclusion:If no BFD control packets are received within2000 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 56
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 aGR 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 60
You are troubleshooting a downed BGP session.

Exhibit

```
user@router> show bgp summary | match "AS|Connect|Active"
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.10.55.2             1111          0          0       0      10 1w0d 3:51:17 Connect
```

Referring to the exhibit, what is the cause of the problem?

A. The UDP session between the peers has not been established.
B. The local peer has sent an Open message but not received one from the remote peer.
C. The TCP session between the peers has not been established.
D. The local peer has sent an Update message but not received one from the remote peer.

Answer: C

Explanation:
The BGP session in the exhibit shows the state as Connect, which indicates that the TCP session between the BGP peers has not been fully established.
Step-by-Step Breakdown:
BGP State 'Connect':
The Connect state is the second stage in the BGP finite state machine (FSM). At this stage, BGP is trying to establish a TCP session with the peer, but the session has not yet been successfully established.
A successful TCP three-way handshake (SYN, SYN-ACK, ACK) is required before BGP can progress to the OpenSent state, where the peers exchange BGP Open messages.
Possible Causes:
A firewall blocking TCP port 179.
Incorrect IP addresses or network connectivity issues between the BGP peers.
Juniper Reference:
BGP Troubleshooting: In Junos, if a BGP session is stuck in the Connect state, the issue is likely due to a failure in establishing the underlying TCP connection.

NEW QUESTION 64
When evaluating BGP routes, what will be evaluated first?

A. The local preference value
B. The AS path
C. The MED value
D. The origin value

Answer: A

Explanation:
InBGP (Border Gateway Protocol), when evaluating multiple routes to the same destination, the first attribute that is considered is thelocal preferencevalue.
Thelocal preferenceis a BGP attribute used to influence outbound routing decisions within an Autonomous System (AS).

Step-by-Step Breakdown:
Local Preference:Thelocal preferenceattribute is used to determine which path is preferred for traffic leaving the AS. The higher the local preference value, the more preferred the route.
BGP Path Selection:The BGP path selection process evaluates the following attributes in this order:
Local Preference(higher is preferred)
AS Path(shorter is preferred)
Origin(IGP > EGP > incomplete)
MED (Multi-Exit Discriminator)(lower is preferred)
Juniper Reference:
BGP Path Selection: In Junos, thelocal preferenceattribute is the first to be evaluated when determining the best path for outbound traffic.

NEW QUESTION 68
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 71
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 75
Exhibit:

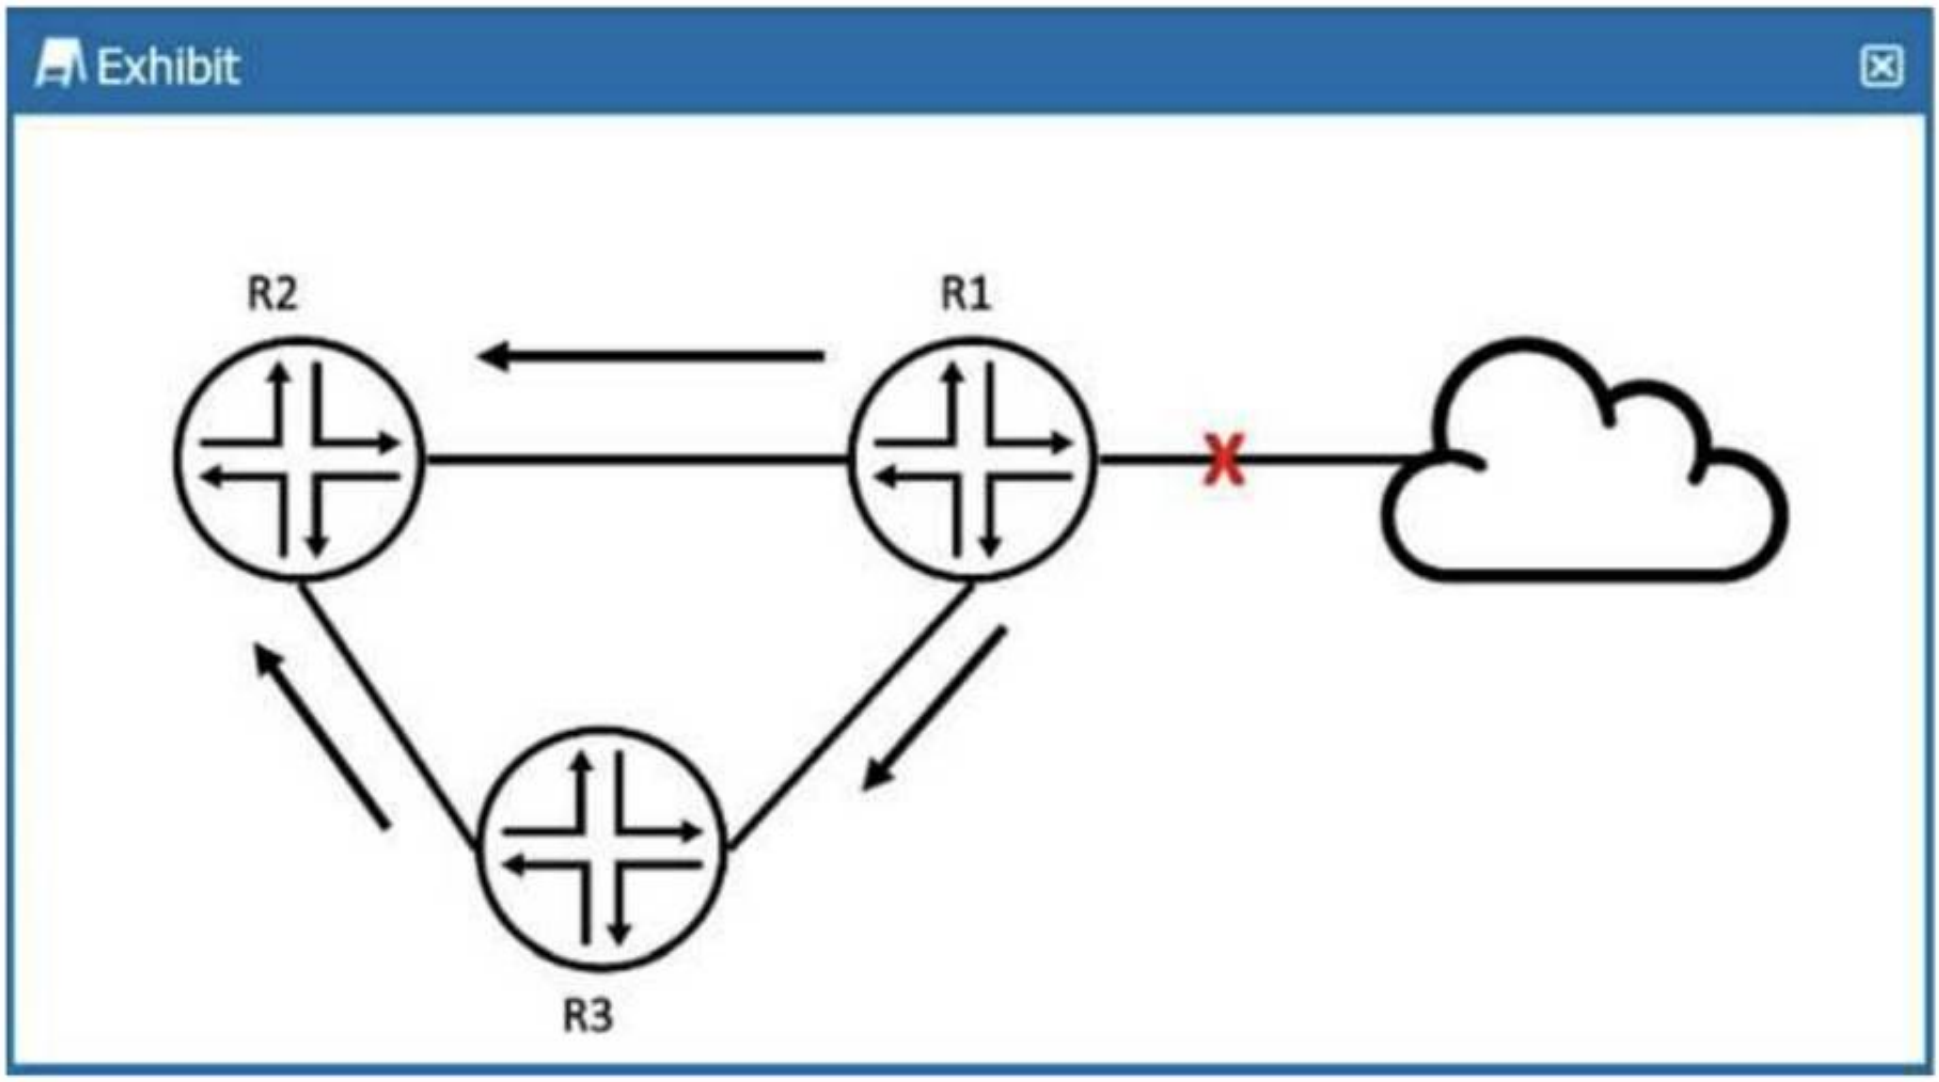

R2 received an OSPF update from R1, and it received the same update from R3.
Referring to the exhibit, what will R2 do?

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 80
Exhibit:

Exhibit

```
[edit routing-options]
user@router# show
static {
    defaults {
        preference 7;
    }
    route 0.0.0.0/0 {
        next-hop 172.25.20.254;
        qualified-next-hop 172.25.20.200 {
            preference 6;
        }
    }
}
```

Referring to the exhibit, which next hop will be preferred in the routing table?

A. Next hop IP address 172.25.20.254 will be preferred.
B. Neither next hop will be preferred.
C. Next hop IP address 172.25.20.200 will be preferred.
D. Both next hops will be preferred.

Answer: C

Explanation:
In the exhibit, we see a static route configuration with two possible next hops for the default route (0.0.0.0/0):
next-hop 172.25.20.254with the default preference of 7.
qualified-next-hop 172.25.20.200with a preference of 6.
Step-by-Step Breakdown:
Preference Value:In Junos OS, thepreferencevalue is used to determine which route should be preferred in the routing table. The lower the preference value, the higher the priority for the route.
Comparison:In this case:
The next hop172.25.20.254has a preference of7.
Thequalified-next-hop 172.25.20.200has a preference of6.
Preferred Next Hop:Since172.25.20.200has a lower preference (6) compared to172.25.20.254(7), it will be the preferred next hop in the routing table, assuming both next hops are reachable.
Juniper Reference:
Qualified Next Hop: In Junos, static routes with multiple next-hop options are selected based on the preference value, with the lower value being preferred.

NEW QUESTION 84
.....

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