

Exam : JN0-281

**Title : Data Center, Associate
(JNCIA-DC)**

Version : DEMO

1. A switch receives an Ethernet frame that contains source and destination MAC addresses that are not in the Ethernet switching table.

In this scenario, which two actions does the switch perform? Choose two.

- A. The switch floods the frame out every port in the broadcast domain except the ingress port.
- B. The switch adds the source MAC address to the Ethernet switching table.
- C. The switch drops the frame.
- D. The switch forwards the frame to the Routing Engine.

Answer: A, B

Explanation:

On Junos-based Ethernet switching platforms used in data centers, Layer 2 forwarding is handled primarily in the forwarding plane. When a switch receives a frame, it first performs source MAC learning. If the source MAC address is not already present in the Ethernet switching table for that VLAN or bridge domain, the switch creates a new entry that maps the source MAC to the ingress interface and the associated VLAN context. This learning step is fundamental to building the MAC table dynamically and enables efficient forwarding for subsequent frames destined back to that source.

Next, the switch attempts to forward the frame based on the destination MAC lookup in the same VLAN or bridge domain. Because the destination MAC is also not in the Ethernet switching table, the frame is treated as an unknown unicast. The default behavior for unknown unicast in a Layer 2 broadcast domain is to flood the frame out all other interfaces that belong to that VLAN or bridge domain, excluding the ingress interface. Flooding ensures the frame has the best chance of reaching the correct destination host. When the destination responds, the switch then learns that MAC address as a source on the return traffic, allowing future traffic to be forwarded as known unicast instead of flooded.

The switch does not drop the frame by default, and it does not forward the frame to the Routing Engine because this is normal Layer 2 bridging behavior, not a control-plane routing decision.

2. Which two statements are correct about configuring VLANs? Choose two.

- A. You must assign an IRB interface to each VLAN.
- B. You must assign a VLAN name or ID and a Layer 2 interface to the VLAN.
- C. You can assign one or more VLANs to a trunk mode interface.
- D. You can assign one or more VLANs to an access mode interface.

Answer: B, C

Explanation:

On Junos switching platforms commonly used in data centers, a VLAN is a Layer 2 construct that defines a broadcast domain. To make a VLAN usable, you define the VLAN using a name and typically a VLAN ID, then associate Layer 2 interfaces with it so traffic entering those interfaces is placed into that VLAN. Without membership on interfaces, the VLAN exists in configuration but does not carry user traffic, because no ports participate in that broadcast domain.

Trunk mode interfaces are specifically designed to carry traffic for multiple VLANs over a single physical link, such as between switches, to servers using tagging, or to other network devices that understand VLAN tags. In Junos, trunking is implemented by allowing a list of VLAN IDs on the trunk so the interface accepts and forwards frames for those VLANs. This makes statement C correct.

An IRB interface is not mandatory for every VLAN. IRB is used when you want Layer 3 routing for a VLAN, typically to provide a default gateway and enable inter VLAN routing. Pure Layer 2 VLANs do not require IRB, which makes statement A incorrect.

Access mode interfaces are intended to connect to a single endpoint and carry traffic for a single VLAN, so assigning multiple VLANs to an access interface is not correct in standard access mode behavior, making statement D incorrect.

3.What is a function of an integrated routing and bridging IRB interface?

- A. to route traffic between different VLANs
- B. to encrypt traffic between network segments
- C. to bridge traffic within the same VLAN
- D. to provide Network Address Translation NAT

Answer: A

Explanation:

In Junos-based data center switching, an IRB interface is the Layer 3 gateway that is logically associated with a Layer 2 VLAN or bridge domain. The VLAN provides Layer 2 bridging inside the broadcast domain, while the IRB interface provides the routed interface that enables hosts in that VLAN to reach destinations outside their local subnet. This is the standard mechanism used for inter-VLAN routing on Juniper switches and for providing default gateway services to servers connected to access ports or VLAN-tagged trunks.

Operationally, endpoints in a VLAN use the IRB interface IP address as their default gateway. Frames destined to a remote subnet are bridged at Layer 2 to the IRB gateway MAC address, and then the packet is routed at Layer 3 based on the routing table. This allows a single device to perform both bridging within the VLAN and routing between VLANs or to other routed interfaces, which is why the concept is called integrated routing and bridging.

IRB does not encrypt traffic and does not provide NAT by itself; those functions are typically associated with security services features and firewall platforms. IRB is also not the mechanism that performs pure bridging within the same VLAN, because bridging is handled by the VLAN or bridge domain and the Ethernet switching table.

4.You are asked to ensure that traffic and routing information is not interrupted if your primary Routing Engine fails or switches to the backup Routing Engine.

In this scenario, which high availability feature will accomplish this behavior?

- A. nonstop active routing NSR
- B. graceful Routing Engine switchover GRES
- C. link aggregation control protocol LACP
- D. bidirectional forwarding detection BFD

Answer: A

Explanation:

Nonstop active routing is the Junos high availability feature designed to keep routing protocol operation and routing information continuous across a Routing Engine switchover on platforms with redundant Routing Engines. With NSR enabled, the control-plane routing state is replicated so that protocol sessions and routing information can remain stable when the device transitions from the primary to the backup Routing Engine. The goal is a transparent switchover that minimizes or eliminates routing reconvergence caused by a Routing Engine failure.

This is especially important in data center environments where routing stability underpins EVPN VXLAN control-plane operation, underlay BGP or OSPF adjacencies, and service reachability. By maintaining

the routing protocol process state across the switchover, NSR helps prevent neighbor resets and reduces churn in the routing table, which directly protects application traffic paths from disruption that would otherwise occur during a control-plane restart.

GRES is closely related but has a different focus: it preserves forwarding and certain kernel and interface states so that packet forwarding can continue, but by itself it does not preserve the full routing protocol control plane. That is why NSR is the best match when the requirement explicitly includes routing information continuity in addition to traffic continuity. LACP and BFD are valuable availability tools, but they address link bundling and fast failure detection, not Routing Engine stateful failover.

5.You have a problem bringing up an aggregated Ethernet interface between a spine and a leaf.

```
[edit]
user@leaf1# show interfaces ae1
description "to spine1";
mtu 9000;
aggregated-ether-options {
    lacp {
        periodic fast;
    }
}
unit 0 {
    family ethernet-switching {
        interface-mode trunk;
        vlan {
            members [ vn100 vn120 ];
        }
    }
}
```

```
[edit]
user@spine1# show interfaces ae2
description "to leaf1";
mtu 9100;
aggregated-ether-options {
    lacp {
        periodic fast;
    }
}
unit 0 {
    family ethernet-switching {
        interface-mode trunk;
        vlan {
            members [ vn100 vn101 vn120 vn122 ];
        }
    }
}
```

Referring to the exhibit, what is the problem?

A. The active statement must be added to LACP under the aggregated-ether-options hierarchy on one or both sides.

- B. The ae interface numbers are not consistent.
- C. The leaf-and-spine VLAN memberships are not consistent and should be changed to include the additional VLANs defined on spine1.
- D. The leaf-and-spine MTUs are not consistent.

Answer: A

Explanation:

An aggregated Ethernet interface that uses LACP requires at least one side to actively initiate LACP negotiations. In the exhibit, both devices have LACP configured only with periodic fast, but neither side explicitly enables LACP active mode. When both ends operate in passive behavior, each side waits for the other to send LACP Data Units, and no negotiation begins. As a result, the LACP state does not progress to collecting and distributing, and the aggregated link fails to form as expected. Adding the active statement under the LACP hierarchy on one or both ends ensures that LACP frames are transmitted and the bundle can be negotiated and brought up.

The other options are not the root cause for bringing the bundle up. The aggregated Ethernet interface number does not need to match across devices because the bundle is locally significant on each system. VLAN membership differences on a trunk do not prevent LACP from establishing the aggregate; they only affect which tagged VLANs are allowed to pass once the link is operational. MTU differences can cause data plane issues such as fragmentation or drops for jumbo frames, but they do not typically prevent LACP formation because control frames are small and the physical link can still come up.