Juniper

Exam Questions JN0-351

Enterprise Routing and Switching - Specialist (JNCIS-ENT)

NEW QUESTION 1
You are attempting to configure the initial two aggregated Ethernet interfaces on a router but there are no aggregated Ethernet interfaces available.
In this scenario, which configuration will enable these interfaces on this router?
A)

```
user@router# show chassis
aggregated-devices {
    ethernet {
        lacp {
            system-priority 10;
    }
}
```

B)

```
user@router# show chassis
aggregated-devices {
    ethernet {
        device-count 10;
    }
}
```

C)

```
user@router# show chassis
maximum-ecmp 16;
aggregated-devices {
    ethernet {
        device-count 1;
    }
}
```

D)

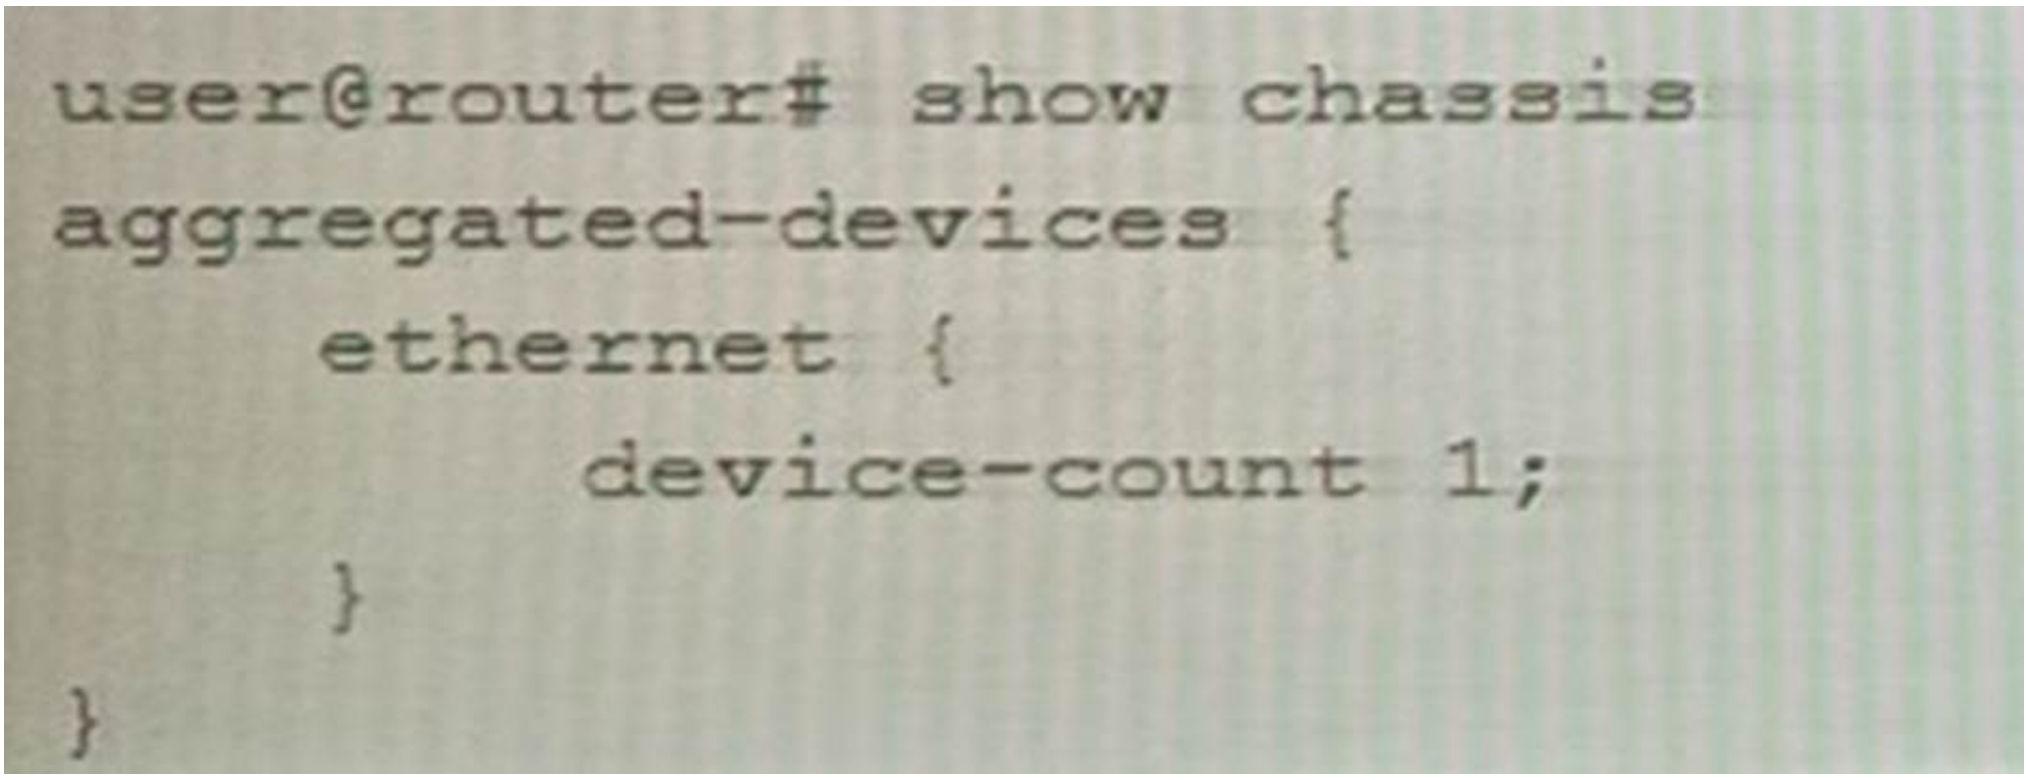

A. Option A
B. Option B
C. Option C
D. Option D

Answer: C

Explanation:
The correct answer to your question is C. Option C. Here is why:
? Option C shows the configuration of the chassis statement, which defines the properties of the router chassis, such as the number of aggregated Ethernet interfaces, the number of FPCs, and the number of PICs1.
? To enable aggregated Ethernet interfaces on a router, you need to specify the aggregated-devices statement under the chassis statement and set the ethernet parameter to the desired number of interfaces2. For example, to enable two aggregated Ethernet interfaces, you can use the following configuration: chassis { aggregated-devices { ethernet { device-count 2; } } }
? Option C shows this configuration with the device-count set to 2, which will enable two aggregated Ethernet interfaces on the router. The other options do not show this configuration and will not enable any aggregated Ethernet interfaces on the router.
? Therefore, option C is the correct answer to your question.

NEW QUESTION 2
Exhibit.

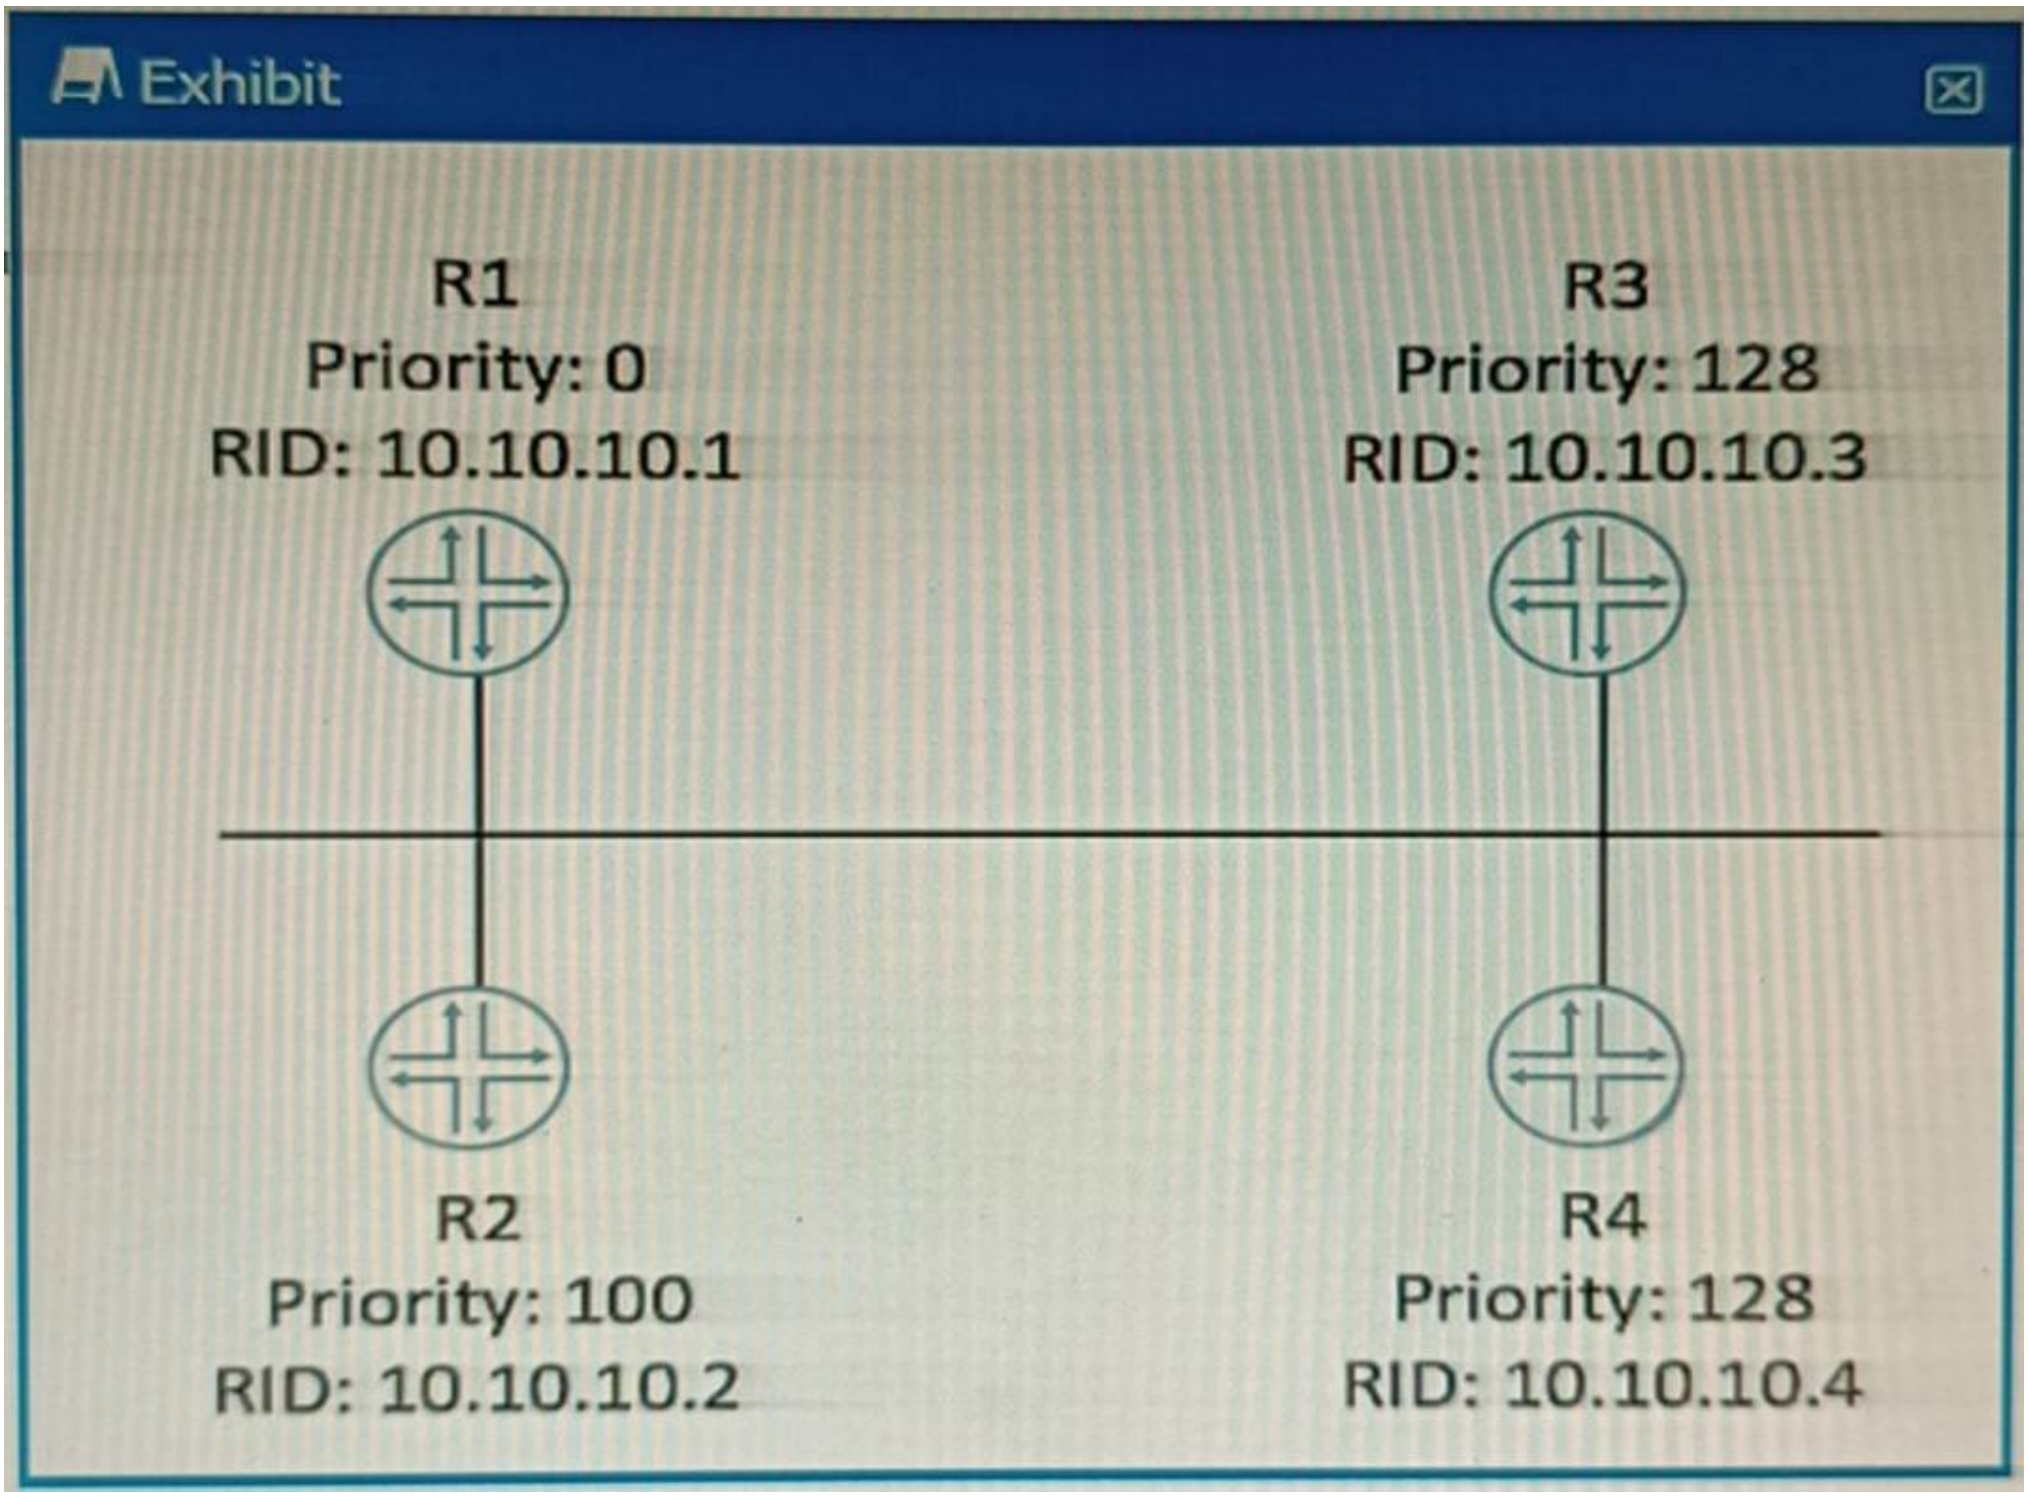

Which router will become the OSPF BDR if all routers are powered on at the same time?

A. R4
B. R1
C. R3
D. R2

Answer: A

Explanation:
OSPF DR/BDR election is a process that occurs on multi-access data links. It is intended to select two OSPF nodes: one to be acting as the Designated Router (DR), and another to be acting as the Backup Designated Router (BDR).The DR and BDR are responsible for generating network LSAs for the multi-access network and synchronizing the LSDB with other routers on the same network1.
The DR/BDR election is based on two criteria: the OSPF priority and the router ID. The OSPF priority is a value between 0 and 255 that can be configured on each interface participating in OSPF. The default priority is 1. A priority of 0 means that the router will not participate in the election and will never become a DR or BDR. The router with the highest priority will become the DR, and the router with the second highest priority will become the BDR. If there is a tie in priority, then the router ID is used as a tie-breaker. The router ID is a 32-bit number that uniquely identifies each router in an OSPF domain.It can be manually configured or automatically derived from the highest IP address on a loopback interface or any active interface2.
In this scenario, all routers have the same priority of 1, so the router ID will determine the outcome of the election. The router IDs are shown in the exhibit as RID values. The highest
RID belongs to R4 (10.10.10.4), so R4 will become the DR. The second highest RID belongs to R3 (10.10.10.3), so R3 will become the BDR.
References:
1:OSPF DR/BDR Election: Process, Configuration, and Tuning2:OSPF Designated Router (DR) and Backup Designated Router (BDR)

NEW QUESTION 3
Which two statements about redundant trunk groups on EX Series switches are correct? (Choose two.)

A. Redundant trunk groups use spanning tree to provide loop-free redundant uplinks.
B. Redundant trunk groups load balance traffic across two designated uplink interfaces.
C. Layer 2 control traffic is permitted on the secondary link.
D. If the active link fails, then the secondary link automatically takes over.

Answer: CD

Explanation:
? C is correct because Layer 2 control traffic is permitted on the secondary link of a redundant trunk group (RTG) on EX Series switches. Layer 2 control traffic includes protocols such as LLDP, LACP, and STP, which are used to exchange information and coordinate actions between switches1. According to the Juniper Networks documentation2, Layer 2 control traffic is allowed to pass through both the active and the secondary links of an RTG, but data traffic is only forwarded through the active link. This allows the switches to maintain their Layer 2 adjacencies and monitor the link status on both links.

? D is correct because if the active link fails, then the secondary link automatically
takes over in an RTG on EX Series switches. An RTG consists of two trunk links: an active or primary link, and a secondary or backup link2. The active link is used to forward data traffic, while the secondary link is in standby mode. If the active link fails or becomes unavailable, the secondary link immediately transitions to a forwarding state and takes over the data traffic without waiting for normal STP convergence2. This provides fast recovery and redundancy for the network.

NEW QUESTION 4
Exhibit

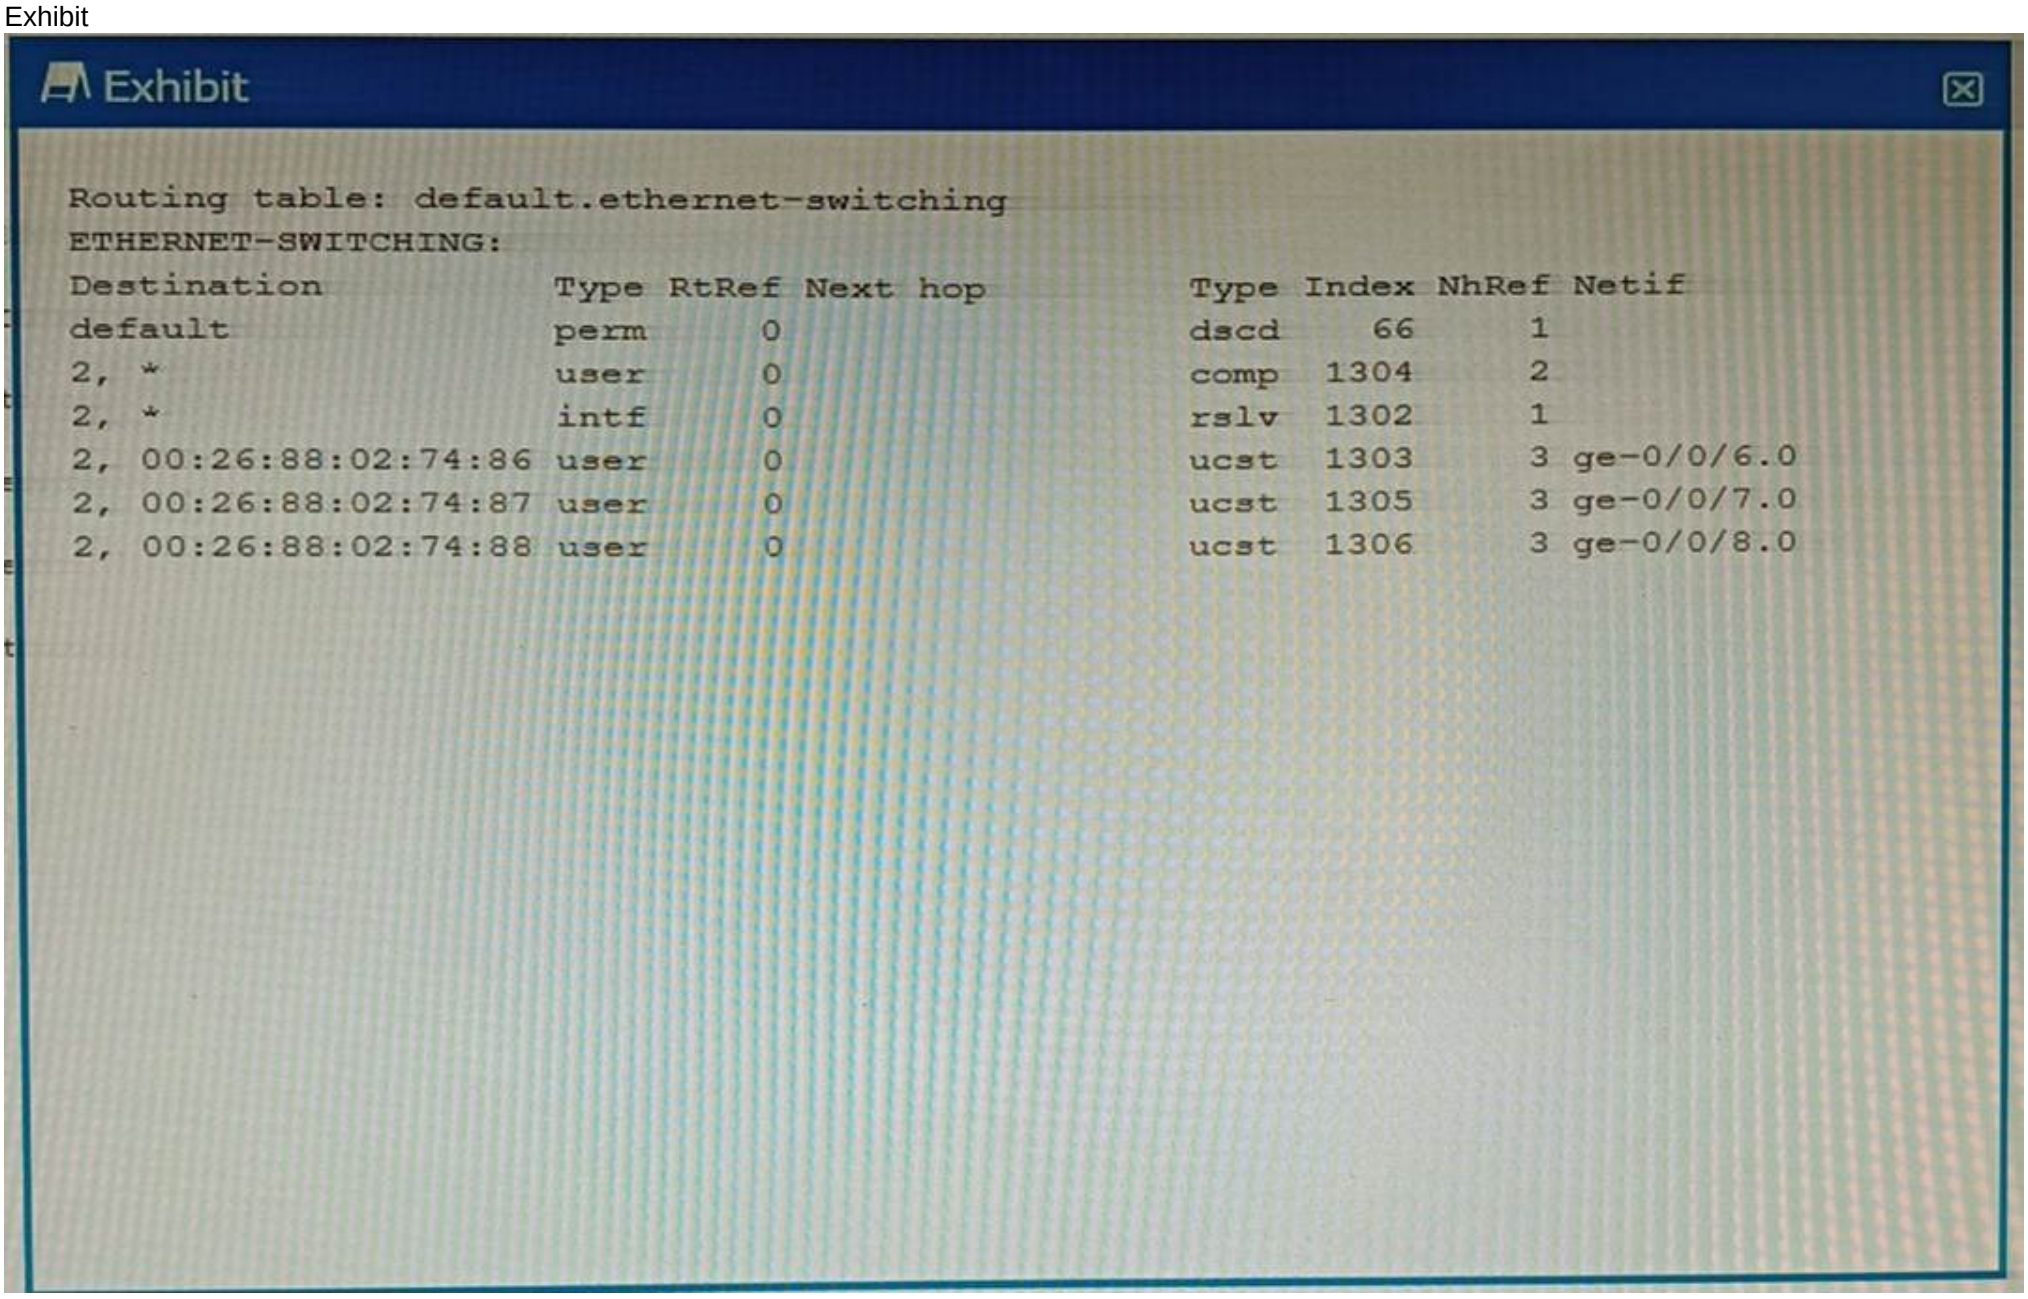

Which command displays the output shown in the exhibit?

A. show route forwarding-table
B. show ethernet-switching table
C. show ethernet—switching table extensive
D. show route forwarding—table family ethernet-switching

Answer: B

Explanation:
? The output shown in the exhibit is a brief display of the Ethernet switching table, which shows the learned Layer 2 MAC addresses for each VLAN and interface1.
? The command show ethernet-switching table displays the Ethernet switching table with brief information, such as the destination MAC address, the VLAN name, the forwarding state, and the interface name1.
? The command show route forwarding-table displays the routing table information for each protocol family, such as inet, inet6, mpls, iso, and so on2. It does not show the Ethernet switching table or the MAC addresses.
? The command show ethernet-switching table extensive displays the Ethernet switching table with extensive information, such as the destination MAC address, the VLAN name, the forwarding state, the interface name, the VLAN index, and the tag type1. It shows more details than the brief output shown in the exhibit.
? The command show route forwarding-table family ethernet-switching displays the routing table information for the ethernet-switching protocol family, whichshows the destination MAC address, the next-hop MAC address, and the interface name3. It does not show the VLAN name or the forwarding state.

NEW QUESTION 5
You are receiving multiple BGP routes from an upstream neighbor and only want to advertise a single summarized prefix to your internal OSPF neighbors. This route should only be advertised when you are receiving these BGP routes from this neighbor.
In this scenario, which type of route should you create?

A. aggregate route
B. static route using the resolve feature
C. generate route
D. static route using qualified next hops

Answer: A

Explanation:
In this scenario, you should create an aggregate route1. Aggregate routes are used for advertising summarized network prefixes1. They help minimize the number of routing tables in an IP network by consolidating selected multiple routes into a single route advertisement1. This approach is in contrast to non-aggregation routing, in which every routing table contains a unique entry for each route1.
Therefore, option A is correct. Options B, C, and D are not correct because:

? Static route using the resolve feature: This type of route uses the resolve feature to install a static route in the routing table only if a specific condition is met1. However, it does not provide the capability to summarize multiple routes into a single prefix.
? Generate route: This type of route generates a route that is always present in the routing table and can be used to summarize routes. However, it does not have the capability to only advertise the route when specific BGP routes are being received from a neighbor1.
? Static route using qualified next hops: This type of route allows for the specification of multiple next-hop addresses for a static route1. However, it does not provide the capability to summarize multiple routes into a single prefix.

NEW QUESTION 6
Exhibit.

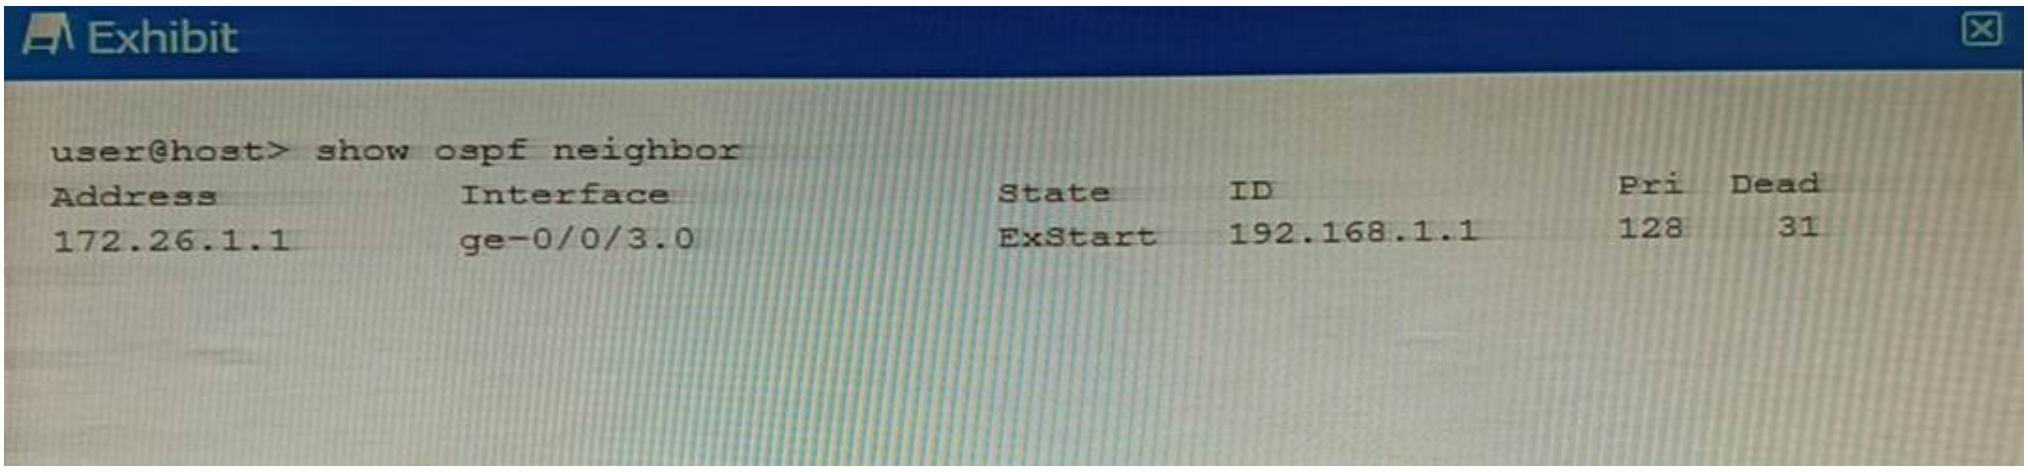

Why is this OSPF adjacency remaining in this state?

A. A subnet mask mismatch exists between the OSPF neighbors.
B. An MTU mismatch exists between the OSPF neighbors.
C. A hello interval mismatch exists between the OSPF neighbors.
D. An area ID mismatch exists between the OSPF neighbors

Answer: B

Explanation:
? The exhibit shows the output of the command show ospf neighbor, which displays information about the OSPF neighbors on a router1.
? The output shows that the OSPF neighbor with the address 172.26.1.1 and the interface ge-0/0/3.0 is in the Exstart state1.
? The Exstart state is the fourth state in the OSPF neighbor formation process, after Down, Init, and 2-Way states2. In this state, the OSPF neighbors establish a master-slave relationship and exchange database description (DBD) packets, which contain summaries of their link-state databases2.
? The most common reason for OSPF neighbors to be stuck in the Exstart state is an MTU mismatch between the interfaces3. MTU stands for maximum transmission unit, which is the largest size of a packet that can be transmitted on a network segment4. If the MTU values of two OSPF neighbors are different, theymay not be able to exchange DBD packets successfully, as some packets may be dropped or fragmented due to their size exceeding the MTU limit3.
? To solve this problem, you need to ensure that the MTU values of both OSPF neighbors are the same or compatible. You can use the command show interfaces to display the MTU value of an interface5. You can also use the command ping with the do-not-fragment option to test the MTU size between two routers. You can change the MTU value of an interface by using the command set interfaces interface-name mtu mtu-value in configuration mode5.

NEW QUESTION 7
You have DHCP snooping enabled but no entries are automatically created in the snooping database for an interface on your EX Series switch. What are two reasons for the problem? (Choose two.)

A. The device that is connected to the interface has performed a DHCPRELEASE.
B. MAC limiting is enabled on the interface.
C. The device that is connected to the interface has a static IP address.
D. Dynamic ARP inspection is enabled on the interface.

Answer: BC

Explanation:
The DHCP snooping feature in Juniper Networks?? EX Series switches works by building a binding database that maps the IP address, MAC address, lease time, binding type, VLAN number, and interface information1. This database is used to filter and validate DHCP messages from untrusted sources1.
However, there are certain conditions that could prevent entries from being automatically created in the snooping database for an interface:
? MAC limiting: If MAC limiting is enabled on the interface, it could potentially
interfere with the operation of DHCP snooping. MAC limiting restricts the number of MAC addresses that can be learned on a physical interface to prevent MAC flooding attacks1. This could inadvertently limit the number of DHCP clients that can be learned on an interface, thus preventing new entries from being added to the DHCP snooping database.
? Static IP address: If the device connected to the interface is configured with a
static IP address, it will not go through the DHCP process and therefore will not have an entry in the DHCP snooping database1. The DHCP snooping feature relies on monitoring DHCP messages to build its database1, so devices with static IP addresses that do not send DHCP messages will not have their information added.
Therefore, options B and C are correct. Options A and D are not correct because performing a DHCPRELEASE would simply remove an existing entry from the database1, and Dynamic ARP inspection (DAI) uses the information stored in the DHCP snooping binding database but does not prevent entries from being created1.

NEW QUESTION 8
You are troubleshooting a BGP routing issue between your network and a customer router and are reviewing the BGP routing policies. Which two statements are correct in this scenario? (Choose two.)

A. Export policies are applied to routes in the RIB-In table.
B. Import policies are applied to routes in the RIB-Local table.
C. Import policies are applied after the RIB-In table.
D. Export policies are applied after the RIB-Local table.

Answer: CD

Explanation:
In BGP, routing policies are used to control the flow of routing information between BGP peers1.
Option C suggests that import policies are applied after the RIB-In table. This is correct because import policies in BGP are applied to routes that are received from a BGP peer, before they are installed in the local BGP Routing Information Base (RIB-In)1. The RIB-In is a database that stores all the routes that are received from all peers1.
Option D suggests that export policies are applied after the RIB-Local table. This is correct because export policies in BGP are applied to routes that are being advertised to a BGP peer, after they have been selected from the local BGP Routing Information Base (RIB- Local)1. The RIB-Local is a database that stores all the routes that the local router is using1.
Therefore, options C and D are correct.

NEW QUESTION 9
Two routers share the same highest priority and start time.

A. In this situation, what is evaluated next when determining the designated router? The router with the lowest router ID become the DR.
B. The router with the highest router ID becomes the DR
C. The routers perform another DR election.
D. The router with the highest MAC address become the DR

Answer: B

Explanation:
? According to the OSPF protocol, the designated router (DR) is the router that acts as the focal point for exchanging routing information on a multi-access network segment, such as a LAN1. The DR election process is based on the following criteria, in order of precedence1:
? In your scenario, two routers share the same highest priority and start time. This means that they have equal chances of becoming the DR based on the first and third criteria. Therefore, the second criterion will be used to break the tie, which is the router ID. The router with the highest router ID will become the DR, and the other router will become the backup designated router (BDR), which is ready to take over the role of DR if it fails1.

NEW QUESTION 10
Which two types of tunnels are able to be created on all Junos devices? (Choose two.)

A. STP
B. GRE
C. IP-IP
D. IPsec

Answer: BD

Explanation:
Junos devices support various types of tunnels for different purposes12.
? Option B is correct. Generic Routing Encapsulation (GRE) is a tunneling protocol that can encapsulate a wide variety of network layer protocols inside virtual point- to-point links over an Internet Protocol network1. Junos devices support GRE tunnels1.
? Option D is correct. IPsec (Internet Protocol Security) is a protocol suite for securing Internet Protocol (IP) communications by authenticating and encrypting each IP packet of a communication session1. Junos devices support IPsec tunnels1.
? Option A is incorrect. Spanning Tree Protocol (STP) is not a type of tunnel. It??s a network protocol designed to prevent loops in a bridged Ethernet local area network2.
? Option C is incorrect. While Junos devices do support IP-IP (also known as IP tunneling), it??s not supported on all Junos devices1.

NEW QUESTION 10
In RSTP, which three port roles are associated with the discarding state? (Choose three.)

A. root
B. backup
C. alternate
D. disabled
E. designated

Answer: BCD

Explanation:
In Rapid Spanning Tree Protocol (RSTP), there are several port roles that determine the behavior of the port in the spanning tree123. The roles include root, designated, alternate, backup, and disabled123.
The discarding state is associated with the backup, alternate, and disabled roles123. In a stable topology with consistent port roles throughout the network, RSTP ensures that every root port and designated port immediately transition to the forwarding state while all alternate and backup ports are always in the discarding state2. Disabled ports are also in the discarding state3.
Therefore, options B, C, and D are correct.

NEW QUESTION 11
You are asked to connect an IP phone and a user computer using the same interface on an EX Series switch. The traffic from the computer does not use a VLAN tag, whereas the traffic from the IP phone uses a VLAN tag.
Which feature enables the interface to receive both types of traffic?

A. native VLAN
B. DHCP snooping
C. MAC limiting
D. voice VLAN

Answer: D

Explanation:
The feature that enables an interface on an EX Series switch to receive both untagged traffic (from the computer) and tagged traffic (from the IP phone) is the voice VLAN12.
The voice VLAN feature in EX-series switches enables access ports to accept both data (untagged) and voice (tagged) traffic and separate that traffic into different VLANs12. This allows the switch to differentiate between voice and data traffic, ensuring that voice traffic can be treated with a higher priority12. Therefore, option D is correct.

NEW QUESTION 13
Which two events cause a router to advertise a connected network to OSPF neighbors? (Choose two.)

A. When an OSPF adjacency is established.
B. When an interface has the OSPF passive option enabled.
C. When a static route to the 224.0.0.6 address is created.
D. When a static route to the 224.0.0.5 address is created.

Answer: AD

Explanation:
? A is correct because when an OSPF adjacency is established, a router will advertise a connected network to OSPF neighbors. An OSPF adjacency is a logical relationship between two routers that agree to exchange routing information using the OSPF protocol1. To establish an OSPF adjacency, the routers must be in the same area, have compatible parameters, and exchange hello packets1. Once an OSPF adjacency is formed, the routers will exchange database description (DBD) packets, which contain summaries of their link-state databases (LSDBs)1. The LSDBs include information about the connected networks and their costs2. Therefore, when an OSPF adjacency is established, a router will advertise
a connected network to OSPF neighbors through DBD packets.
? D is correct because when a static route to the 224.0.0.5 address is created, a router will advertise a connected network to OSPF neighbors. The 224.0.0.5 address is the multicast address for all OSPF routers3. A static route to this address can be used to send OSPF hello packets to all OSPF neighbors on a network segment3. This can be useful when the network segment does not support multicast or when the router does not have an IP address on the segment3. When a static route to the 224.0.0.5 address is created, the router will send hello packets to this address and establish OSPF adjacencies with other routers on the segment3. As explained above, once an OSPF adjacency is formed, the router will advertise a connected network to OSPF neighbors through DBD packets.

NEW QUESTION 17
What is a purpose of using a spanning tree protocol?

A. to look up MAC addresses
B. to eliminate broadcast storms
C. to route IP packets
D. to tunnel Ethernet frames

Answer: B

Explanation:
? A broadcast storm is a network condition where a large number of broadcast packets are sent and received by multiple devices, causing congestion and performance degradation1. A broadcast storm can occur when there are loops in the network topology, meaning that there are multiple paths between two devices2.
? A spanning tree protocol is a network protocol that prevents loops from being formed when switches or bridges are interconnected via multiple paths. It does this by creating a logical tree structure that spans all the devices in the network, and disabling or blocking the links that are not part of the tree, leaving a single active path between any two devices3.
? By eliminating loops, a spanning tree protocol also eliminates broadcast storms, as broadcast packets will not be forwarded endlessly along the looped paths. Instead, broadcast packets will be sent only along the tree structure, reaching each device once and avoiding congestion3.

NEW QUESTION 19
Which statement is correct about the storm control feature?

A. The storm control feature is enabled in the factory-default configuration on EX Series switches.
B. The storm control feature requires a special license on EX Series switches.
C. The storm control feature is not supported on aggregate Ethernet interfaces.
D. The storm control configuration only applies to traffic being sent between the forwarding and control plane.

Answer: A

Explanation:
? Option A is correct. The storm control feature is enabled in the factory-default configuration on EX Series switches12. On EX2200, EX3200, EX3300, EX4200, and EX6200 switches, the factory default configuration enables storm control for broadcast and unknown unicast traffic on all switch interfaces2. On EX4300 switches, the factory default configuration enables storm control on all Layer 2 switch interfaces1.
? Option B is incorrect. The storm control feature does not require a special license on EX Series switches34.
? Option C is incorrect. There??s no information available that suggests the storm control feature is not supported on aggregate Ethernet interfaces.
? Option D is incorrect. The storm control configuration applies to traffic at the ingress of an interface5, not just between the forwarding and control plane.

NEW QUESTION 20
Exhibit.

Exhibit

```
user@R1> show route receive-protocol bgp 10.36.1.4
inet.0: 33 destinations, 57 routes (33 active, 0 holddown, 0 hidden)
  Prefix    Nexthop          MED     Lclpref    AS path
* 10.30.100.8/32           10.36.1.4                          65401 65520 I
* 10.30.100.9/32           10.36.1.4                          65401 65521 I
* 10.30.189.0/30           10.36.1.4                          65401 65521 I
  10.32.1.0/30             10.36.1.4                          65401 I
* 10.32.2.0/30             10.36.1.4                          65401 I
* 10.32.12.0/30            10.36.1.4                          65401 I
* 10.52.100.2/32           10.36.1.4                          65401 I
```

You want to verify prefix information being sent from 10.36.1.4.
Which two statements are correct about the output shown in the exhibit? (Choose two.)

A. The routes displayed have traversed one or more autonomous systems.
B. The output shows routes that were received prior to the application of any BGP import policies.
C. The output shows routes that are active and rejected by an import policy.
D. The routes displayed are being learned from an I BGP peer.

Answer: AB

Explanation:
The output shown in the exhibit is the result of the command ??show ip bgp neighbor 10.36.1.4 received-routes??, which displays all received routes (both accepted and rejected) from the specified neighbor.
Option A is correct, because the routes displayed have traversed one or more autonomous systems. This can be seen from the AS_PATH attribute, which shows the sequence of AS numbers that the route has passed through. For example, the route 10.0.0.0/8 has an AS_PATH of 65001 65002, which means that it has traversed AS 65001 and AS 65002 before reaching the local router.
Option B is correct, because the output shows routes that were received prior to the application of any BGP import policies. This can be seen from the fact that some routes have a status code of ??r??, which means that they are rejected by an import policy. The??received-routes?? keyword shows the routes coming from a given neighbor before the inbound policy has been applied. To see the routes after the inbound policy has been applied, the ??routes?? keyword should be used instead.
Option C is incorrect, because the output does not show routes that are active and rejected by an import policy. The status code of ??r?? means that the route is rejected by an import policy, but it does not mean that it is active. The status code of ??>?? means that the route is active and selected as the best path. None of the routes in the output have both ??>?? and ??r?? status codes.
Option D is incorrect, because the routes displayed are not being learned from an IBGP peer. An IBGP peer is a BGP neighbor that belongs to the same AS as the local router. The output shows that the neighbor 10.36.1.4 has a remote AS of 65001, which is different from the local AS of 65002. Therefore, the neighbor is an EBGP peer, not an IBGP peer.

NEW QUESTION 22
You deployed a new EX Series switch with DHCP snooping enabled and you do not see any entries in the snooping databases for an interface. Which two Juniper configurations for that interface caused this issue? (Choose two.)

A. The interface is configured as a disabled port.
B. MAC limiting is enabled on the interface.
C. The interface is configured as a trunk port.
D. Dynamic ARP inspection is enabled on the interface.

Answer: AC

Explanation:
? A is correct because the interface is configured as a disabled port. A disabled port does not forward any traffic, including DHCP packets. Therefore, DHCP snooping cannot learn any MAC addresses or lease information from a disabled port1.
? C is correct because the interface is configured as a trunk port. By default, all trunk ports on the switch are trusted for DHCP snooping2. This means that DHCP snooping does not inspect or filter any DHCP packets received on a trunk port. Therefore, DHCP snooping does not add any entries to the snooping database for a trunk port2.

NEW QUESTION 23
Which statement is correct about graceful Routing Engine switchover (GRES)?

A. The PFE restarts and the kernel and interface information is lost.
B. GRES has a helper mode and a restarting mode.
C. When combined with NSR, routing is preserved and the new master RE does not restart rpd.
D. With no other high availability features enabled, routing is preserved and the new master RE does not restart rpd.

Answer: C

Explanation:
The Graceful Routing Engine Switchover (GRES) feature in Junos OS enables a router with redundant Routing Engines to continue forwarding packets, even if one Routing Engine fails1. GRES preserves interface and kernel information, ensuring that traffic is not interrupted1. However, GRES does not preserve the control plane1.
To preserve routing during a switchover, GRES must be combined with either Graceful
Restart protocol extensions or Nonstop Active Routing (NSR)1. When GRES is combined with NSR, nearly 75 percent of line rate worth of traffic per Packet Forwarding Engine remains uninterrupted during GRES1. Any updates to the primary Routing Engine are replicated to the backup Routing Engine as soon as they occur1.
Therefore, when GRES is combined with NSR, routing is preserved and the new master RE does not restart rpd1.

NEW QUESTION 26
Exhibit

Exhibit

```
user@host# show
  protocols {
        oam {
              gre-tunnel {
                    interface gr-1/1/10.1 {
                          keepalive-time 10;
                                hold-time 10;
                    }
              }
        }
        lldp {
              interface all;
        }
  }
```

You have configured a GRE tunnel. To reduce the risk of dropping traffic, you have configured a keepalive OAM probe to monitor the state of the tunnel; however, traffic drops are still occurring.
Referring to the exhibit, what is the problem?

A. For GRE tunnels, the OAM protocol requires that the BFD protocols also be used.
B. The "event link-adjacency-loss" option must be set.
C. LLDP needs to be removed from the gr-1/1/10.1 interface.
D. The hold-time value must be two times the keepalive-time value

Answer: D

Explanation:
A keepalive OAM probe is a mechanism that can be used to monitor the state of a GRE tunnel and detect any failures in the tunnel path. A keepalive OAM probe consists of sending periodic packets from one end of the tunnel to the other and expecting a reply. If no reply is received within a specified time, the tunnel is considered down and the line protocol of the tunnel interface is changed to down1.
To configure a keepalive OAM probe for a GRE tunnel, you need to specify two parameters: the keepalive-time and the hold-time. The keepalive-time is the interval between each keepalive packet sent by the local router. The hold-time is the maximum time that the local router waits for a reply from the remote router before declaring the tunnel down2.
According to the Juniper Networks documentation, the hold-time value must be two times the keepalive-time value for a GRE tunnel2. This is because the hold-time value must account for both the round-trip time of the keepalive packet and the processing time of the remote router. If the hold-time value is too small, it may cause false positives and unnecessary tunnel flaps.
In the exhibit, the configuration shows that the keepalive-time is set to 10 seconds and the hold-time is set to 15 seconds for the gr-1/1/10.1 interface. This means that the local router will send a keepalive packet every 10 seconds and will wait for 15 seconds for a reply from the remote router. However, this hold-time value is

not two times the keepalive-time value, which violates the recommended configuration. This may cause traffic drops if the remote router takes longer than 15 seconds to reply.
Therefore, option D is correct, because the hold-time value must be two times the keepalive-time value for a GRE tunnel. Option A is incorrect, because BFD is not required for GRE tunnels; BFD is another protocol that can be used to monitor tunnels, but it is not compatible with GRE keepalives3. Option B is incorrect, because the ??event link- adjacency-loss?? option is not related to GRE tunnels; it is an option that can be used to trigger an action when a link goes down4. Option C is incorrect, because LLDP does not need to be removed from the gr-1/1/10.1 interface; LLDP is a protocol that can be used to discover neighboring devices and their capabilities, but it does not interfere with GRE tunnels5.
References:
1: Configuring Keepalive Time and Hold time for a GRE Tunnel Interface 2: keepalive | Junos OS | Juniper Networks 3: Configuring Bidirectional Forwarding Detection 4: event link-adjacency-loss | Junos OS | Juniper Networks 5: Understanding Link Layer Discovery Protocol

NEW QUESTION 28
You are a network operator who wants to add a second ISP connection and remove the default route to the existing ISP You decide to deploy the BGP protocol in the network.
What two statements are correct in this scenario? (Choose two.)

A. IBGP updates the next-hop attribute to ensure reachability within an AS.
B. IBGP peers advertise routes received from EBGP peers to other IBGP peers.
C. IBGP peers advertise routes received from IBGP peers to other IBGP peers.
D. EBGP peers advertise routes received from IBGP peers to other EBGP peers.

Answer: AB

Explanation:
? A is correct because IBGP updates the next-hop attribute to ensure reachability within an AS. This is because the next-hop attribute is the IP address of the router that advertises the route to a BGP peer. If the next-hop attribute is not changed by IBGP, it would be the IP address of an external router, which may not be reachable by all routers within the AS. Therefore, IBGP updates the next-hop attribute to the IP address of the router that received the route from an EBGP peer1.
? B is correct because IBGP peers advertise routes received from EBGP peers to other IBGP peers. This is because BGP follows the rule of advertising only the best route to a destination, and EBGP routes have a higher preference than IBGP routes. Therefore, IBGP peers advertise routes learned from an EBGP peer to all BGP peers, including both EBGP and IBGP peers1.

NEW QUESTION 29
Which statement about aggregate routes is correct?

A. Aggregate routes can only be used for static routing but not for dynamic routing protocols.
B. Aggregate routes are automatically generated for all of the subnets in a routing table.
C. Aggregate routes are always preferred over more specific routes, even when the specific routes have a better path.
D. Aggregate routes are used for advertising summarized network prefixes.

Answer: D

Explanation:
Aggregate routes are used for advertising summarized network prefixes12. They help minimize the number of routing tables in an IP network by consolidating selected multiple routes into a single route advertisement1. This approach is in contrast to non-aggregation routing, in which every routing table contains a unique entry for each route1.
Therefore, option D is correct. Options A, B, and C are not correct because:
? Aggregate routes can be used with both static routing and dynamic routing protocols1.
? Aggregate routes are not automatically generated for all of the subnets in a routing table. They need to be manually configured1.
? Aggregate routes are not always preferred over more specific routes. The route selection process in Junos OS considers several factors, including route preference and metric, before determining the active route1.

NEW QUESTION 32
Which two statements about BGP facilitate the prevention of routing loops between two autonomous systems? (Choose two.)

A. EBGP routers will append their AS number when advertising routes to their neighbors.
B. EBGP routers will only accept routes that contain their own AS number in the AS_PATH.
C. EBGP routers will drop routes that contain their own AS number in the AS_PATH
D. EBGP routers will prepend their AS number when advertising routes to their neighbors

Answer: AC

Explanation:
BGP (Border Gateway Protocol) is a protocol designed to exchange routing and reachability information among autonomous systems (AS) on the internet1.
? Option A is correct. When an EBGP router advertises routes to its neighbors, it appends its AS number to the AS_PATH attribute1. This is a key mechanism in BGP to prevent routing loops1.
? Option C is correct. BGP has a built-in loop prevention mechanism whereby if a BGP router detects its own AS in the AS_PATH attribute, it will drop the prefix and will not continue to advertise it2. This helps to prevent routing loops2.
? Option B is incorrect. EBGP routers do not accept routes that contain their own AS number in the AS_PATH2. Instead, they drop such routes as part of the loop prevention mechanism2.
? Option D is incorrect. While it??s true that EBGP routers append their AS number
when advertising routes, they do not prepend their AS number1. The term ??prepend?? in BGP usually refers to a technique used to influence path selection by artificially lengthening the AS_PATH3.

NEW QUESTION 34
Exhibit

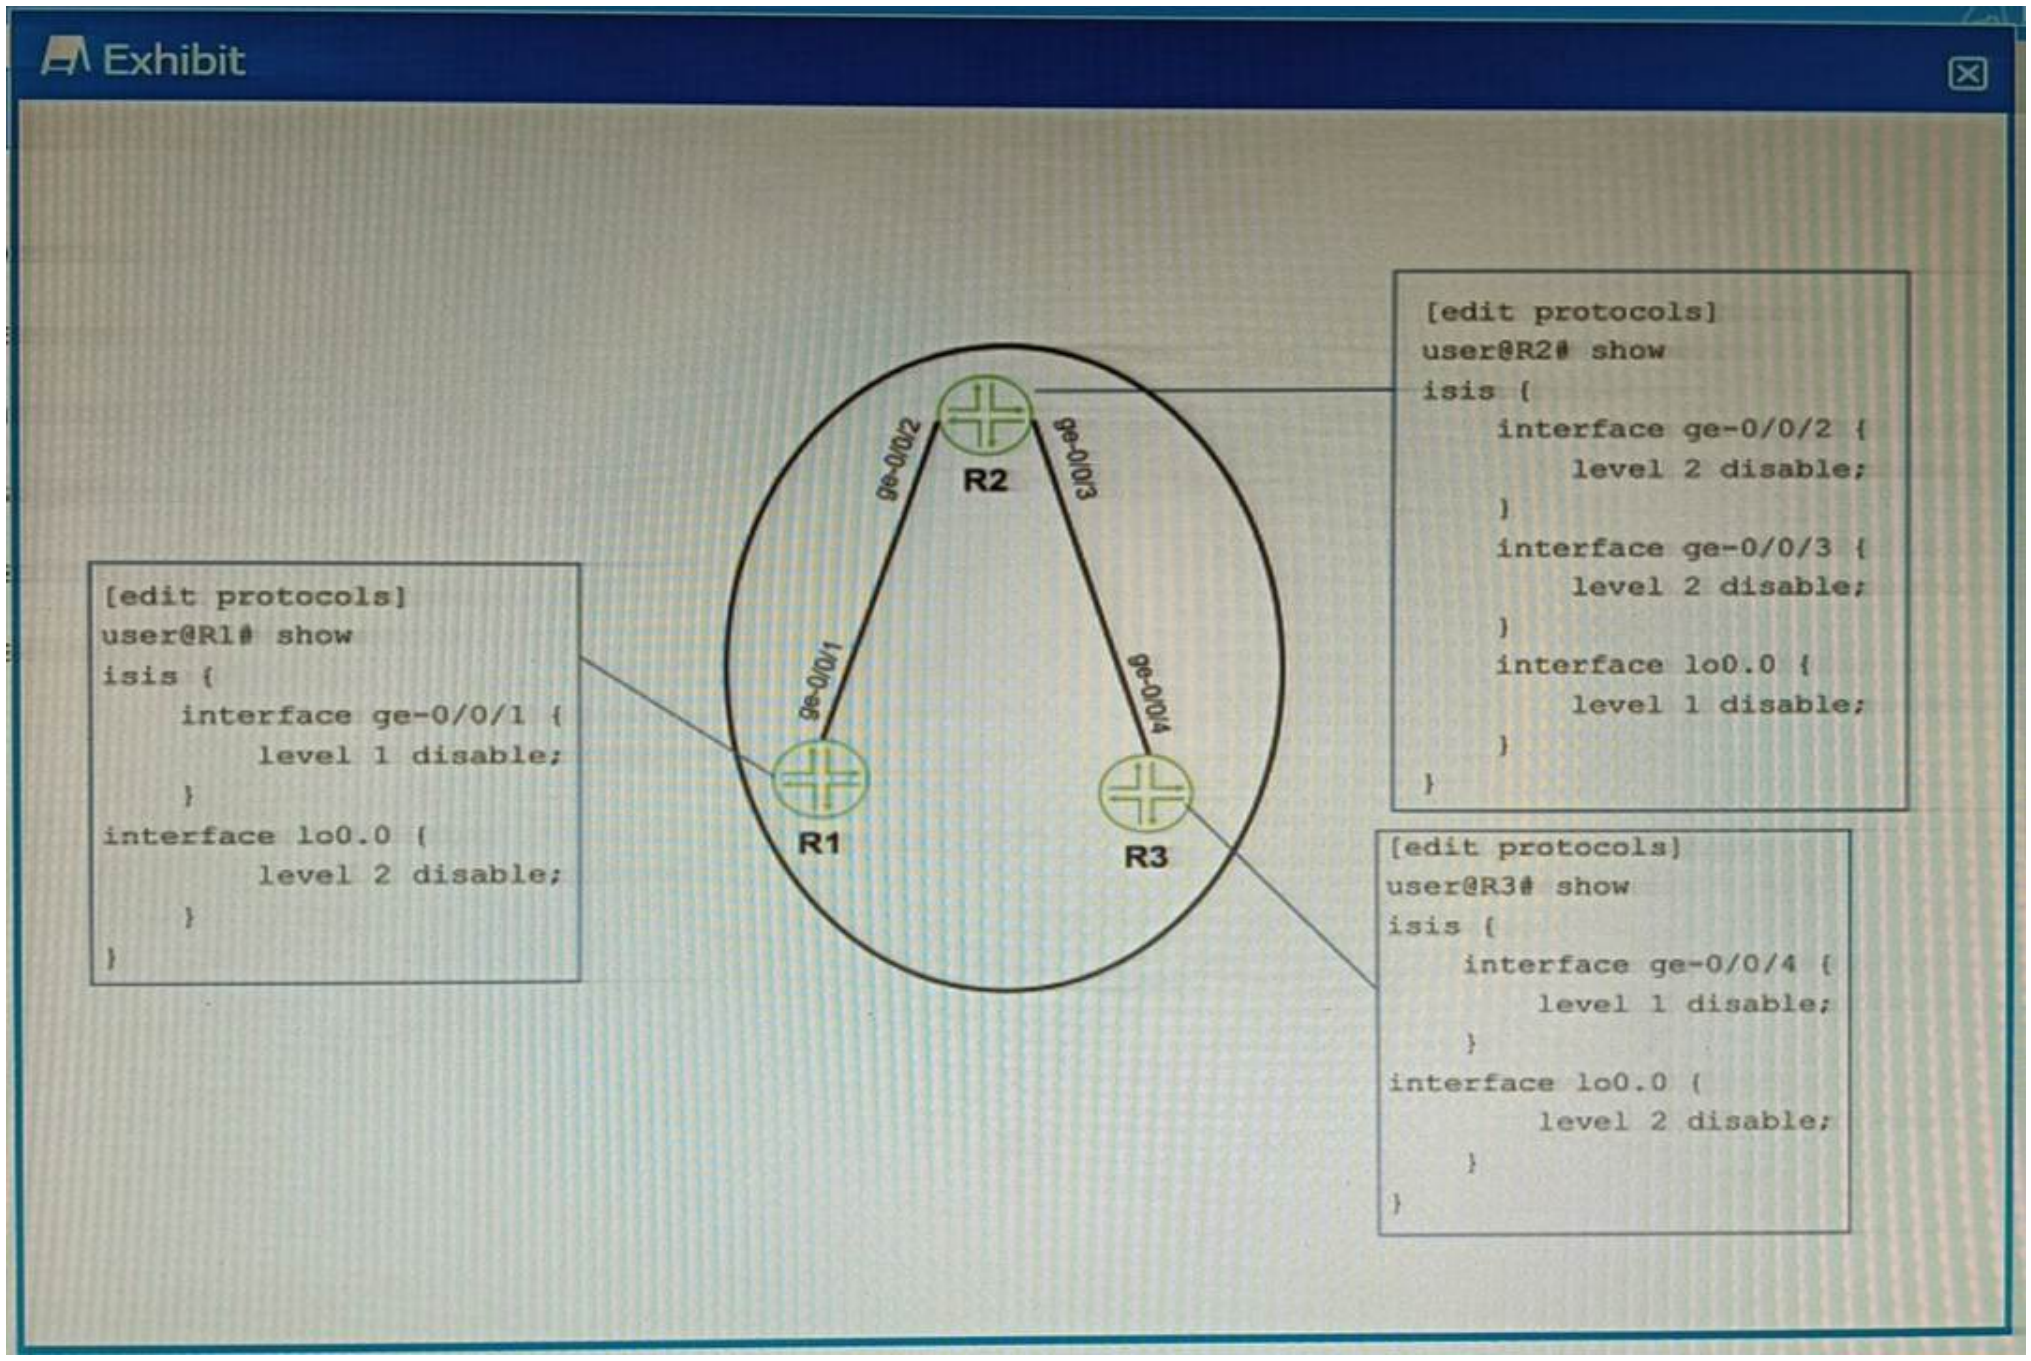

Referring to the exhibit, which two configuration changes must you apply for packets to reach from R1 to R3 using IS-IS? (Choose two.)

A. On R1, enable Level 1 on the ge-0/0/1 interface.
B. On R3 disable Level 2 on the ge-0/0/4 interface.
C. On R1, disable Level 2 on the ge-0/0/1 interface.
D. On R3 enable Level 1 on the ge-0/0/4 interface

Answer: AD

Explanation:
A. On R1, enable Level 1 on the ge-0/0/1 interface. In IS-IS, both levels (Level 1 and Level 2) are enabled by default when you enable IS-IS on an interface1. Level 1 systems route within an area2. If the destination is outside an area, Level 1 systems route toward a Level 2 system2. Therefore, enabling Level 1 on the ge-0/0/1 interface on
R1 would allow packets to reach from R1 to R3.
* D. On R3 enable Level 1 on the ge-0/0/4 interface Similarly, enabling Level 1 on the ge- 0/0/4 interface on R3 would allow packets to reach from R1 to R3.
These explanations are based on the IS-IS configuration documents and learning resources available at Juniper Networks1 and Cisco34.

NEW QUESTION 35
.....

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

JN0-351 Practice Exam Features:

* JN0-351 Questions and Answers Updated Frequently

* JN0-351 Practice Questions Verified by Expert Senior Certified Staff

* JN0-351 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* JN0-351 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year