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 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 4
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 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 statement is correct about an IRB interface?

A. An IRB interface switches traffic within the same VLAN.
B. An IRB interface trunks together VLANs on different switches.
C. An IRB interface is a physical Layer 3 interface that connects VLANs together.
D. An IRB interface is a Layer 3 interface that can be used to route between VLANs.

Answer: D

Explanation:
AnIRB (Integrated Routing and Bridging)interface provides routing functionality between VLANs at Layer 3, allowing devices in different VLANs to communicate with each other.
Step-by-Step Breakdown:

- IRB Functionality:
- The IRB interface enables routing between different VLANs by acting as a Layer 3 gateway.

Traffic within the same VLAN is handled by Layer 2 switching, while traffic between VLANs is routed through the IRB interface.

- Layer 3 Routing Between VLANs:
- Each VLAN can be assigned an IP address on the IRB interface, which allows traffic to flow between VLANs based on Layer 3 IP routing.

Juniper Reference:

- IRB Interface Configuration: Juniper supports IRB for inter-VLAN routing on devices like the EX and QFX series switches, facilitating Layer 3 communication in data centers.

NEW QUESTION 7
What is the behavior of the default export policy for OSPF?

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

Answer: B

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

NEW QUESTION 8
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 9
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 10
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 10
Referring to the exhibit, why are the BGP routes hidden?

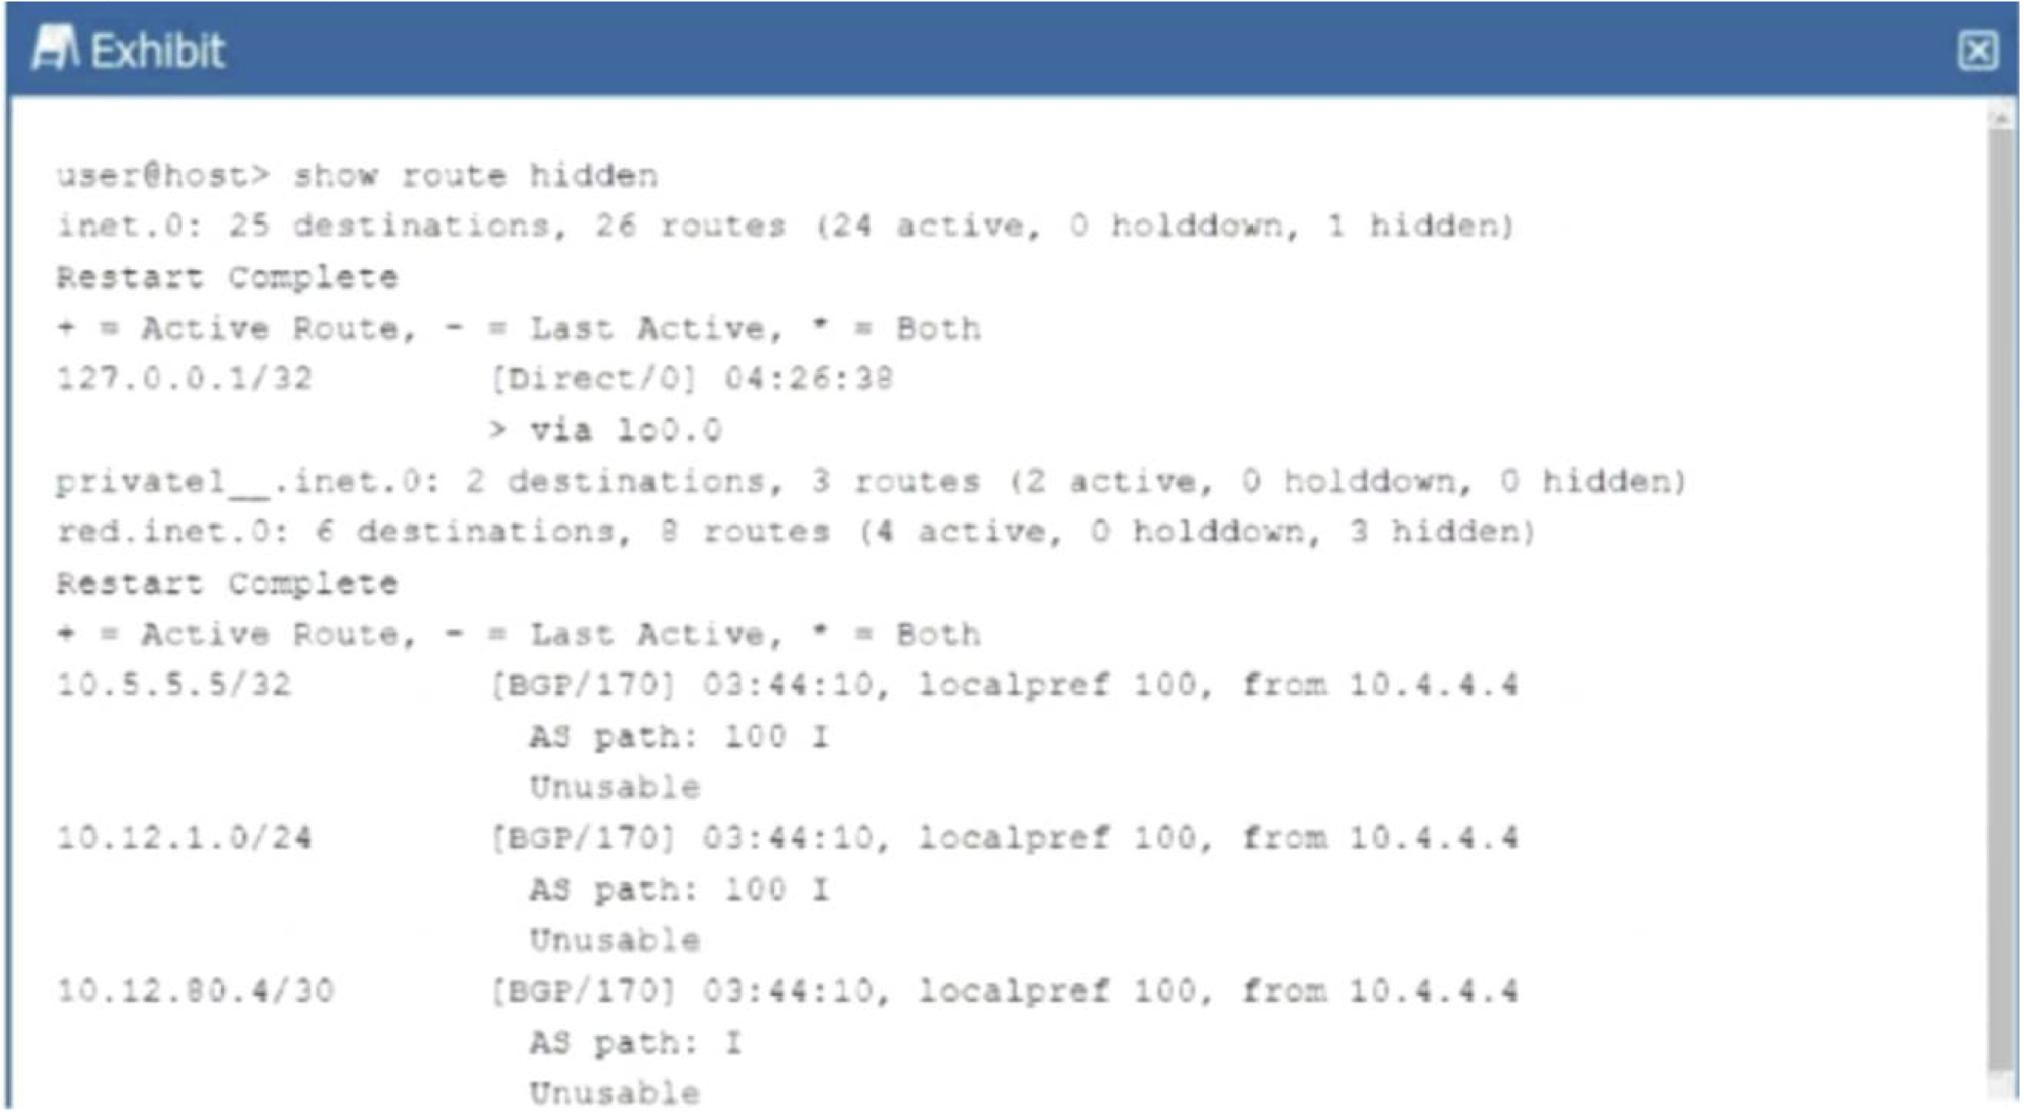

```
user@host> show route hidden
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
127.0.0.1/32          [Direct/0] 04:26:38
                      > via lo0.0
private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.5.5.5/32           [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                        AS path: 100 I
                        Unusable
10.12.1.0/24          [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                        AS path: 100 I
                        Unusable
10.12.80.4/30         [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                        AS path: I
                        Unusable
```

A. Load balancing is not enabled.
B. There are too many hops to the destination.
C. The BGP next hop is unreachable.
D. Other routes are selected because of better metrics.

Answer: C

Explanation:
In the exhibit, the BGP routes are marked ashidden. This typically happens when the routes are not considered valid for use, but they remain in the routing table for reference. One common reason for BGP routes being hidden is that thenext hopfor these routes is unreachable.
Step-by-Step Breakdown:
BGP Next Hop:In BGP, when a route is received from a neighbor, thenext hopis the IP address that must be reachable for the route to be used. If the next hop is unreachable (i.e., the router cannot find a path to the next-hop IP), the route is marked as hidden.
Analyzing the Exhibit:The exhibit shows that the BGP next hop for all hidden routes is 10.4.4.4. If this IP is unreachable, the BGP routes from that neighbor will not be considered valid, even though they appear in the routing table.
Verification:
Use the command show route 10.4.4.4 to check if the next-hop IP is reachable.
If the next-hop is not reachable, the BGP routes will be hidden. Resolving the next-hop reachability issue (e.g., fixing an IGP route or an interface) will allow the BGP routes to become active.
Juniper Reference:
Junos Command: show route hidden displays routes that are not considered for forwarding.
Troubleshooting: Check the next hop reachability for hidden BGP routes using show route .

NEW QUESTION 14
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 16
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 18
MACsec provides protection against which two types of threats? (Choose two.)

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

Answer: BD

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

NEW QUESTION 23
......