JN0-280 Dumps

Data Center Associate (JNCIA-DC)

https://www.certleader.com/JN0-280-dumps.html

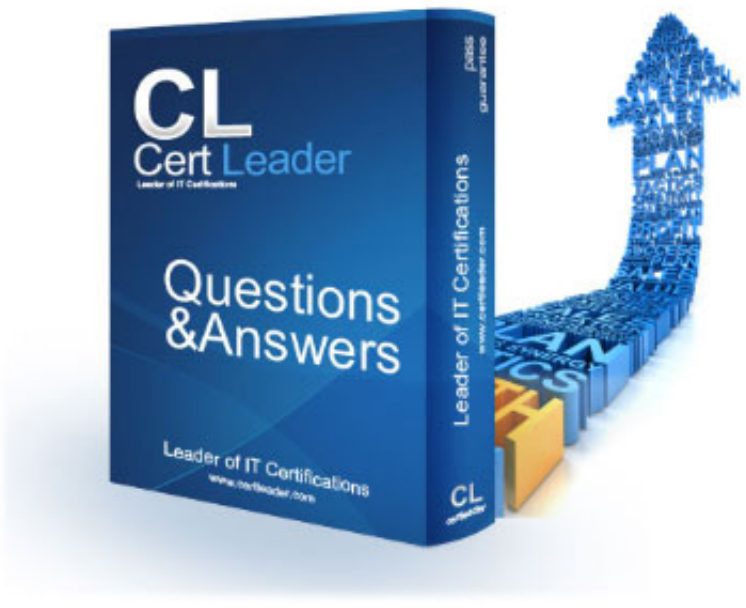

NEW QUESTION 1
What are three correct layer names used in legacy hierarchical network design? (Choose three.)

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
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 3
What are two requirements for an IP fabric? (Choose two.)

A. a Layer 3 routing protocol
B. a single connection between each spine and leaf
C. a single connection between each leaf
D. a Layer 2 switching protocol

Answer: AB

Explanation:
AnIP fabricis a network architecture commonly used in data centers to provide scalable, high-throughput connectivity using aspine-leaf topology.
Step-by-Step Breakdown:

- Layer 3 Routing Protocol:An IP fabric relies on aLayer 3 routing protocol, typically BGP or OSPF, to provide routing between the leaf and spine switches. This ensures efficient traffic forwarding across the network.
- Single Connection Between Spine and Leaf:In an IP fabric, each leaf switch connects toevery spine switchwith a single connection. This ensures that traffic between any two leaf switches can travel through the spine layer in just two hops.

Juniper Reference:

- Spine-Leaf Design: Juniper??s IP fabric implementations are designed for scalability and low-latency routing, often using protocols like BGP for Layer 3 control.

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 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 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 themapping between VLAN IDs and portson 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 interfacesassociated 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
Exhibit:

Exhibit

```
[edit]
user@router# show protocols
bgp {
    group 1 {
        local-as 65101;
        }
        neighbor 172.16.1.1 {
            peer-as 65201;
        }
    }
}
[edit]
user@router# show routing-options
router-id 192.168.100.1;
autonomous-system 65000;
```

Referring to the exhibit, which statement is correct?

A. The configuration will commit successfully and BGP group 1 will operate as IBGP.
B. The configuration will commit successfully and BGP group 1 will operate as EBGP.
C. BGP group 1 requires a type external parameter.
D. BGP group 1 requires a type internal parameter.

Answer: B

Explanation:
In the exhibit, BGP is configured withlocal AS 65101and a neighbor at172.16.1.1inpeer AS 65201. This setup involves two different Autonomous Systems (AS), indicating anExternal BGP (EBGP)configuration.
Step-by-Step Breakdown:

- EBGP vs. IBGP:
- EBGPis used between routers in different ASes. In this case, the local AS is65101and the peer AS is65201, meaning the BGP session isEBGP.
- IBGPis used between routers within the same AS, which is not applicable here as the AS numbers are different.
- BGP Group Configuration:
- The configuration does not require a type external parameter because Junos OSautomatically recognizes the session asEBGPwhen the local and peer AS numbers are different.
- The BGP session will operate as EBGP, and the configuration will commit successfully.

Juniper Reference:

- BGP Configuration: In Juniper, EBGP is automatically recognized when the local and peer AS numbers differ, without needing to specify type external.

NEW QUESTION 8
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 9
Which signaling protocol is used for EVPN?

A. OSPF
B. PIM
C. IS-IS
D. BGP

Answer: D

Explanation:
EVPN (Ethernet Virtual Private Network)is a standard protocol used for building Layer 2 and Layer 3 VPNs over an IP or MPLS network. Thesignaling protocolused for EVPN isBGP(Border Gateway Protocol).
Step-by-Step Breakdown:
BGP as the EVPN Signaling Protocol:EVPN uses BGP to exchange MAC address reachability information between routers (PE devices). This enables devices to learn which MAC addresses are reachable through which PE devices, facilitating Layer 2 forwarding across an IP or MPLS core.
BGP Extensions for EVPN:BGP is extended with new address families (e.g., EVPN NLRI) to carry both MAC and IP address information, allowing for scalable and efficient multi-tenant network solutions.
Juniper Reference:
Junos EVPN Configuration: Juniper uses BGP as the control plane for EVPN to exchange MAC and IP route information between different data center devices.

NEW QUESTION 10
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 10
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 balancingensures 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:Inper-flow load balancing, the hashing algorithm ensures that all packets in a particular flow use thesame 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 13
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 thecostof 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 thenumerically lowest total costis chosen as the best path.
Cost Calculation:The OSPF cost can be manually adjusted or automatically calculated using the default formula:
Cost=Reference BandwidthLink Bandwidth\text{Cost} = \frac{\text{Reference Bandwidth}}{\text{Link Bandwidth}}Cost=Link BandwidthReference 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 17
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 21
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 23
What information in the Ethernet header is used to populate the bridging table?

A. destination address
B. source address
C. type
D. protocol

Answer: A

Explanation:
Thesource MAC addressin the Ethernet header is used to populate thebridging table(also called the MAC address table) on a switch. When a frame arrives at a switch, the switch examines the source MAC address and records it along with the ingress port in its MAC address table.
Step-by-Step Breakdown:
Learning Process:When an Ethernet frame arrives on a switch port, the switch looks at thesource MAC addressand adds this MAC address to the MAC table along with the port it was received on. This process is calledMAC learning.
Purpose:The switch uses this information to determine the correct port to send frames destined for that MAC address in future transmissions, thus ensuring efficient Layer 2 forwarding.
Juniper Reference:
Ethernet Switching: Juniper switches use source MAC addresses to build and maintain the MAC address table, which is essential for Layer 2 switching.

NEW QUESTION 27
Which static routing parameter will silently drop the packet if it is set as the next hop?

A. Reject
B. Resolve
C. Readvertise
D. Discard

Answer: D

Explanation:
When thediscardoption is configured as the next hop for a static route, it silently drops any packets that match the route without sending any notification to the sender.
Step-by-Step Breakdown:
Discard Behavior:
If a route uses thediscardnext hop, the router drops the packet without generating any ICMP message or error back to the sender. This is useful for creating null routes to prevent routing loops or blackhole traffic intentionally.
Reject vs. Discard:
Therejectnext hop, in contrast, drops the packet but sends anICMP Destination Unreachablemessage back to the source.
Juniper Reference:
Static Route Behavior: In Junos, thediscardoption ensures packets matching a static route are dropped silently, providing a way to discard traffic without alerting the source.

NEW QUESTION 29
Which two statements are correct about VLAN tags? (Choose two.)

A. VLAN tags carry a VLAN ID and priority.
B. VLAN tags are required on access ports.
C. VLAN tags require multiple forwarding tables.
D. VLAN tags can be inserted or removed by trunk interfaces.

Answer: AD

Explanation:
VLAN tagsare used in Ethernet frames to identify and differentiate traffic between multiple VLANs. They are especially important for devices like switches that handle multiple VLANs on the same physical link.
Step-by-Step Breakdown:
VLAN Tag Contents:
VLAN ID: The tag contains a 12-bit VLAN ID field that identifies the VLAN to which the frame belongs.
Priority: The tag also includes a 3-bit priority field (also known as 802.1p priority) used for QoS (Quality of Service) to prioritize traffic.
Trunk Ports and VLAN Tagging:
Trunk Portsare used to carry traffic for multiple VLANs across a single link. These interfaces insert (tag) VLAN identifiers into frames when they leave the switch and remove (untag) them when frames enter the switch.
Access Ports:VLAN tags are typically not used onaccess ports(ports that connect to end devices) since those ports are configured to be part of a single VLAN, and the traffic doesn't need VLAN tags.
Juniper Reference:
VLAN Tagging: Juniper switches support VLAN tagging and ensure that frames are tagged or untagged as they traverse trunk or access ports, respectively.

NEW QUESTION 34
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 37
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 41

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 45
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 50
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 isinet6.0.
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 theinet6.0table. 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 routingprotocols like OSPFv3 or BGP are properly configured for IPv6 traffic handling.

NEW QUESTION 52
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 56
Referring to the exhibit, you notice that after committing the configuration, the ae0 and ae1 interfaces appear in a link down state.

Exhibit

```
[edit]
user@switch# show chassis
aggregated-devices {
    ethernet {
        device-count 2;
    }
}
[edit]
user@switch# run show interfaces terse | match ae
ae0                     up    down
ae1                     up    down
```

Which statement is correct in this scenario?

A. No operational interfaces have been added to the LAG interfaces.
B. No traffic is traversing the LAG interfaces.
C. The LAG interfaces are in a passive state.
D. The LAG interfaces are in aggressive mode.

Answer: A

Explanation:
In the exhibit, the ae0 and ae1 interfaces are in a link down state. This occurs when no physical interfaces (member interfaces) have been added to the LAG (Link Aggregation Group) interfaces, or the member interfaces are not operational.
Step-by-Step Breakdown:
LAG Configuration:
A LAG interface (aggregated Ethernet interface) is a logical interface that combines multiple physical interfaces for redundancy and increased bandwidth. The LAG will only be operational if at least one member interface is active and configured correctly.
No Operational Member Interfaces:
If no member interfaces are added or if the member interfaces are down, the LAG will remain in a down state, as shown in the exhibit for ae0 and ae1.
Resolution:
Verify that physical interfaces have been added to the LAG using commands like:

LAG Interface Status: In Juniper, the link status of the LAG depends on its member interfaces, which must be operational for the LAG to function.

NEW QUESTION 61
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 63
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your JN0-280 Exam with Our Prep Materials Via below:

https://www.certleader.com/JN0-280-dumps.html