Juniper

Exam Questions JN0-280

Data Center Associate (JNCIA-DC)

NEW QUESTION 1
You are creating an IP fabric underlay and want to use OSPF as your routing protocol. In this scenario, which statement is correct?

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

Answer: C

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

OSPF and Interface Speeds:OSPF calculates the cost of a link based on its bandwidth. The default cost calculation in OSPF is:
A black text on a white background Description automatically generated

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

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

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

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

NEW QUESTION 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
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 5
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 6
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 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 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
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 10
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
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 12
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 15
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 20
Referring to the exhibit, why are the BGP routes hidden?

Exhibit

```
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 21
Which statement is correct about member interfaces when creating a LAG?

A. The interface's MTU settings must match on all member interfaces.
B. The interface's duplex settings and link speed must be the same on all member interfaces.
C. Member interfaces must all be allocated on the same chassis when using a Virtual Chassis.
D. Member interfaces must all be allocated on the same PFE.

Answer: B

Explanation:
When creating aLAG (Link Aggregation Group)in Junos, theduplex settingsandlink speedmust be the same across all member interfaces.
Step-by-Step Breakdown:
LAG Overview:A LAG combines multiple physical interfaces into a single logical interface to increase bandwidth and provide redundancy. All member links must act as a single cohesive unit.
Interface Requirements:
Duplex: All member interfaces must operate in the same duplex mode (either full-duplex or half-duplex). Mismatched duplex settings can cause performance issues, packet drops, or interface errors.
Link Speed: All interfaces in the LAG must have the same link speed (e.g., all interfaces must be 1 Gbps or 10 Gbps). Mismatched speeds would prevent the interfaces from functioning correctly within the LAG.
Configuration and Validation: Ensure that all member interfaces have identical settings before adding them to the LAG. These settings can be checked using the show interfaces command, and the LAG can be configured using:
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ge-0/0/1 ether-options 802.3ad ae0
Juniper Reference:
LAG Configuration: Duplex and link speed must be consistent across member interfaces to ensure proper LAG operation in Juniper devices.

NEW QUESTION 26
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 29
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 33
.....

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