Juniper

Exam Questions JN0-280

Data Center Associate (JNCIA-DC)

NEW QUESTION 1
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 2
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 3
Which statement is correct about IBGP?

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

Answer: B

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

Logical Full Mesh:

In an IBGP setup, routers do not re-advertise routes learned from one IBGP peer to another IBGP peer. This rule is in place to prevent routing loops within the AS.

To ensure full route propagation, alogical full meshis required, meaning every IBGP router must peer with every other IBGP router in the AS. This can be done either directly or via route reflection or confederation.

Physical Full Mesh Not Required:The physical topology does not need to be a full mesh, but the BGP peering relationships must form a logical full mesh. Techniques like route reflectors or BGP confederations can reduce the need for manual full-mesh peering.
Juniper Reference:

IBGP Configuration: IBGP logical full mesh requirements can be simplified usingroute reflectorsto avoid the complexity of manually configuring many IBGP peers.

NEW QUESTION 4

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 5
How many stages are shown in the exhibit?

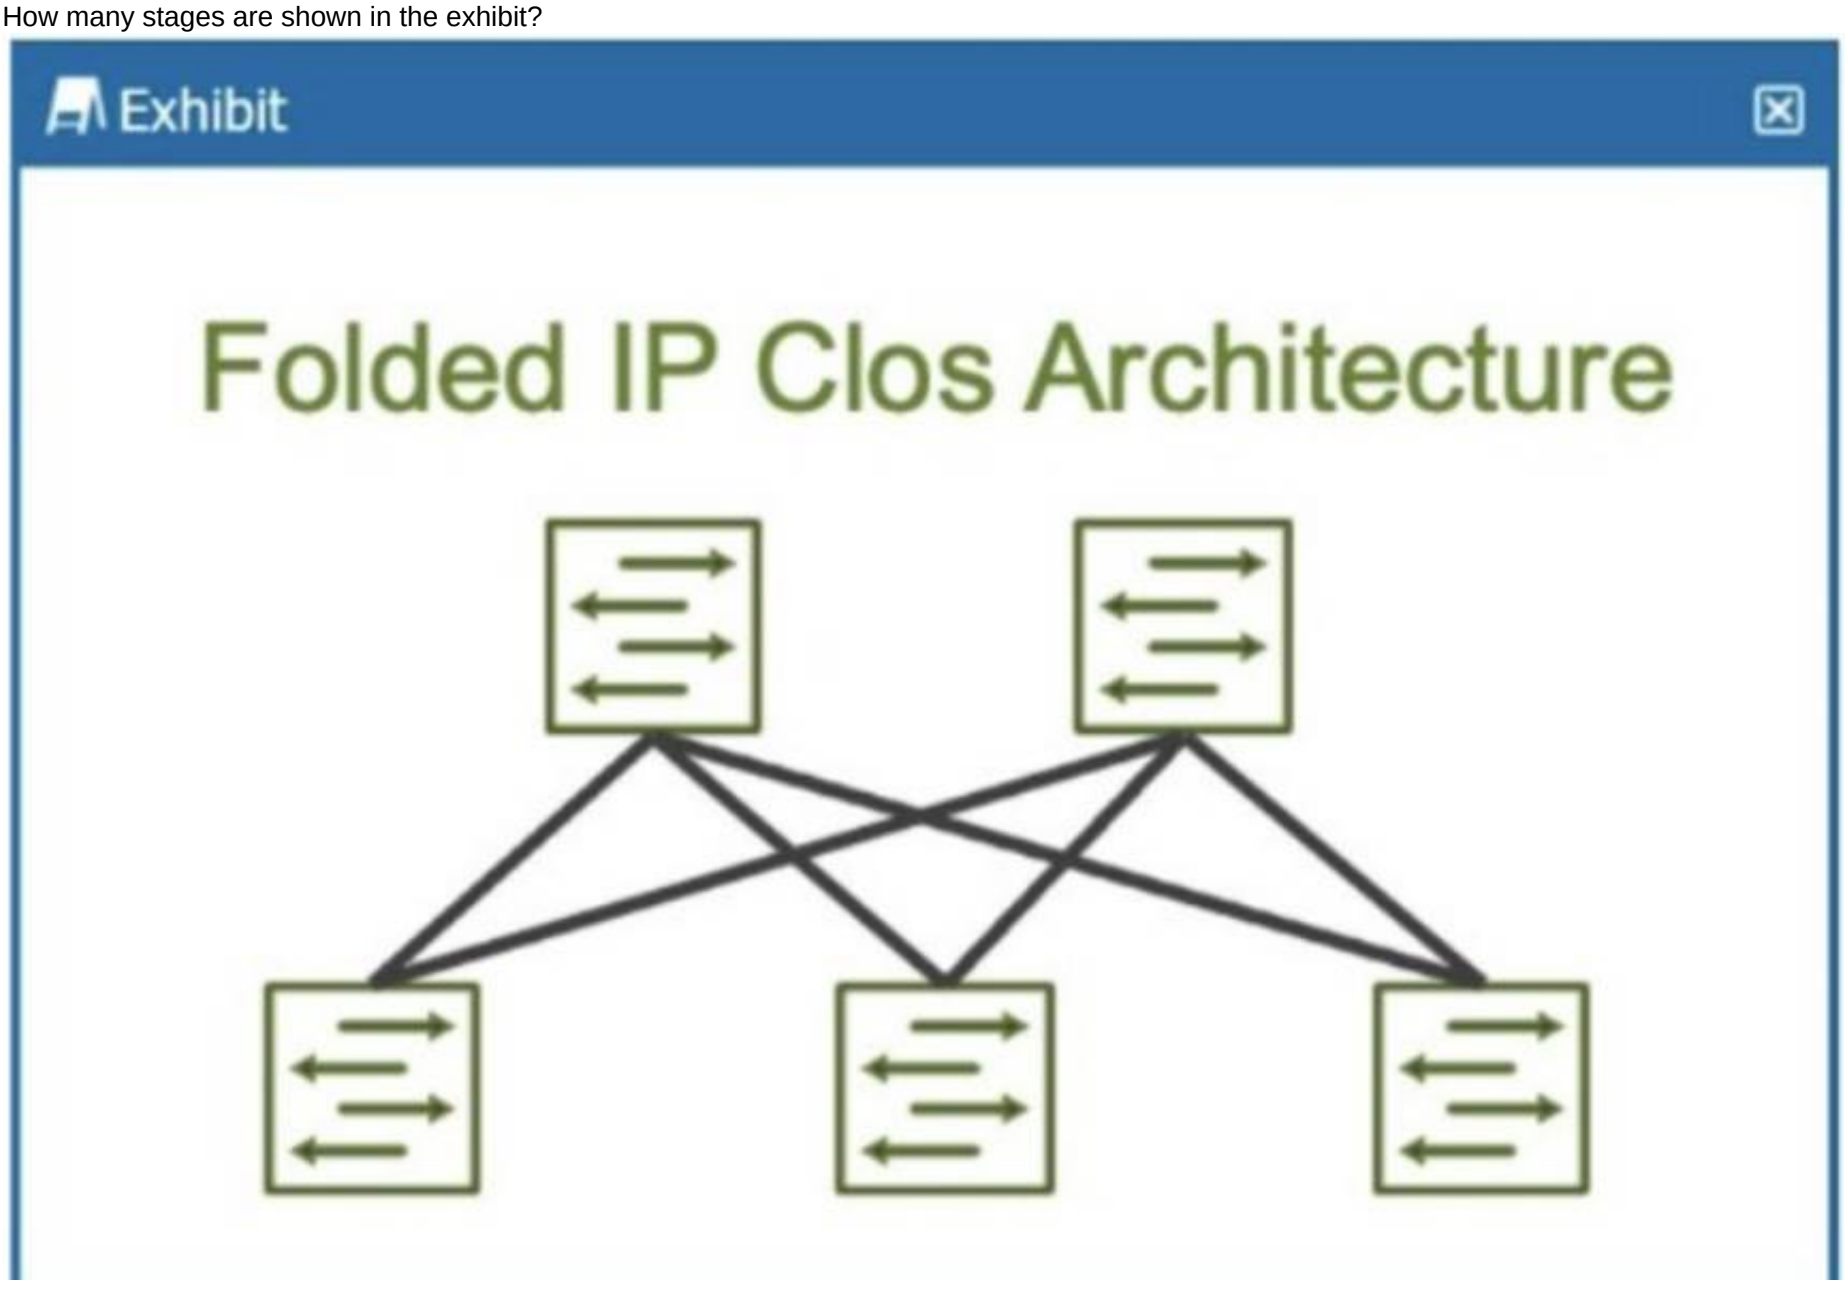

A. 2
B. 5
C. 6
D. 3

Answer: D

Explanation:
The exhibit shows aFolded IP Clos Architecture, which is also referred to as a3-stage Closnetwork design. This architecture typically consists of two layers of switches:
Spine Layer: The top row of switches.
Leaf Layer: The bottom row of switches.
Step-by-Step Breakdown:
Clos Architecture:A 3-stage Clos network has two types of devices: spine and leaf. In this design, each leaf switch connects to every spine switch, providing a high level of redundancy and load balancing.
Stage Explanation:
Stage 1: The first set of leaf switches.
Stage 2: The spine switches.
Stage 3: The second set of leaf switches.
TheFolded Closarchitecture shown here effectively "folds" the 3-stage design by combining the ingress and egress leaf layers into one, reducing it to two visible layers, but still maintaining the overall3-stage architecture.
Juniper Reference:
IP Clos Architecture: The 3-stage Clos design is commonly used in modern data centers for high availability, redundancy, and scalability.

NEW QUESTION 6
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 7
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 8
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 9
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 10
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 10
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 13
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 14
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 19
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 20
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 24
By default, which two statements are correct about BGP advertisements? (Choose two.)

A. BGP peers advertise routes received from EBGP peers to other IBGP peers.
B. BGP peers advertise routes received from IBGP peers to other IBGP peers.
C. BGP peers advertise routes from EBGP peers to other IBGP peers using its own address as thenext hop.
D. BGP peers advertise routes from IBGP peers to EBGP peers using its own address as the next hop.

Answer: AD

Explanation:
BGP (Border Gateway Protocol)has specific rules for route advertisement between peers.
Step-by-Step Breakdown:
EBGP to IBGP Route Propagation:
BGP peers advertise routes learned fromEBGP peerstoIBGP peerswithin the same AS. This ensures that routes learned from external networks are propagated internally within the AS.
IBGP to EBGP Route Propagation:
Routes learned fromIBGP peerscan be advertised toEBGP peers, but when advertising these routes, the router uses its own IP address as thenext hop.
IBGP Split Horizon:
By default, IBGP peers do not advertise routes learned from one IBGP peer to another IBGP peer. This rule (IBGP split horizon) prevents routing loops within an AS.
Juniper Reference:
BGP Advertisement Rules: Junos adheres to BGP standards, where IBGP peers do not propagate routes to other IBGP peers, but EBGP peers receive IBGP routes with the advertising router as the next hop.

NEW QUESTION 29
When using spine and leaf fabric architectures, what is the role of each device? (Choose two.)

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

Answer: BD

Explanation:
In a spine-leaf fabric architecture, which is commonly used in data center designs, each device has a distinct role to ensure efficient and scalable network traffic flow.
Step-by-Step Breakdown:
Spine Nodes:
The spine nodes form the backbone of the fabric and are responsible for transit traffic between leaf nodes. They connect to every leaf switch and provide multiple paths for traffic between leaf nodes, ensuring redundancy and load balancing.
Leaf Nodes:
The leaf nodes are used for host connectivity. These switches connect to servers, storage, or edge routers. They also connect to the spine switches to reach other leaf switches.
Juniper Reference:
Spine-Leaf Architecture: In Juniper's IP fabric designs, spine switches handle inter-leaf communication, while leaf switches manage host and endpoint connectivity.

NEW QUESTION 31
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 35
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 40
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 41
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 43
.....

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