# Juniper

## Exam Questions JN0-280

Data Center Associate (JNCIA-DC)

# About Exambible

***Your Partner of IT Exam***

# Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

# Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

**NEW QUESTION 1**
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 2**
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 3**
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 4**
Which three technologies improve high availability and convergence in a data center network? (Choose three.)

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

**Answer:** ABE

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

- Graceful Restart (GR):
- GR allows routers to maintain forwarding state during control plane restarts, ensuring continuous packet forwarding while minimizing network disruptions.
- Bidirectional Forwarding Detection (BFD):
- BFD provides fast detection of path failures, allowing routing protocols to converge quickly by detecting link failures much faster than traditional timers.
- Link Aggregation Group (LAG):
- LAG increases both redundancy and bandwidth by combining multiple physical links into one logical link, providing load balancing and fault tolerance.

Juniper Reference:

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

**NEW QUESTION 5**
Within your router, you want to verify that you are learning routes from a remote BGP peer at IP address 10.10.100.1. Which command would satisfy the requirement?

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

**Answer:** A

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

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

show route receive-protocol bgp 10.10.100.1

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

Juniper Reference:

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

**NEW QUESTION 6**
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 7**
How many stages are shown in the exhibit?

Exhibit

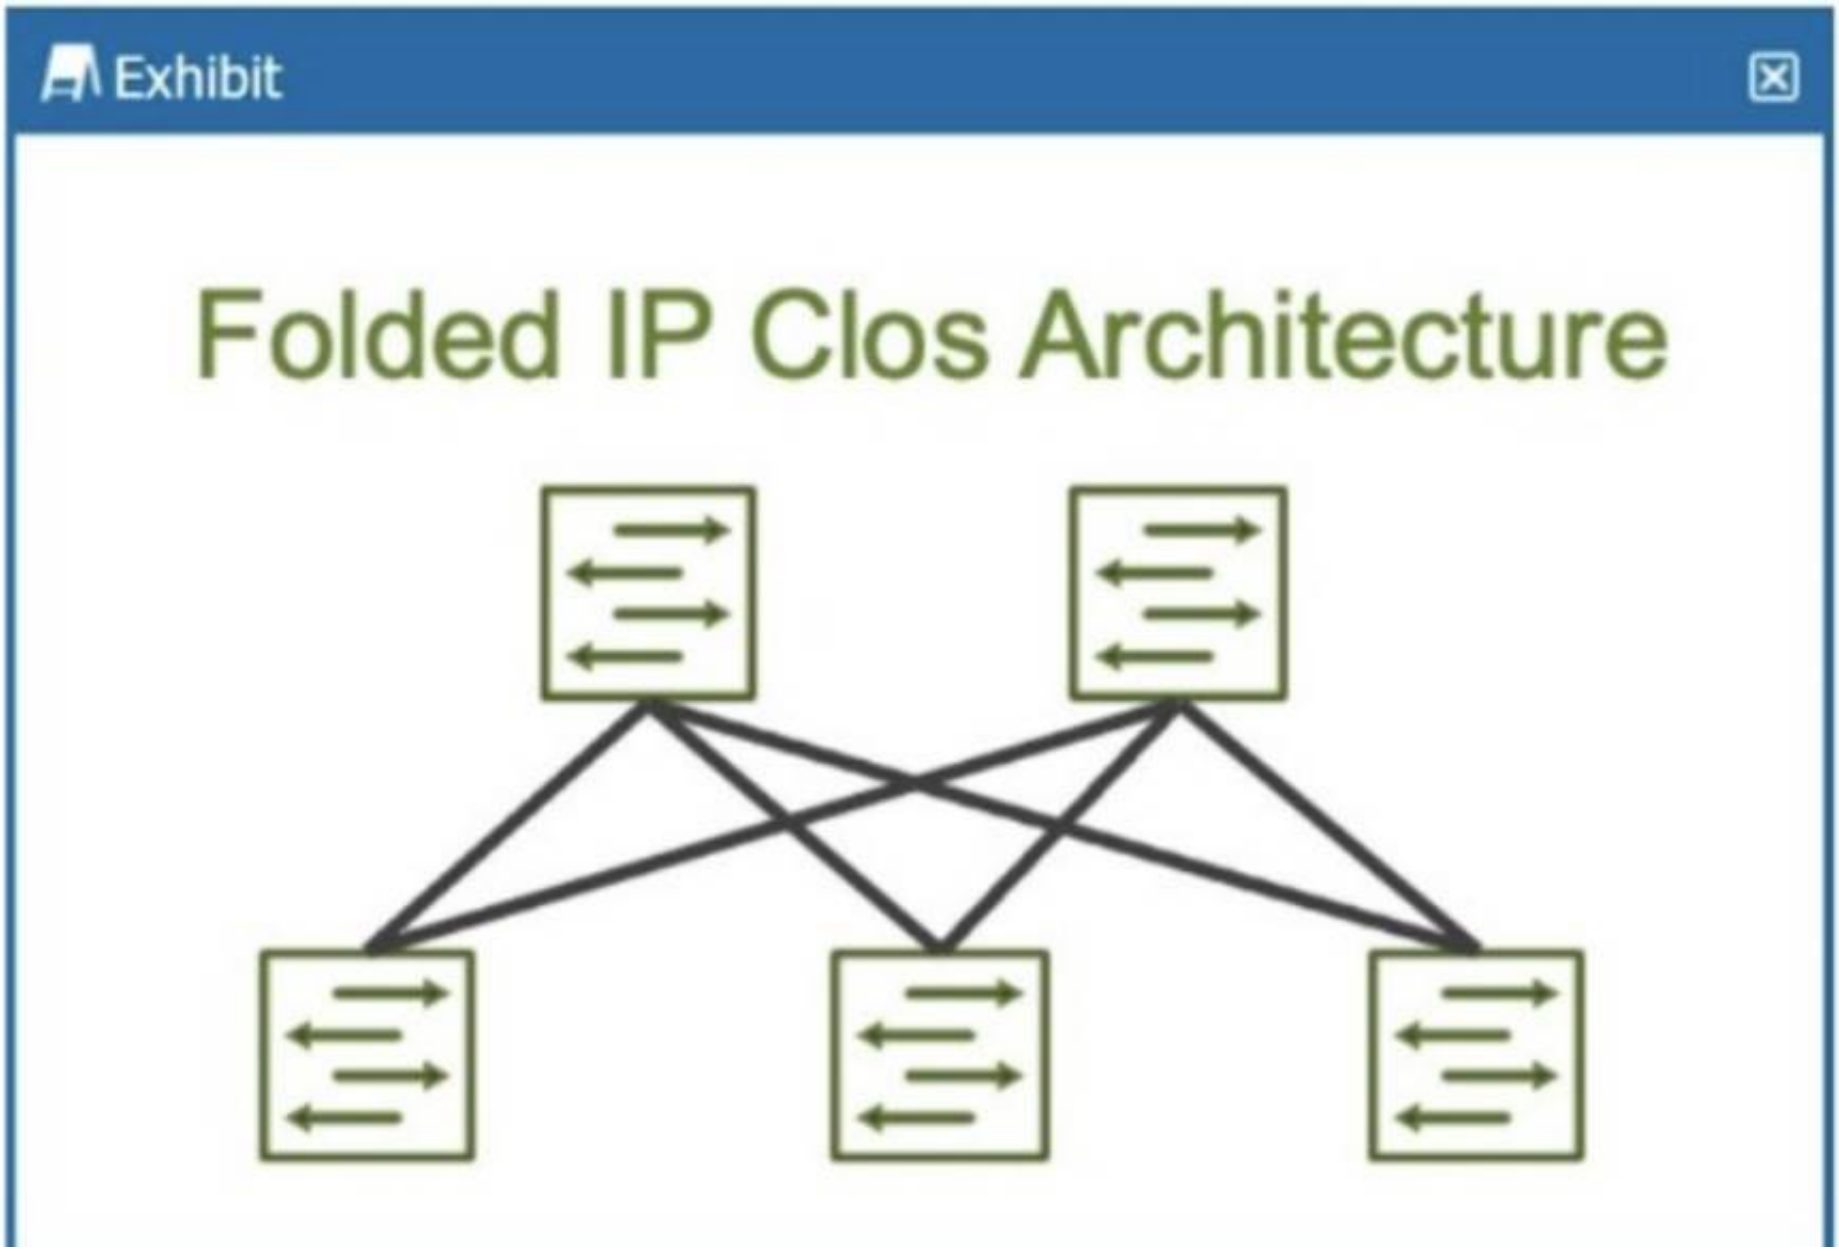


A. 2
B. 5
C. 6
D. 3

**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 8**
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 9**
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 10**
What are two consequences of having all network devices in a single collision domain? (Choose two.)

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

**Answer:** CD

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

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

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

**Answer:** C

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

**NEW QUESTION 13**
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 15**
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 20**
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 23**
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 28**
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 30**
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**
What are two device roles in a five-member Virtual Chassis? (Choose two.)

A. PFE
B. Control-board
C. Line card
D. Routing-engine

**Answer:** CD

**Explanation:**

In aVirtual Chassis(VC) configuration, multiple Juniper switches are interconnected to form a single logical device. Each member switch in the Virtual Chassis plays a specific role.
Step-by-Step Breakdown:
Line Card Role:
Member switches acting asline cardsprovide additional ports for traffic forwarding but do not perform control or routing functions. These switches depend on the routing engine to handle control-plane tasks.
Routing Engine Role:
A switch in therouting-enginerole is responsible for control-plane operations such as routing protocol management and control of the Virtual Chassis.
Virtual Chassis Roles:
Master Routing Engine: Handles control-plane functions and manages the entire Virtual Chassis.
Backup Routing Engine: Takes over if the master fails.
Line Card: Provides additional ports and handles data-plane operations.
Juniper Reference:
Virtual Chassis: In a five-member Virtual Chassis, multiple switches act asline cards, while one or more switches are designated as therouting engines(master and backup).

**NEW QUESTION 39**
Exhibit:

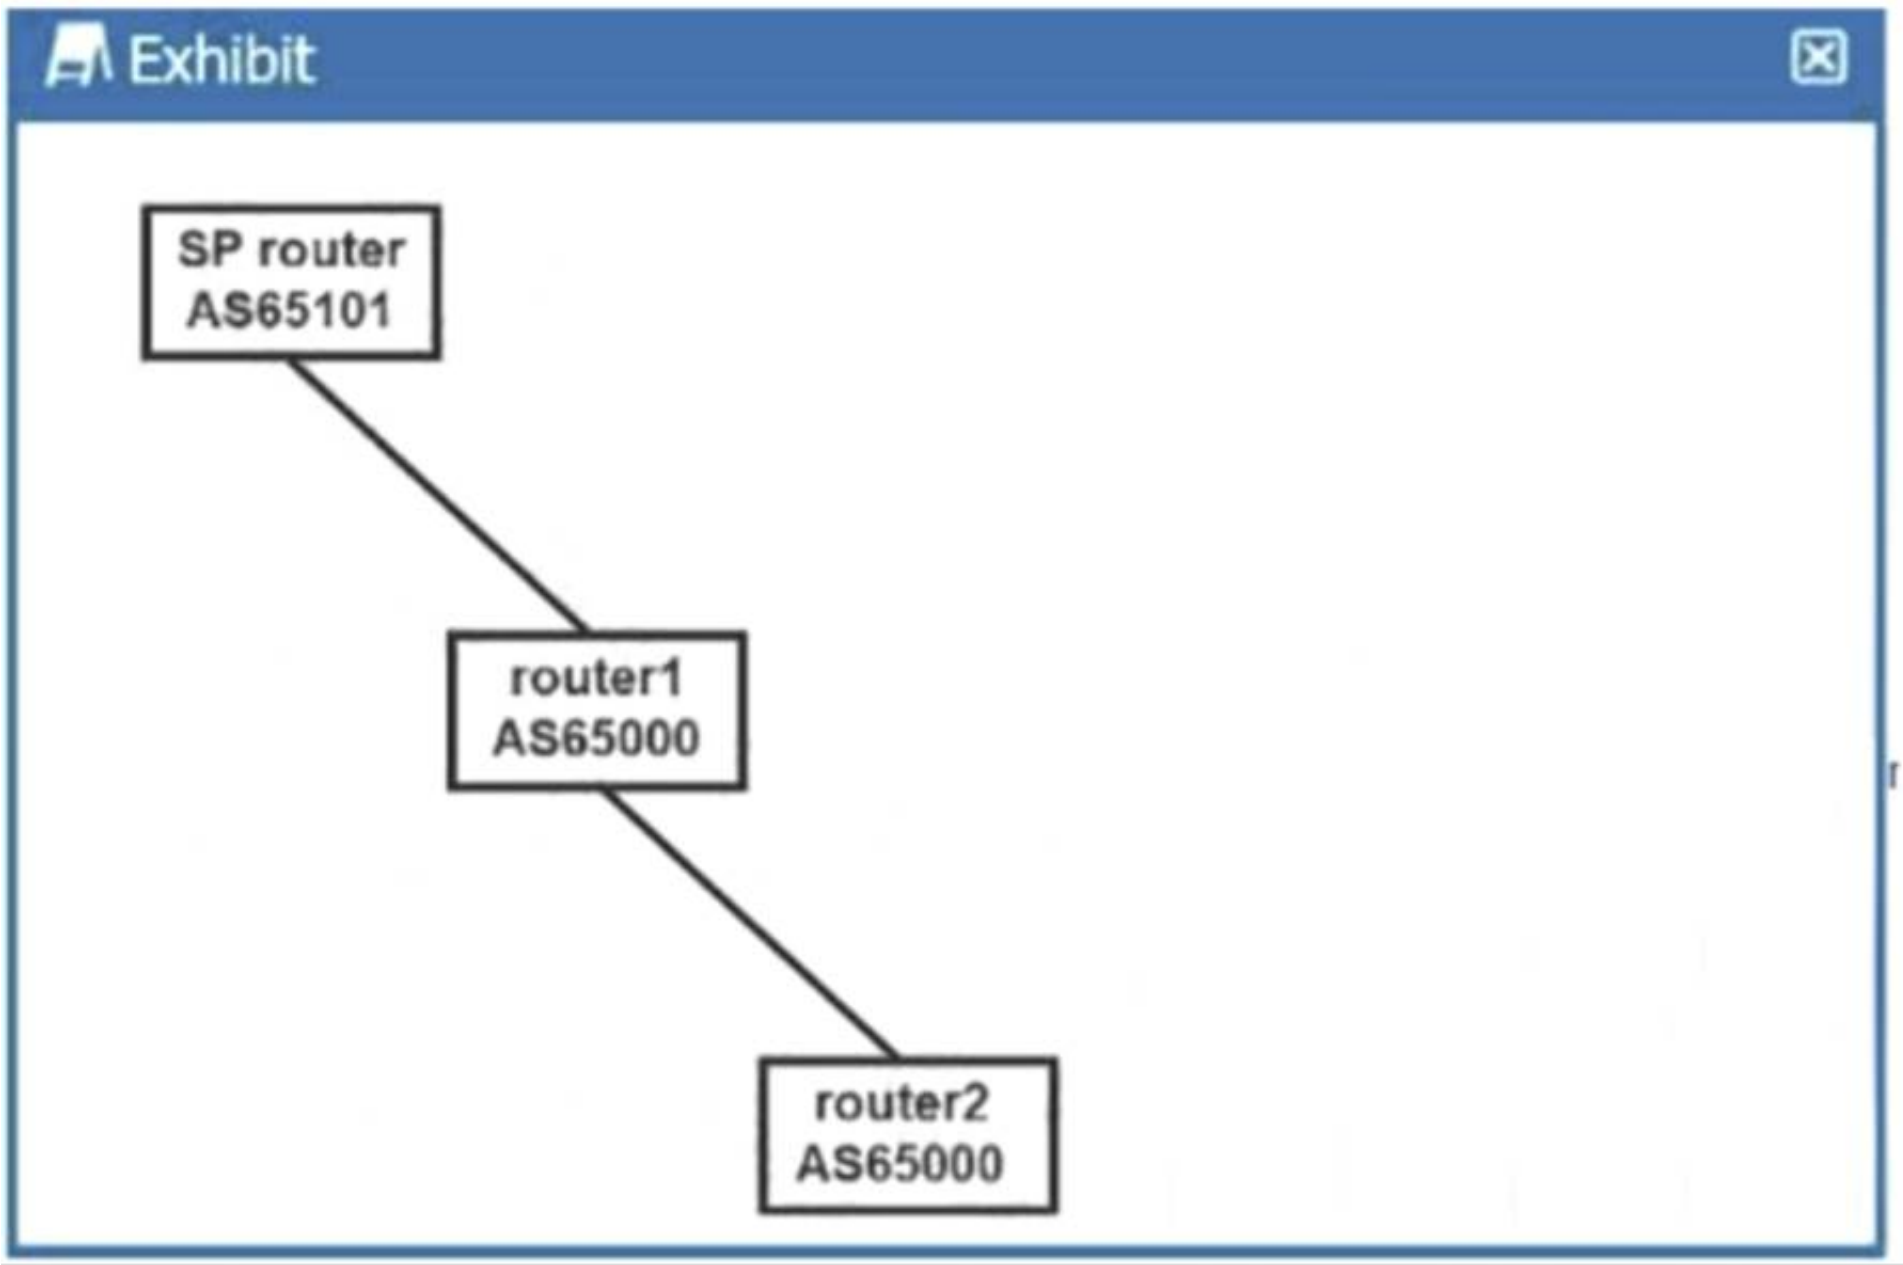


Referring to the exhibit, which two statements are correct about default BGP advertisements? (Choose two.)

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

**Answer:** BD

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

**NEW QUESTION 44**

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 49**
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 53**
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 54**
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 55**
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 60**
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 61**
When troubleshooting an OSPF neighborship, you notice that the router stopped at the ExStart state. What is the cause of this result?

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

**Answer:** D

**Explanation:**
When an OSPF (Open Shortest Path First) neighborship is stuck in the ExStart state, it usually points to a mismatch in Maximum Transmission Unit (MTU) settings between two routers trying to establish the adjacency. The ExStart state is where OSPF routers negotiate the master-slave relationship and exchange DBD (Database Description) packets.
Step-by-Step Breakdown:
OSPF Neighbor States: OSPF goes through several states to establish an adjacency with a neighbor:
Down: No hello packets have been received.
Init: Hello packets are received, but bidirectional communication isn't confirmed.
2-Way: Bidirectional communication is established.
ExStart: The routers are negotiating who will be the master and who will be the slave, and begin to exchange DBD packets.
Exchange: The routers start exchanging the database information.
Loading: The routers process the Link-State Advertisements (LSAs).
Full: The adjacency is fully established.
MTU Mismatch Issue:
During the ExStart state, both OSPF routers must agree on their MTU values. If there is an MTU mismatch between the two routers, OSPF neighbors will fail to move from the ExStart to the Exchange state. The router with the larger MTU setting will not accept DBD packets from the router with a smaller MTU because the packets may exceed the smaller MTU size.
In Juniper devices, this behavior can be identified by examining the MTU settings using the show interfaces command and ensuring both routers have matching MTU configurations. To resolve this issue, either match the MTU settings on both routers or configure OSPF to ignore MTU mismatches using the command set protocols ospf ignore-mtu.

**NEW QUESTION 65**
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 70**
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 71**
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 72**
.....

# Relate Links

**100% Pass Your JN0-280 Exam with Exambible Prep Materials**

https://www.exambible.com/JN0-280-exam/

# Contact us

**We are proud of our high-quality customer service, which serves you around the clock 24/7.**

**Viste -** https://www.exambible.com/