Free Questions for JN0-280

Shared by Mcclure on 30-09-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

What is the behavior of the default export policy for OSPF?

Options:

A- Accept all routes.

B- Reject all routes.

C- Redistribute all routes.

D- Forward all routes.

Answer:

B

Explanation:

In Junos, the default export policy for OSPF is to reject all routes from being exported.

Step-by-Step Breakdown:

Default Export Policy:

By default, OSPF in 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 custom export policy that 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.

Question 2

Question Type: **MultipleChoice**

You are troubleshooting a downed BGP session.

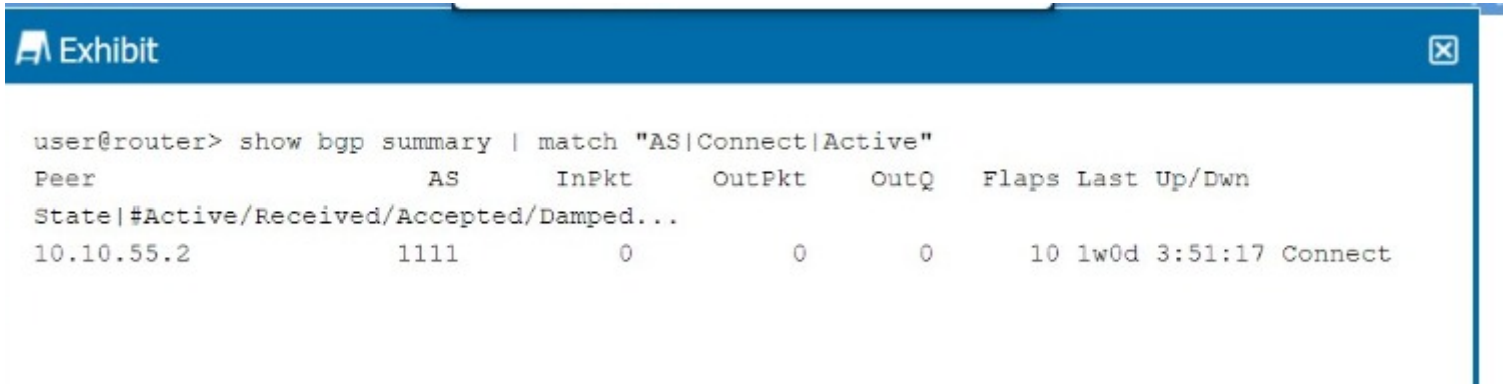

Referring to the exhibit, what is the cause of the problem?

Options:

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.

Question 3

Question Type: MultipleChoice

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

Options:

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:

B, D

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.

Question 4

Question Type: **MultipleChoice**

Exhibit:

Exhibit

```
[edit routing-options]
user@Router# show
static {
    route 0.0.0.0/0 {
        next-hop 172.25.11.254;
        qualified-next-hop 172.25.11.200 {
            preference 140;
        }
    }
}
```

Referring to the exhibit, what is the route preference of the 172.25.11.254 next hop?

Options:

A- 5

B- 10

C- 130

D- 140

Answer:

A

Explanation:

In the exhibit, we see two next-hop addresses for the default static route (0.0.0.0/0):

The first next hop is 172.25.11.254, with no specified preference.

The second next hop is 172.25.11.200, with a specified preference of 140.

Step-by-Step Breakdown:

Default Static Route Preference:

If no preference is explicitly set for a next hop in Junos, it defaults to 5 for static routes.

Determining Preference:

In this case, the next hop 172.25.11.254 does not have an explicit preference defined, so it will use the default value of 5. The second next hop has a preference of 140, which is higher, meaning it will only be used if the primary next hop is unavailable.

Juniper Reference:

Static Route Preference: In Junos, the default preference for static routes is 5, and this value is applied unless overridden by the preference parameter.

Question 5

Question Type: **MultipleChoice**

Exhibit:

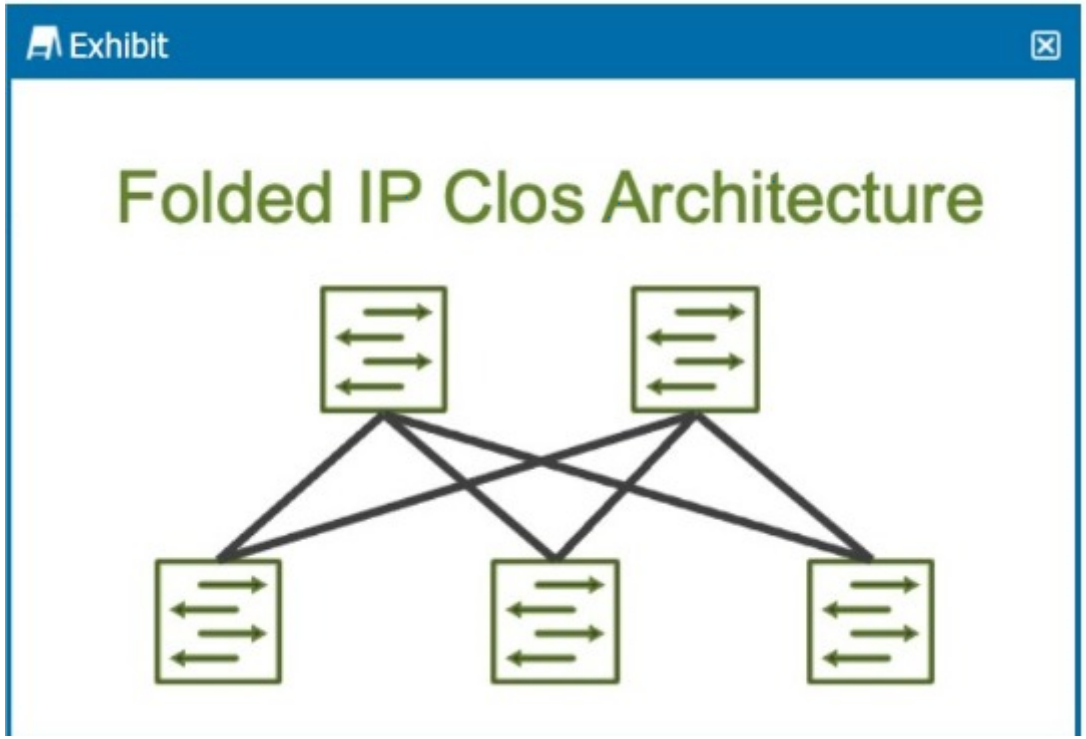

How many stages are shown in the exhibit?

Options:

A- 2

B- 5

C- 6

D- 3

Answer:

D

Explanation:

The exhibit shows a Folded IP Clos Architecture, which is also referred to as a 3-stage Clos network 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.

The Folded Clos architecture 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 overall 3-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.

Question 6

Question Type: **MultipleChoice**

Which route is preferred by the Junos OS software routing tables?

Options:

A- Static

B- Aggregate

C- Direct

D- BGP

Answer:

C

Explanation:

In Junos OS, direct routes are 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 a preference 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 routes are always preferred over other routes, ensuring that the router forwards traffic through locally connected networks.

Question 7

Question Type: MultipleChoice

Which two statements are correct about aggregate routes and generated routes? (Choose two.)

Options:

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:

A, C

Explanation:

Aggregate routes and generated routes are used to create summarized routes in Junos, but they behave differently in terms of forwarding.

Step-by-Step Breakdown:

Aggregate Routes:

An aggregate route summarizes 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:

A generated route also summarizes specific routes, but it has a forwarding next hop that 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.

Question 8

Question Type: **MultipleChoice**

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

Options:

A- Access layer

B- Modular layer

C- Aggregation layer

D- Core layer

E- Function layer

Answer:

A, C, D

Explanation:

In legacy hierarchical network design, three key layers are used to create a scalable and structured network:

Step-by-Step Breakdown:

Access Layer:

The access layer is where end devices, such as computers and IP phones, connect to the network. It typically involves switches that provide connectivity for devices at the edge of the network.

Aggregation Layer (Distribution Layer):

The aggregation layer (also called the distribution layer) aggregates traffic from multiple access layer devices and applies policies such as filtering and QoS. It also provides redundancy and load balancing.

Core Layer:

The core layer provides high-speed connectivity between aggregation layer devices and facilitates traffic within the data center or between different network segments.

Juniper Reference:

Legacy Hierarchical Design: Juniper networks often follow the traditional three-layer design (Access, Aggregation, and Core) to ensure scalability and high performance.

Question 9

Question Type: **MultipleChoice**

You are creating an IP fabric underlay and want to use OSPF as your routing protocol.

In this scenario, which statement is correct?

Options:

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 an IP fabric underlay using OSPF as the routing protocol, consistent interface speeds are 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:

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

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.