

Reliable exam torrent & valid VCE PDF materials & dumps torrent files

Input your exam code ...



Our high-quality valid VCE PDF materials & dumps torrent files guarantee you pass exam 100% for sure. Our reliable exam torrent will be the best help for your exams and will give you a new start, a new life.

[All Products](#)

[Contact now](#)



Quality and Value

VCETorrent Practice Exams are written to the highest standards of technical accuracy, using only certified subject matter experts and published authors for development - no all dumps.



Tested and Approved

We are committed to the process of vendor and third party approvals. We believe professionals and executives alike deserve the confidence of quality coverage these authorizations provide.



Easy to Pass

If you prepare for the exams using our VCETorrent testing engine, It is easy to succeed for all certifications in the first attempt. You don't have to deal with all dumps or any free torrent / rapidshare all stuff.



Try Before Buy

VCETorrent offers free demo of each product. You can check out the interface, question quality and usability of our practice exams before you decide to buy.

19728

CUSTOMERS

32727

DOWNLOADS

21728

TEAM
MEMBERS

26394

SHARES

<http://www.vcetorrent.com>

Reliable exam torrent & valid VCE PDF materials & dumps torrent files

Exam : **JN0-364**

Title : Service Provider Routing and Switching, Specialist (JNCIS-SP)

Vendor : Juniper

Version : DEMO

NO.1 Referring to the exhibit. Which prefix in the output shown in the exhibit is an external prefix injected by an OSPF router?

```

user@R2> show ospf route
Topology default Route Table:

```

Prefix	Path Type	Route Type	NH Type	Metric	NextHop Interface	NextHop addr/label
192.168.1.1	Intra	AS BR	IP	1	ge-0/0/3.0	172.26.1.1
192.168.1.3	Intra	Area BR	IP	1	ge-0/0/1.0	172.26.2.2
172.18.1.0/24	Ext2	Network	IP	0	ge-0/0/3.0	172.26.1.1
172.26.1.0/30	Intra	Network	IP	1	ge-0/0/3.0	
172.26.2.0/30	Intra	Network	IP	1	ge-0/0/1.0	
172.26.3.0/30	Intra	Network	IP	100	ge-0/0/2.0	
172.26.4.0/30	Inter	Network	IP	2	ge-0/0/1.0	172.26.2.2
192.168.1.1/32	Ext2	Network	IP	1	ge-0/0/3.0	172.26.1.1
192.168.1.2/32	Intra	Network	IP	0	lo0.0	
192.168.1.3/32	Intra	Network	IP	1	ge-0/0/1.0	172.26.2.2
192.168.1.4/32	Inter	Network	IP	2	ge-0/0/1.0	172.26.2.2

- A. 192.168.1.3
- B. 172.18.1.0/24
- C. 192.108.1.4
- D. 172.26.4.0/30

Answer: B

Explanation:

In the OSPF routing table output, prefixes are marked with different route types. An external prefix injected into OSPF is marked as 'Ext' (External) followed by a number that indicates whether it's an E1 or E2 route. The prefix 172.18.1.0/24 is marked as Ext2, which indicates that it is an external route that has been redistributed into OSPF from another routing protocol or static configuration.

NO.2 Which statement describes integrated routing and bridging (IRB) interfaces?

- A. An IRB interface is an IP gateway for hosts of a bridge domain.
- B. An IRB interface assigns interfaces to VLANs.
- C. An IRB interface enables Layer 2 switching on the router.
- D. An IRB interface defines a bridge domain.

Answer: A

Explanation:

An Integrated Routing and Bridging (IRB) interface is used to provide Layer 3 routing services to hosts within a bridge domain. The IRB acts as a default gateway for hosts in that domain, enabling communication with other networks.

NO.3 The LSP is not establishing correctly.

Referring to the exhibit, what should you do to solve the problem?

```
Exhibit
user@router> show mpls lsp ingress detail
Ingress LSP: 1 sessions
192.168.0.3
  From: 0.0.0.0, State: Dn, ActiveRoute: 0, LSPname: to-R3
  ActivePath: (none)
  LSPTYPE: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Follow destination IGP metric
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  LSP Self-ping Status : Enabled
  Primary                               State: Dn
  Priorities: 7 0
  SmartOptimizeTimer: 180
  Flap Count: 0
  MBB Count: 0
  Will be enqueued for recomputation in 18 second(s).
  1 Mar 9 23:22:22.998 CSFP: could not determine self
user@router> show ted database
TED database: 0 ISIS nodes 0 INET nodes
[edit protocols]
user@router# show
ospf {
  area 0.0.0.0 {
    interface ge-0/0/2.0;
    interface ge-0/0/4.0;
  }
}
rsvp {
  interface all;
}
bgp {
  group Int {
    type internal;
    local-address 192.168.0.1;
    export nhs;
    neighbor 192.168.0.3;
  }
}
mpls {
  label-switched-path to-R3 {
    to 192.168.0.3;
  }
  interface all;
}
```

- A. Enable traffic engineering for the OSPF protocol.
- B. Enable traffic engineering for the IS-IS protocol.
- C. Enable traffic engineering for the BGP protocol.
- D. Enable traffic engineering for the RSVP protocol.

Answer: A

Explanation:

The exhibit shows that the Label Switched Path (LSP) is down. One common reason for this could be

that the IGP is not providing traffic engineering information to the MPLS process. Since the exhibit shows the OSPF configuration, enabling traffic engineering extensions for OSPF would allow OSPF to distribute the labels and traffic engineering information necessary for LSP establishment.

NO.4 You are bringing a new network online with three MX Series devices enabled for STP. No root bridge priority has been configured. Which statement is true in this scenario?

- A.** The device with the lowest MAC address will be elected as the root bridge.
- B.** The device with the highest MAC address will be elected as the root bridge.
- C.** The device with the lowest numerical lo0 IP address will be elected as the root bridge.
- D.** The device with the highest numerical lo0 IP address will be elected as The bridge.

Answer: A

Explanation:

The root bridge in a spanning-tree network is the bridge with the smallest or the lowest bridge ID. In the absence of a manually configured priority, the Spanning Tree Protocol (STP) elects the root bridge based on the lowest bridge ID, which is a combination of the priority and the MAC address. The device with the lowest MAC address will have the lowest bridge ID and thus be elected as the root bridge.

NO.5 What is a key differentiator of generate routes from aggregate routes?

- A.** Generate routes use a forwarding next hop.
- B.** Generate routes have a default next-hop value of reject.
- C.** Generate routes have a default preference value of 210.
- D.** Generate routes cannot be used as a gateway of last resort.

Answer: A

Explanation:

Generated routes are a type of route that can be created to summarize and generate more specific routes within the routing table. Unlike aggregate routes, which summarize existing routes and inherit a next-hop, generated routes do not necessarily have to match an existing route and will have a next-hop of reject by default unless specified otherwise.

NO.6 Which statement is correct about the FE80::/10 prefix?

- A.** This prefix range is used for the link local address.
- B.** This prefix range is used on the loopback interface.
- C.** This prefix range is reserved for multicast applications
- D.** This prefix range is not reserved.

Answer: A

Explanation:

The FE80::/10 prefix is reserved for IPv6 link-local addresses. These addresses are auto-configured on all IPv6-enabled interfaces and can be used for communication within the local link (subnet) only.

NO.7 You are deploying link aggregation groups. By default, what are two considerations in this scenario? (Choose two.)

- A.** There should only be four member links per LAG.
- B.** All the ports must have the same speed.

- C. Member links are required to be contiguous ports.
- D. Member links can reside on different members within an MC-LAG.

Answer: BD

Explanation:

When deploying Link Aggregation Groups (LAGs), two key default considerations are: All member ports must have the same speed. This ensures consistent performance and proper aggregation since mismatched speeds can cause issues in traffic distribution and link stability. Member links can reside on different members within a Multi-Chassis LAG (MC-LAG). MC-LAG allows member links of the aggregation group to be spread across different physical chassis, providing redundancy and load sharing across devices.

NO.8 Referring to the exhibit. You are asked to assign interface xe-1/0/5 to a virtual switch. What must be accomplished to complete the configuration?

```
[edit]
user@switch# show routing-instances
sw-1 {
    instance-type virtual-switch;
    bridge-domains {
        vlan_1 {
            vlan-id 1;
        }
        vlan_2 {
            vlan-id 2;
        }
    }
}

[edit]
user@switch# show interfaces xe-1/0/5
unit 0 {
    family bridge {
        interface-mode access;
        vlan-id 2;
    }
}
```

- A. Interface xe-1/0/5 must be added to routing-instance sw-1 vlan_2.
- B. Interface xe-1/0/5 must be a trunk port.
- C. Interface xe-1/0/5 must be added to routing-instance sw-1.

D. An IRB interface must be configured to routing-instance sw-1 vlan_2.

Answer: A

Explanation:

The exhibit shows the configuration of a virtual switch called sw-1 with two VLANs defined within it: vlan_1 with VLAN ID 1 and vlan_2 with VLAN ID 2. To add interface xe-1/0/5 to the virtual switch, the interface must be associated with one of these VLANs. Since the interface is already configured with vlan-id 2, it implies that it is intended to be part of vlan_2 within the virtual switch sw-1. Therefore, the correct answer is to add interface xe-1/0/5 to the vlan_2 bridge domain under the sw-1 routing instance.

NO.9 Referring to the exhibit, which statement is true about VRRP?

```

user@R1> show vrrp summary
Interface      State      Group      VR state  VR Mode   Type Address
ge-0/0/4.0    up         10         master   Active   lcl  172.25.100.2
                                       vip  172.25.100.1

user@R2> show vrrp summary
Interface      State      Group      VR state  VR Mode   Type Address
ge-0/0/4.0    up         10         master   Active   lcl  172.25.100.3
                                       vip  172.25.100.1

```

- A.** VRRP communication between the two devices is not functioning correctly.
- B.** Both routers are in the same state because they have the same VRRP priority.
- C.** VRRP is functioning normally in active/active mode.
- D.** The routers should use different virtual IP addresses for VRRP to function correctly.

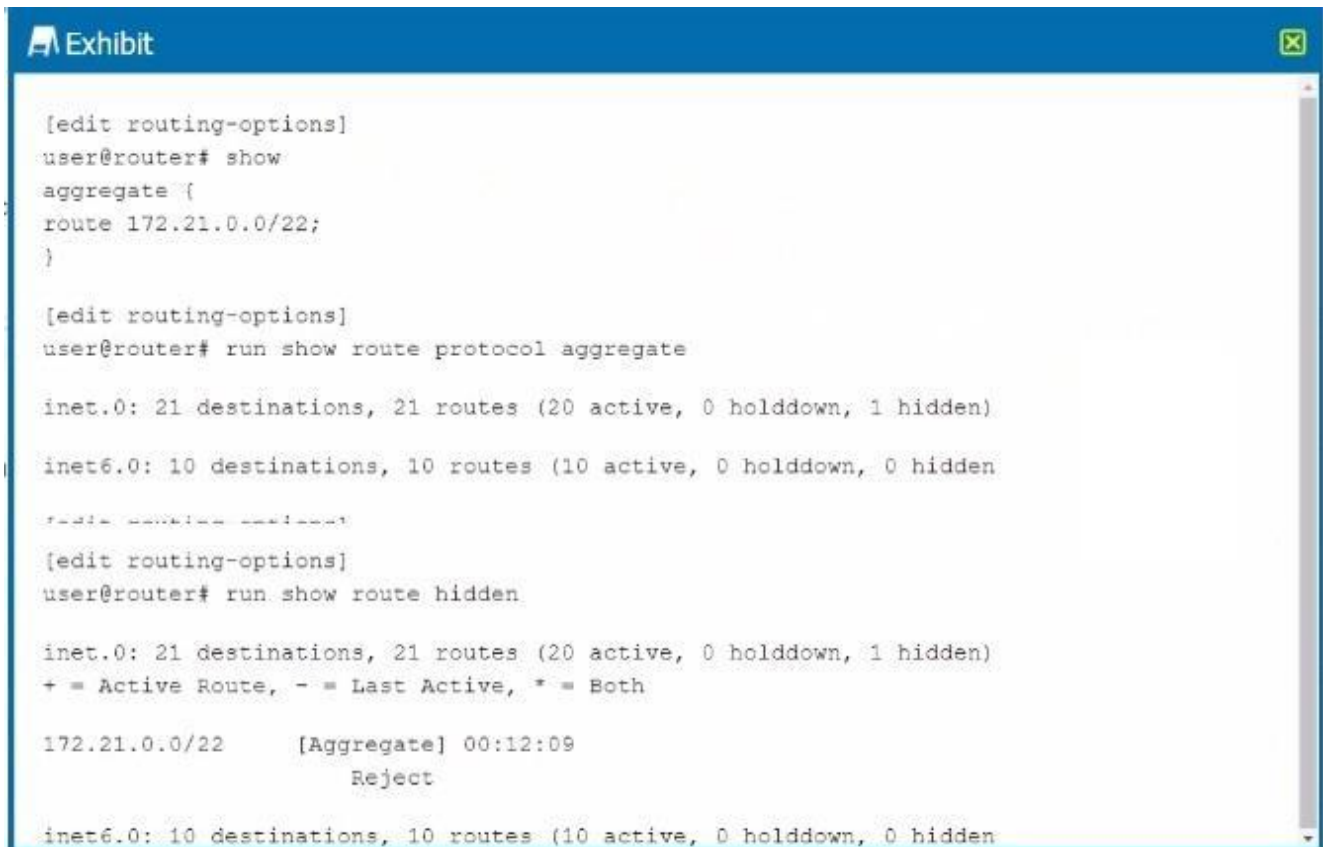
Answer: C

Explanation:

The exhibit shows the show vrrp summary output from two routers, with both routers listing the state of their VRRP group as master. VRRP allows for multiple routers to share a virtual IP address and a virtual MAC address, but usually, only one router is the master at any given time while others are in backup state. However, in some configurations, VRRP can operate in an active/active mode where multiple routers handle traffic simultaneously for load sharing. Since both routers show the state as master and there is no evidence of misconfiguration in the virtual IP addresses or priority, it indicates that VRRP is likely functioning in an active/active setup.

NO.10 Referring to the exhibit, you have configured an aggregate route that represents the 172.21.0.0/24, 172.21.1.0/24, and 172.21.2.0/24 networks. However, when you view the routing table, your new route is hidden.

Which action would you perform to determine the problem?



```
[edit routing-options]
user@router# show
aggregate {
  route 172.21.0.0/22;
}

[edit routing-options]
user@router# run show route protocol aggregate

inet.0: 21 destinations, 21 routes (20 active, 0 holddown, 1 hidden)
inet6.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
redix.unused.routes: 0

[edit routing-options]
user@router# run show route hidden

inet.0: 21 destinations, 21 routes (20 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.21.0.0/22    [Aggregate] 00:12:09
                Reject

inet6.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
```

- A. Verify that you have active contributing routes on the device.
- B. Verify that you have configured a policy on the device to accept aggregate routes.
- C. Verify that you have defined a metric value for the aggregate route.
- D. Verify that you have set the preference to a lower default value.

Answer: D

Explanation:

The exhibit shows an aggregate route configuration for the network 172.21.0.0/22, which would summarize the specific networks 172.21.0.0/24, 172.21.1.0/24, and 172.21.2.0/24. For an aggregate route to be active, it must have contributing routes in the routing table. If the route is hidden, it usually means there are no contributing routes that are active or the policy applied to the aggregate does not match any of the specific routes. Therefore, the first step in troubleshooting would be to verify that there are indeed active contributing routes for the aggregate to be valid.

NO.11 What are two bridging concepts that are used to maintain an Ethernet switching table? (Choose two.)

- A. learning
- B. exporting
- C. aging
- D. timing

Answer: AC

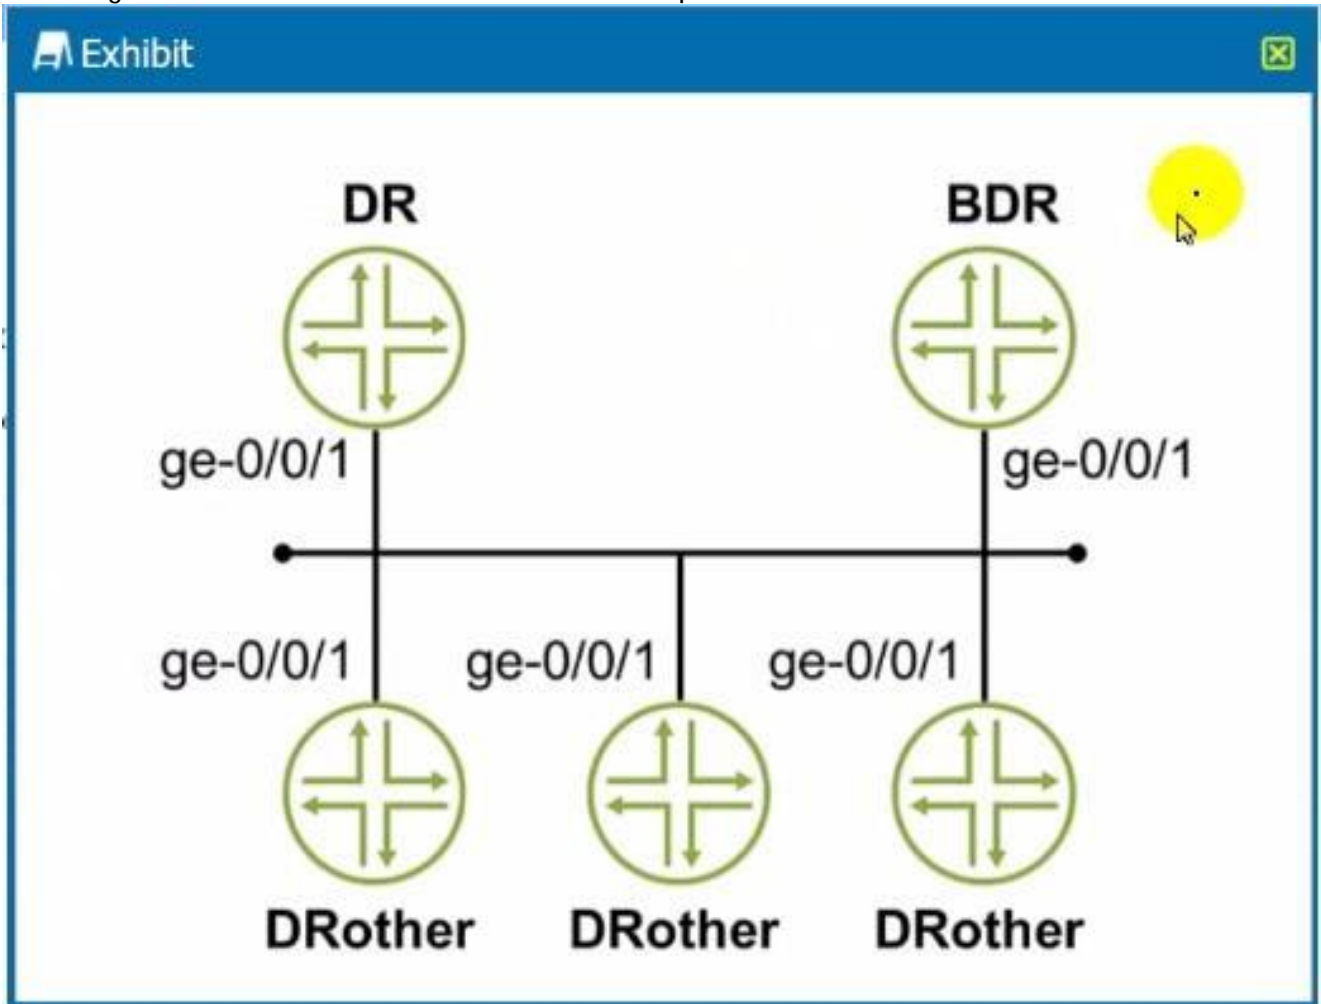
Explanation:

In Ethernet switching, learning and aging are two fundamental concepts for maintaining a dynamic Ethernet switching table (also known as a MAC address table). Learning is the process by which switches listen for frames and learn the MAC addresses by associating them with the incoming port.

Aging is the mechanism that ensures the switch updates its table by removing MAC addresses that have not been seen for a certain amount of time (aging time).

NO.12 You are asked to configure the OSPF environment to prevent the DRother routers from participating in DR/BDR election.

Referring to the exhibit, which command will accomplish this task?



- A. set protocols ospf area 0.0.0.0 interface ge-0/0/1 priority 255
- B. set protocols ospf area 0.0.0.0 interface ge-0/0/1 priority 0
- C. set protocols ospf area 0.0.0.0 interface ge-0/0/1 interface-type nbma
- D. set protocols ospf area 0.0.0.0 interface ge-0/0/1 interface-type p2p

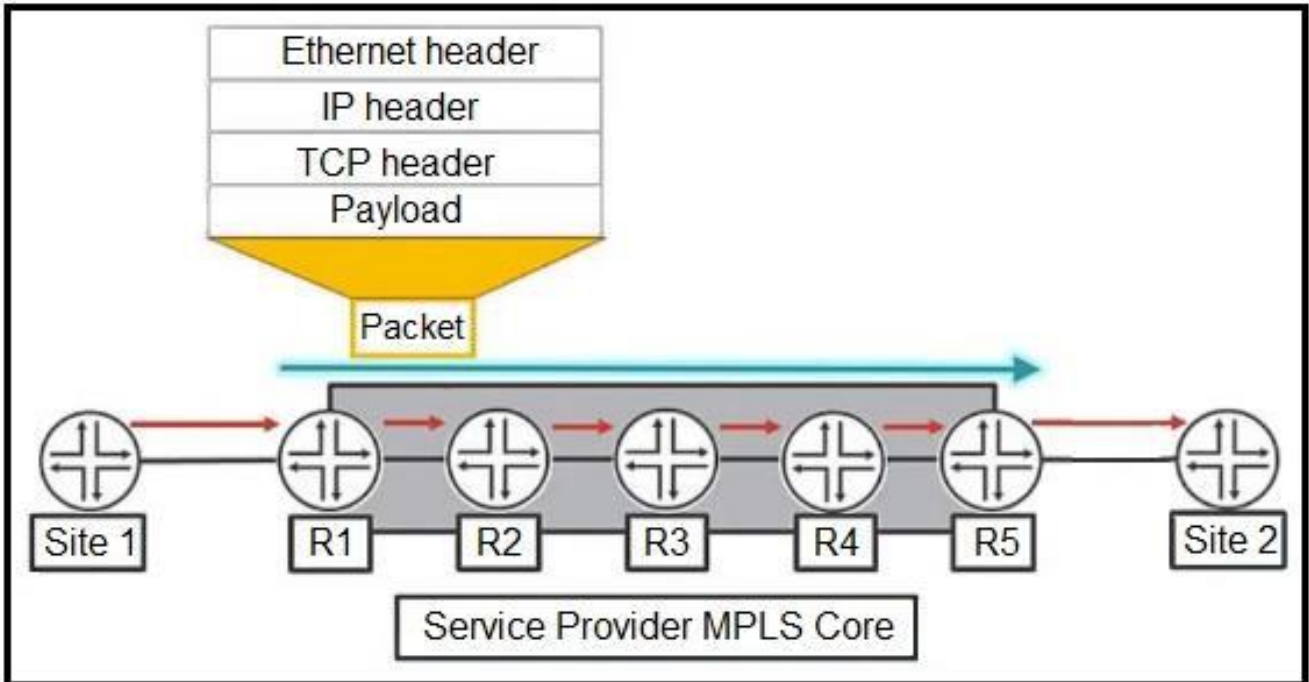
Answer: B

Explanation:

To prevent a router from participating in the OSPF designated router (DR) and backup designated router (BDR) election process, you set the OSPF priority to 0 on the interfaces of that router. This tells OSPF that the router on those interfaces is not to be considered for election as either DR or BDR, effectively making them DRother routers. The exhibit indicates that this is the desired configuration, and therefore the correct command to issue on the Juniper device would be to set the priority to 0 on the relevant interfaces.

NO.13 Referring to the exhibit. Which two statements are correct about the actions taken as the packet traverses the service provider MPLS network from Site 1 to Site 2 as shown in the exhibit?

(Choose two.)



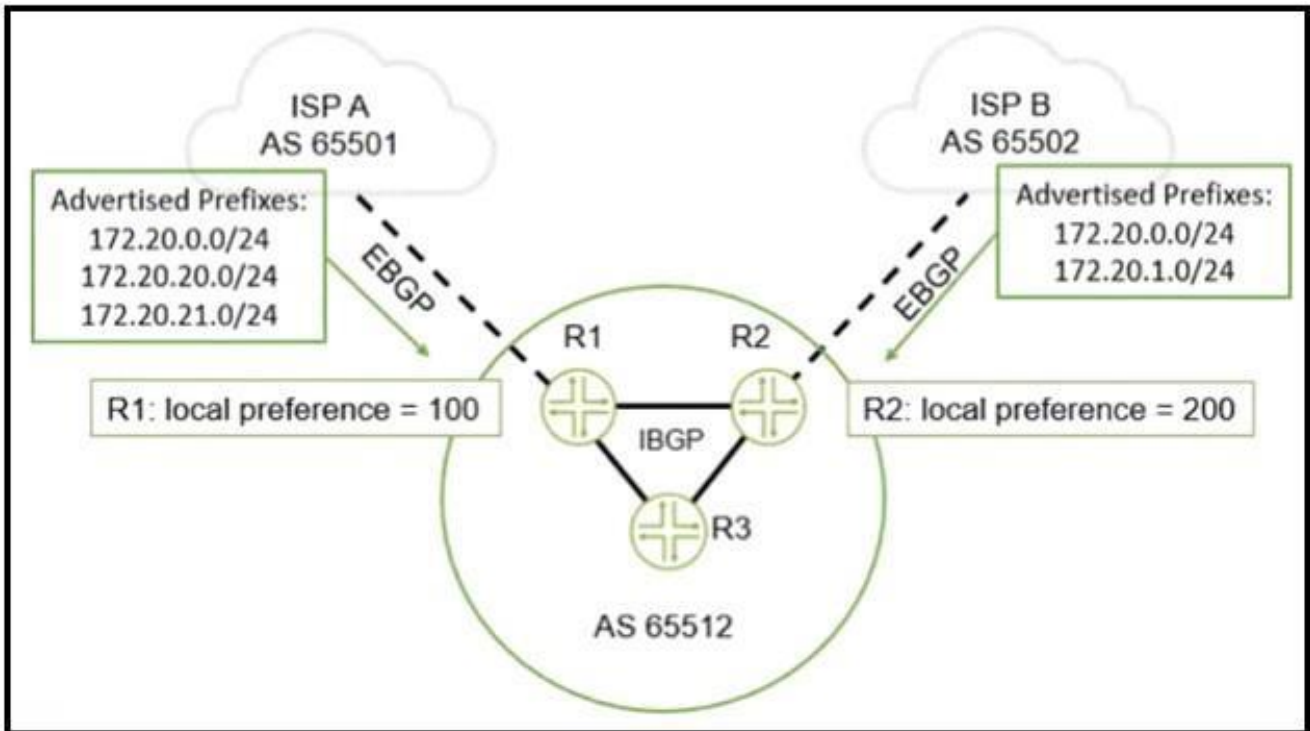
- A. R2 will perform a lookup using the mpls.0 table.
- B. R1 will perform a lookup using the inet.3 table.
- C. R1 will perform a lookup using the mpls.0 table.
- D. R2 will perform a lookup using the inet.3 table.

Answer: AC

Explanation:

In MPLS (Multiprotocol Label Switching) networks, routers use label switching to forward packets. The first router at the edge of the MPLS network (R1) will perform a lookup in the mpls.0 table to determine the label to attach to the packet as it enters the MPLS network. This label informs the next routers in the MPLS network (like R2) on how to forward the packet. Internal MPLS routers, like R2, also perform lookups in their mpls.0 table to determine how to switch the packet toward its destination (label swapping). The inet.3 table is used for resolving next-hop information for labeled routes, but it's the mpls.0 table that is used for label switching decisions.

NO.14 Referring to the exhibit, which two statements are correct? (Choose two.)



- A. Devices in AS 65512 will prefer ISP A for traffic destined to the 172.20.21.0/24 network.
- B. Devices In AS 65512 will prefer ISP A for traffic destined to the 172.20.0.0/24 network.
- C. Devices in AS 65512 will prefer ISP B for traffic destined to the 172.20.21.0/24 network.
- D. Devices In AS 65512 will prefer ISP B for traffic destined to the 172.20.0.0/24 network.

Answer: AB

Explanation:

Local preference is a BGP attribute that dictates which path is preferred when multiple paths to the same destination are available. A higher local preference is preferred over a lower one. Since ISP A (AS 65501) is advertising a local preference of 100 and ISP B (AS 65502) is advertising a local preference of 200 to AS 65512, traffic from devices in AS 65512 will prefer the path via ISP B for all networks that ISP B advertises. However, since ISP A advertises the 172.20.21.0/24 network and ISP B does not, traffic for 172.20.21.0/24 will go through ISP A. For the 172.20.0.0/24 network, which both ISPs advertise, devices in AS 65512 will prefer the path via ISP B due to the higher local preference.

NO.15 Referring to the exhibit. You have an established LSP between your R1 and R5 devices using the configuration shown in the exhibit. You are asked to ensure that MPLS labels are used to forward traffic by all devices within the LSP.

Which action will accomplish this behavior?

```

user@R1> show configuration protocols mpls
label-switched-path R1_TO_R5 {
    to 192.168.1.5;
    no-cspf;
}
interface ge-0/0/0.0;
interface ge-0/0/1.0;

```

- A. Configure the ultimate-hop-popping statement under the R1_TO_R5 label switched path on R1.
- B. Configure the explicit-null statement under the protocol mpls hierarchy on R1.
- C. Delete the no-espaf statement under the R1_TO_R5 label switched path on R1.
- D. Configure the install statement under the R1_TO_R5 label switched path on R1.

Answer: B

Explanation:

The "ultimate-hop-popping" term refers to the action taken by the penultimate router in an LSP to remove the MPLS label before delivering the packet to the ultimate router, which is not desired here. Configuring the "explicit-null" statement causes the penultimate router to replace the top label with a label that has a value of 0, which instructs the ultimate router to perform a label lookup and preserve the label switching for the entire LSP.

NO.16 Referring to the exhibit. Which two statements are correct about the information shown in the exhibit? (Choose two.)

```

user@switch> show spanning-tree bridge
STP bridge parameters
Context ID                : 0
Enabled protocol          : RSTP
  Root ID                  : 8192.59.c5:8d:ae:db:41
  Hello time               : 10 seconds
  Maximum age              : 40 seconds
  Forward delay            : 30 seconds
  Message age              : 0
  Number of topology changes : 6
  Time since last topology change : 781 seconds
  Topology change initiator : ge-0/0/14.0
  Topology change last recvd. from : 2c:6b:f5:31:06:0b
Local parameters
  Bridge ID                : 8192.50:c5:8d:ae:db:41
  Extended system ID       : 0
  Internal instance ID     : 0

```

- A. The root bridge is reachable using the ge-0/0/14 interface.
- B. This switch is the root bridge for this spanning tree topology.
- C. This switch has a bridge priority of 8k.

D. The root bridge's priority is 4k.

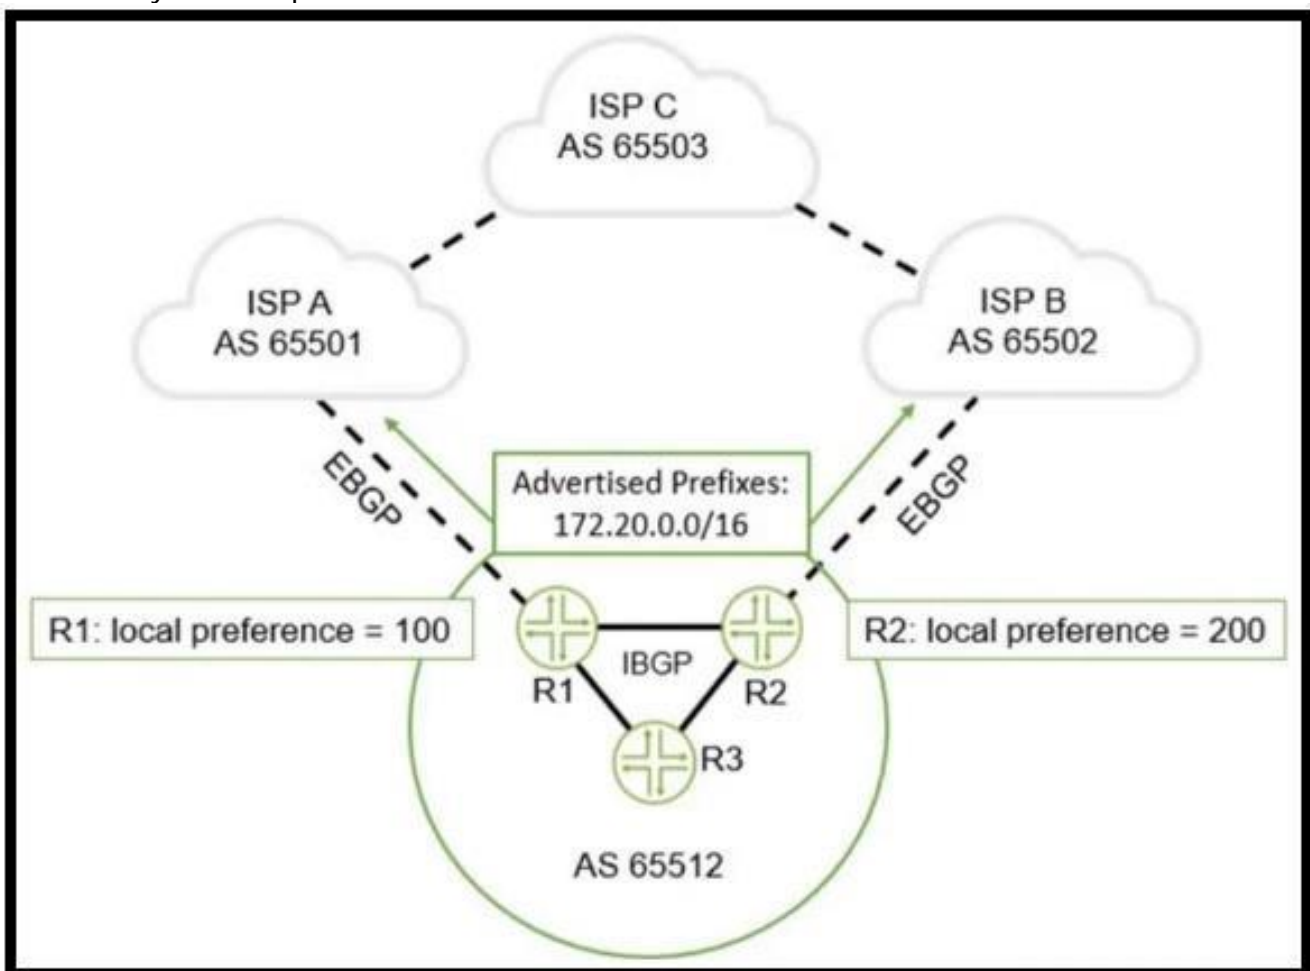
Answer: AC

Explanation:

The exhibit shows the output of the command show spanning-tree bridge, which provides information about the spanning tree status of the switch. From the output, we can see that the switch has a bridge ID different from the root ID, which implies that this switch is not the root bridge. The "Topology change initiator" field shows ge-0/0/14, which indicates that the last topology change occurred on this interface, and this is also the interface used to reach the root bridge.

NO.17 Referring to the exhibit. You are advertising a summary route that represents your local network (172.20.0.0/16) to both ISP A and ISP B. You want to influence all traffic sent to you from ISP C to go through R2.

How would you accomplish this task?



- A. On R1, prepend your AS number three times on the 172.20.0.0/16 route when advertising it to ISP 1.
- B. On R1, change the local preference value to 250.
- C. On R2, prepend your AS number three times on the 172.20.0.0/16 route when advertising it to ISP 2.
- D. On R2, change the local preference value to 50.

Answer: C

Explanation:

To influence ISP C to prefer the path through R2 for the 172.20.0.0/16 network, R2 should make its path less preferable to ISP B by prepending its AS number multiple times. This makes the path through R2 appear longer and less desirable to ISP C, which would receive these updates from ISP B. AS path length is a well-known attribute in BGP and is used to determine the best path; paths with a longer AS path are considered less preferable.

NO.18 Which two statements are correct about the community BGP attribute on a Junos device? (Choose two.)

- A.** The community attribute is a mandatory BGP attribute.
- B.** If the community attribute is present, it is ignored and deleted in the BGP updates.
- C.** If the community attribute is present, it should be passed unchanged in the BGP updates.
- D.** The community attribute is an optional BGP attribute.

Answer: CD

Explanation:

The community attribute in BGP is an optional transitive attribute. It is not mandatory for BGP operations, but when it is present, it should be passed unchanged to other BGP peers unless explicitly modified by a routing policy. The community attribute is used to group destinations in a certain way, often to apply routing policies to those destinations collectively.