KillTest

高品質　より良いサービス



練習問題

https://www.killtest.jp

1年間無料アップデート

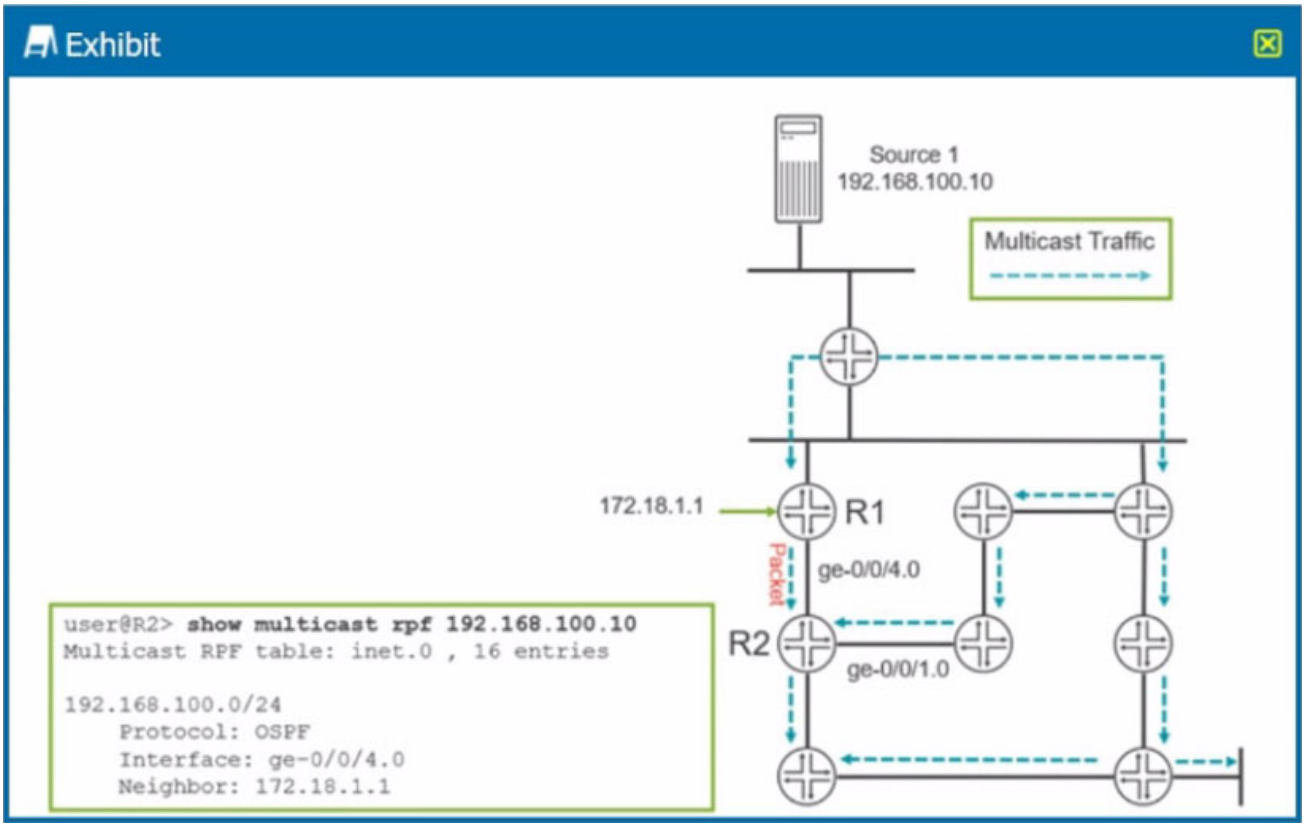
Exam : **JN0-650**

Title : Enterprise Routing and Switching, Professional (JNCIP-ENT)

Version : DEMO

1.You are asked to establish authentication for users connecting to the corporate network. You must ensure that only corporate devices that are identified by MAC addresses are allowed to authenticate and connect Authentication must be handled by a centralized database.
Which authentication method would you implement in this scenario?
A. MAC RADIUS
B. captive portal
C. single-secure supplicant mode
D. single supplicant mode
Answer: A
Explanation:
In this scenario, the requirement is to authenticate devices based on their MAC addresses using a centralized database.
MAC RADIUS (Option A): This is the standard Juniper method for authenticating "headless" or non-802.1X capable devices based on their hardware (MAC) address. The switch acts as a proxy, sending the device's MAC address as the username and password to a centralized RADIUS server (the centralized database). If the MAC address is found in the server's database, the device is granted access to the network. This perfectly matches the user's requirement for MAC-based authentication and a centralized database.
Captive Portal (Option B): This is a web-based authentication method requiring user interaction (login via a browser), which is not based on MAC address identification for the initial authentication phase.
Supplicant Modes (Options C & D): These define how multiple devices are handled on an 802.1X-enabled port, but they are not authentication methods themselves; they rely on 802.1X (EAP-based) or MAC-RADIUS being already configured.

2.Click the Exhibit button.



Referring to the exhibit, R2 receives a packet from the multicast source on its ge-0/0/4 0 interface.
Which two statements are correct in this scenario? (Choose two.)

A. R2 verifies the existing unicast routing table and determines that this interface is not on the reverse path back to the source

B. All interfaces other than ge-0/0/4.0 will be on the outgoing interface list.

C. The packet is discarded

D. R2 determines that the packet is coming in on an interface that is on the reverse path back to the source.

Answer: A, C

Explanation:

The exhibit illustrates a PIM multicast environment where router R2 has received a multicast packet from Source 1 (192.168.100.10) on its ge-0/0/4.0 interface. To prevent loops, PIM-SM uses a mechanism called Reverse Path Forwarding (RPF).

RPF Check Process (Option A): When a multicast packet arrives, the router performs an RPF check by looking up the source's IP address in its unicast routing table (typically inet.0).

The exhibit's command output, show multicast rpf 192.168.100.10, explicitly shows that for the source network $192.168.100.0/24$, the RPF interface is ge-0/0/1.0.

Because the packet actually arrived on ge-0/0/4.0 instead of the expected ge-0/0/1.0 interface, the RPF check fails.

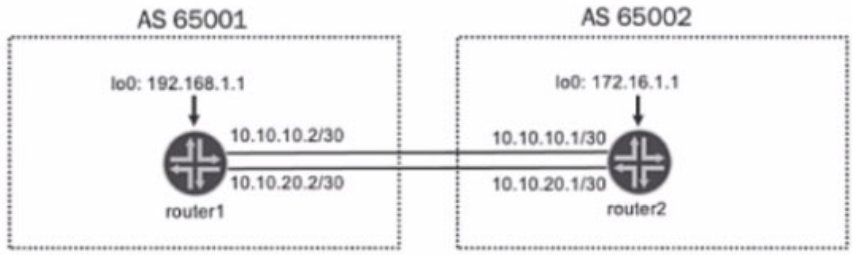
Packet Discard (Option C): According to standard PIM-SM operation in Junos OS 24.4, if a multicast packet fails the RPF check—meaning it arrived on an interface that the router does not use to reach the source via unicast—the packet is discarded. This is a fundamental loop-prevention mechanism that ensures multicast traffic is only accepted from the shortest path toward the source.

Option B is incorrect because the router will not add any interfaces to the Outgoing Interface List (OIL)

for a packet that fails the initial RPF check.
Option D is incorrect because the exhibit clearly shows the interface in the RPF table (ge-0/0/1.0) is different from the interface where the packet was received (ge-0/0/4.0).

3.Click the Exhibit button.

Exhibit



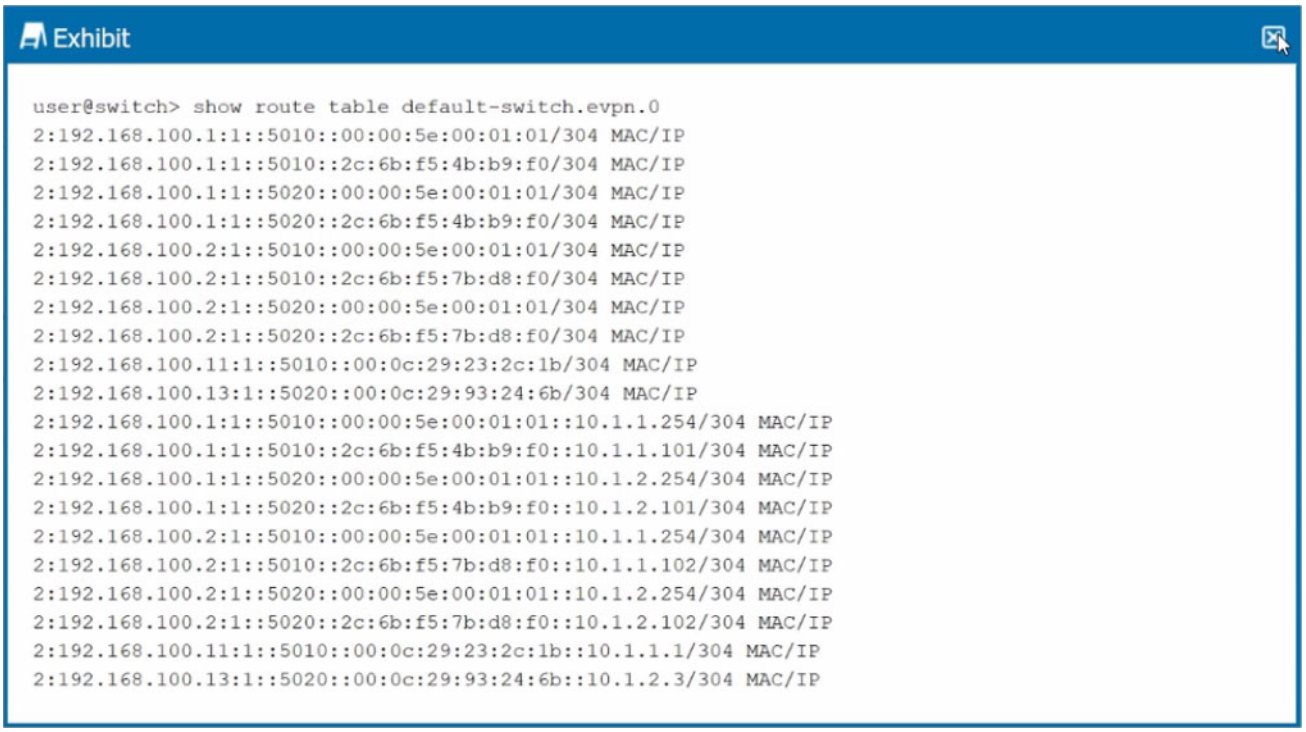

```
user@router1> show configuration protocols bgp group ext-peers
type external;
peer-as 65002;
neighbor 10.10.10.1;
neighbor 10.10.20.1;
}

user@router1> show bgp summary
[...]
Peer         AS      InPkt  OutPkt  OutQ  Flaps Last Up/Dwn State|#Active/Received/Accepted/Damped...
10.10.10.1   65002     53      53     0      0       22:48 4/4/4/0              0/0/0/0
10.10.20.1   65002     54      51     0      0       22:48 0/4/4/0              0/0/0/0
```

Referring to the exhibit, what will enable active routes from both peers?
A. Configure neighbor 172.16.1.1 multipath under the ext-peers hierarchy.
B. Configure multipath under the ext-peers hierarchy
C. Configure neighbor 172.16.1.1 under the ext-peers hierarchy.
D. Configure neighbor 10.10.20.2 multipath under the ext-peers hierarchy.
Answer: B
Explanation:
The exhibit shows router1 (AS 65001) connected to router2 (AS 65002) via two parallel physical links. The show bgp summary output indicates that sessions are established with two neighbors: 10.10.10.1 and 10.10.20.1. Currently, for the second neighbor (10.10.20.1), there are 0 active routes despite having 4 accepted routes, which indicates that BGP has selected only one "best path" via the first neighbor for forwarding.
BGP Best Path Selection: By default, BGP only selects a single best path for any given destination prefix and installs that one path into the forwarding table. In a topology with parallel links to the same AS, this leads to underutilization of available bandwidth.
Multipath Solution (Option B): To enable active routes from both peers and allow for load-balancing (ECMP) across both links, you must enable the multipath feature.
When the multipath statement is configured under the protocols bgp group ext-peers hierarchy, it tells Junos OS to install multiple equal-cost BGP paths into the routing table and subsequent forwarding table.
Since both neighbors belong to the same peer group (ext-peers) and the same AS (65002), configuring multipath at the group level will apply to both sessions, allowing paths from both neighbors to be marked

as "active".
Incorrect Options:
Option A: 172.16.1.1 is the loopback address of router2. The exhibit shows peering is currently done using physical interface addresses (10.10.10.1 and 10.10.20.1), so this address is irrelevant to the current active sessions.
Option C: Configuring the neighbor address alone without the multipath parameter will not change the best-path selection behavior.
Option D: 10.10.20.2 is the local interface IP of router1, not the neighbor's IP. BGP multipath must be configured to point to remote peer paths.
Configuration Example for Junos OS 24.4: To implement this, apply the following command: set protocols bgp group ext-peers multipath

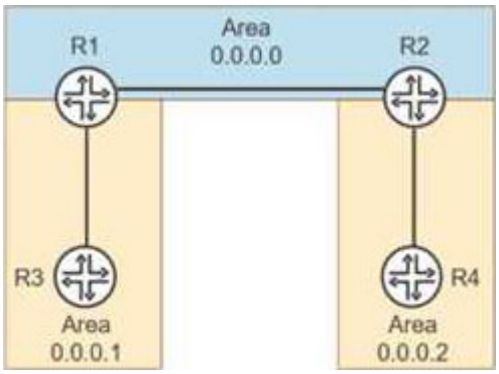
4.Click the Exhibit button.



Referring to the exhibit output, which statement is correct?
A. The device with an IP address of 192.168.100 2 is attached to VNI 5010
B. The output shows the router ID of the device sending the EVPN Type 2 routes
C. The output shows the route distinguisher of the device sending the EVPN Type 2 routes
D. IP addresses 10.1.1.1 and 10.1 2.3 are connected to the same VNI.
Answer: C
Explanation:
The exhibit displays the default-switch.evpn.0 routing table, which is used on Juniper leaf devices to store EVPN Type 2 (MAC/IP) routes.
Route Distinguisher (Option C): In EVPN, the Route Distinguisher (RD) is an 8-byte prefix added to a route to make it unique within the BGP control plane. The RD format in the exhibit is <IP-Address>:<Identifier>.
For example, the prefix 2:192.168.100.1:1::5010::... indicates an EVPN Type 2 route (2:) where

192.168.100.1:1 is the Route Distinguisher.
This RD identifies the specific routing instance on the originating VTEP that advertised the MAC/IP address.
Option A is incorrect: The RD 192.168.100.2:1 does not necessarily mean the host device has that IP; it means the originating switch has that router ID/IP used for its RD.
Option B is incorrect: While the RD often incorporates the router ID, the RD itself is the full string (e.g., 192.168.100.1:1), which is distinct from the raw Router ID used in the BGP summary.
Option D is incorrect: Looking at the entries for 10.1.1.1 and 10.1.2.3, they are associated with different identifiers in the RD strings (5010 and 5020 respectively), which typically map to different VNIs or bridge domains.

5.Click the Exhibit button.



You are troubleshooting suboptimal routing in an enterprise OSPFv2 deployment:
-- R1 is an ABR between Area 0 and Area 1.
-- R2 is an ABR between Area 0 and Area 2.
-- R3 in Area 1 advertises a network that should be reachable by routers in Area 2.
-- The ABRs are connected in the backbone (Area 0).
On router R2, the show ospf database does not display router LSAs or network LSAs from Area 1, but an interarea route to R3's network exists in the routing table. R2 does not have Type-1 or Type-2 LSAs from Area 1 in its OSPF database.
Referring to the exhibit, which statement is correct?
A. R2 is filtering them using the OSPF import policy.
B. Junos disables Type-1 and Type-2 LSAs by default in Interarea LSDBs.
C. R2 is running OSPFv3, which does not use Type-1 or Type-2 LSAs.
D. Type-1 and Type-2 LSAs are summarized into Type-3 LSAs by ABRs and do not cross area boundaries.
Answer: D
Explanation:
The behavior described is the fundamental operational principle of OSPF's hierarchical area structure.
LSA Flooding Scope: In OSPFv2, Type-1 (Router) and Type-2 (Network) Link-State Advertisements (LSAs) are restricted to a local area flooding scope. They contain detailed topology information (such as which routers are connected to which links) that is only relevant for calculating the Shortest Path First (SPF) tree within that specific area.
ABR Functionality: Area Border Routers (ABRs), such as R1 in this scenario, act as the gateway between areas. The ABR does not forward Type-1 or Type-2 LSAs from one area to another. Instead, it

extracts the prefix information (the network address and mask) from those LSAs and generates a Type-3 (Summary) LSA.
Interarea Reachability: This Type-3 LSA is then flooded into the backbone (Area 0) and subsequently into other non-backbone areas, such as Area 2. This explains why R2 sees an "interarea route" to R3’s network in its routing table and has the corresponding Type-3 LSA in its database, but cannot see the original Type-1 or Type-2 LSAs from Area 1.
Why other options are incorrect: Option A is incorrect because this is standard protocol behavior, not the result of a manual filter.
Option B is incorrect as there is no such "disable" feature; the protocol design inherently prevents these LSAs from crossing boundaries.
Option C is incorrect because OSPFv3 still uses Type-1 and Type-2 LSAs (renamed but functionally similar for topology), and the prompt specifically identifies the deployment as OSPFv2.