Juniper

Exam Questions JN0-351

Enterprise Routing and Switching - Specialist (JNCIS-ENT)

NEW QUESTION 1
Exhibit

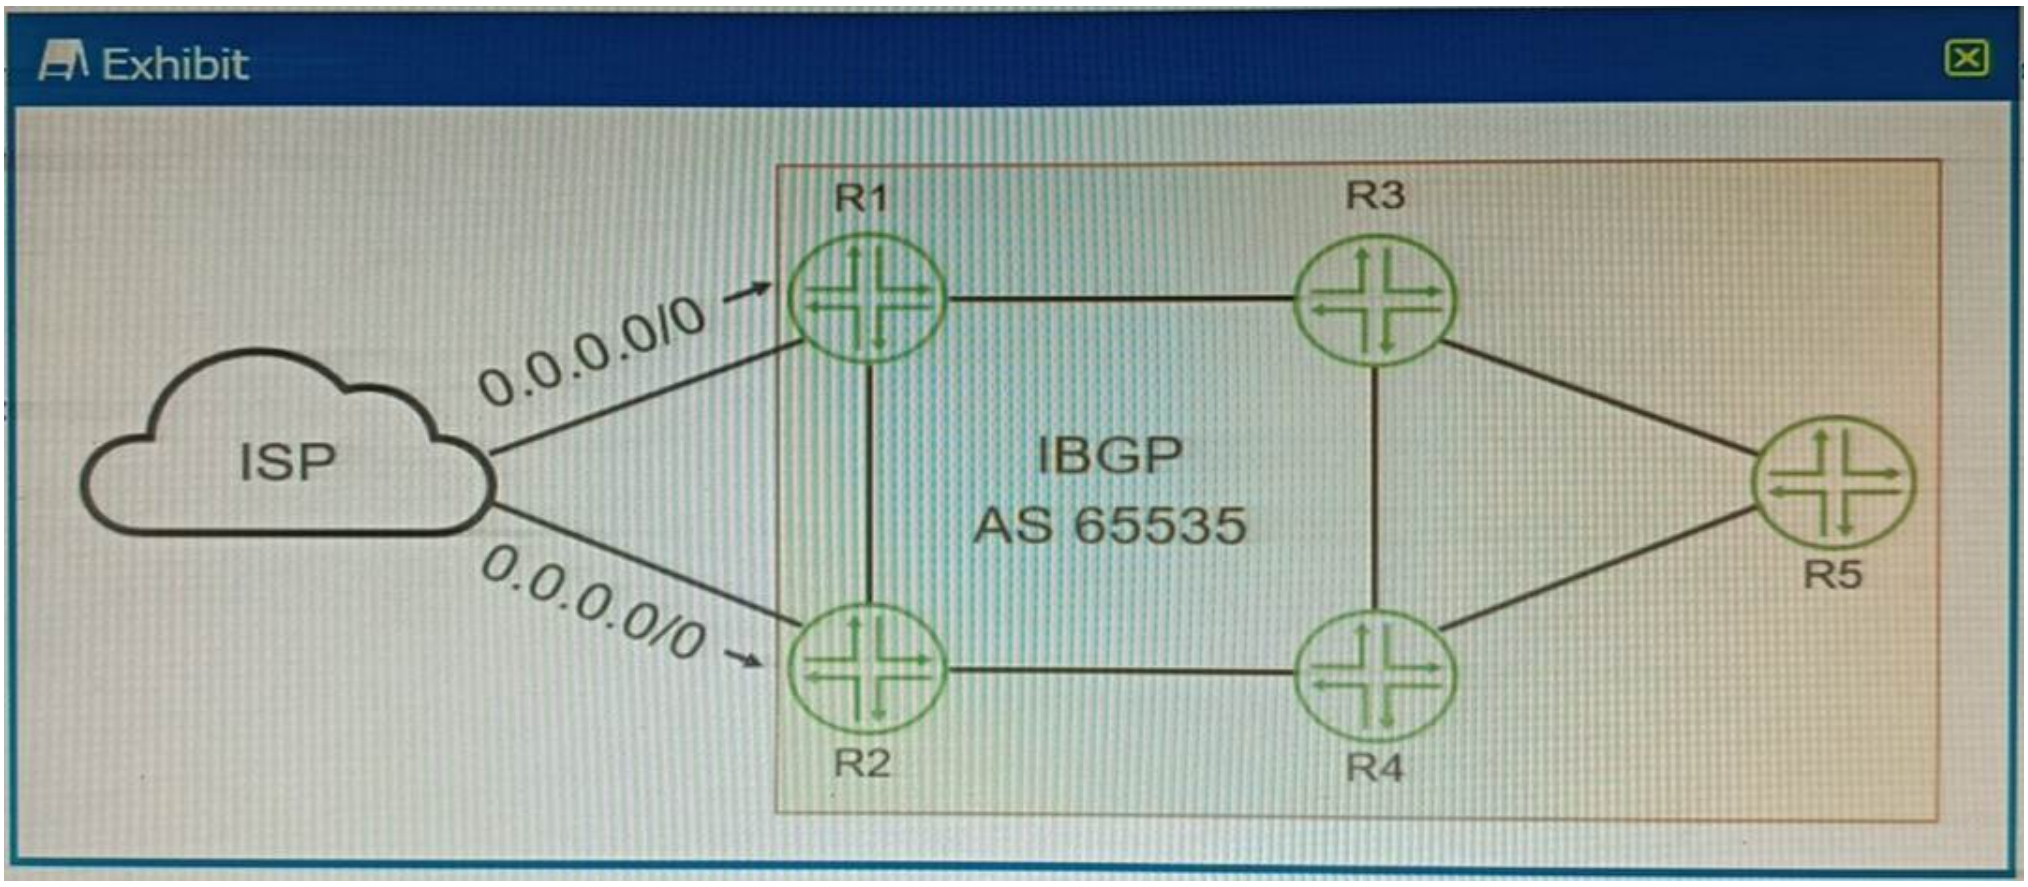

Your ISP is announcing a default route to both R1 and R2. You want your network routers to forward all Internet traffic through the R1 device
Which BGP attribute would you use?

A. MED
B. next-hop
C. local preference
D. origin

Answer: C

Explanation:
The BGP attribute that you would use to forward all Internet traffic through the R1 device is the local preference1.
The local preference is an attribute that is used within an autonomous system (AS) and exchanged between iBGP routers1. It is used to select an exit point from the AS1. The path with the highest local preference is preferred1. By setting a higher local preference for the routes received from R1, you can make R1 the preferred exit point for all Internet traffic1.

NEW QUESTION 2
Exhibit.

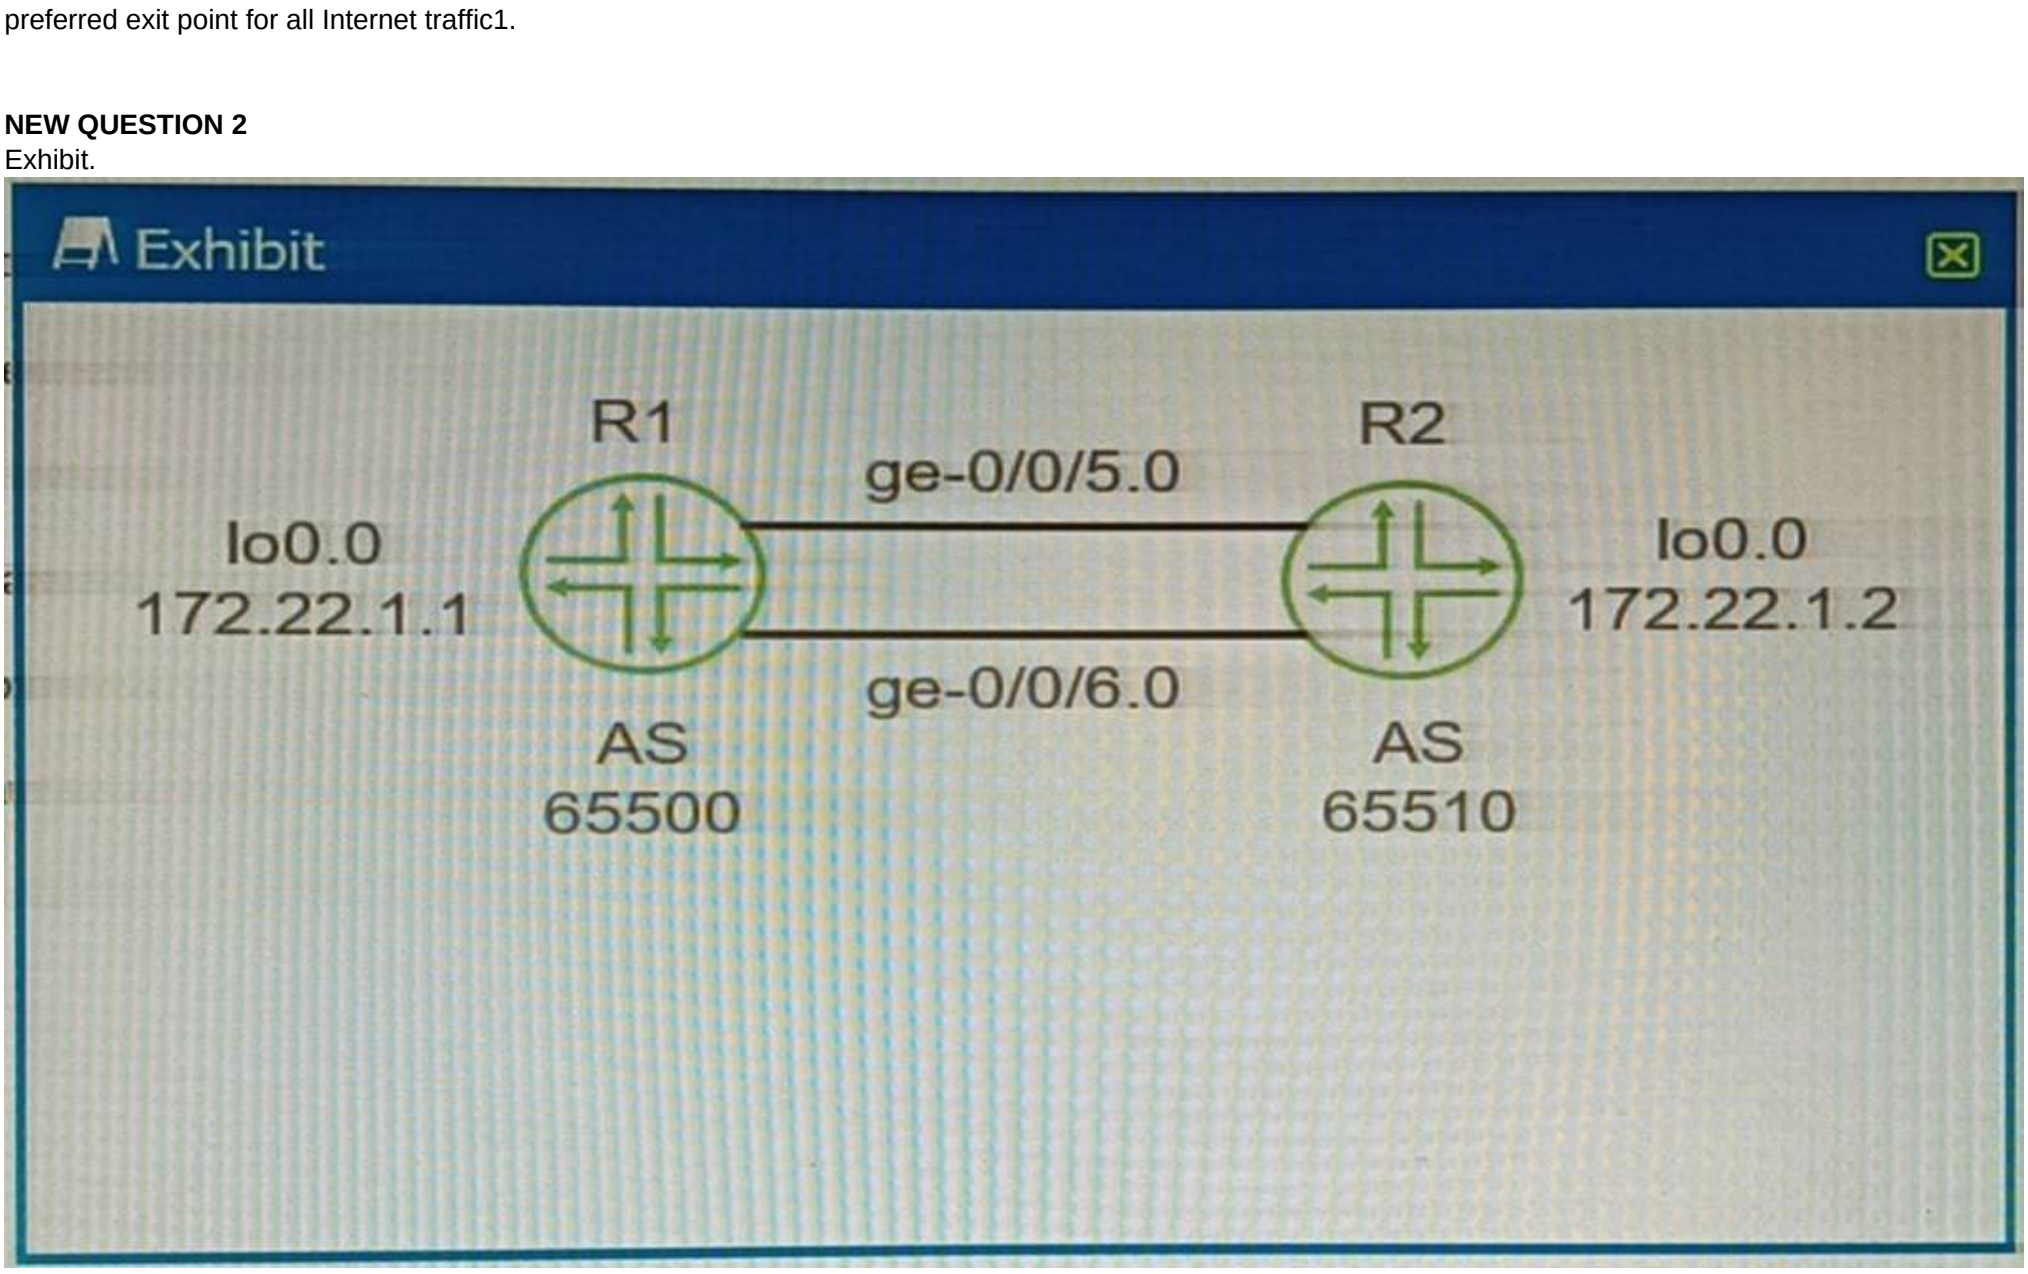

You want to enable redundancy for the EBGP peering between the two routers shown in the exhibit. Which three actions will you perform in this scenario? (Choose three.)

A. Configure BGP multihop.
B. Configure loopback interface peering.
C. Configure routes for the peer loopback interface IP addresses.
D. Configure an MD5 peer authentication.
E. Configure a cluster ID.

Answer: ABC

Explanation:
? A is correct because you need to configure BGP multihop to enable redundancy for the EBGP peering between the two routers. BGP multihop is a feature that allows BGP peers to establish a session over multiple hops, instead of requiring them to be directly connected1. By default, EBGP peers use a time-to-live (TTL) value of 1 for their packets, which means that they can only reach adjacent neighbors1. However, if you configure BGP multihop with a higher TTL value, you can allow EBGP peers to communicate over multiple routers in between1. This can provide redundancy in case of a link failure or a router failure between the EBGP peers.
? B is correct because you need to configure loopback interface peering to enable redundancy for the EBGP peering between the two routers. Loopback interface peering is a technique that uses loopback interfaces as the source and destination addresses for BGP sessions, instead of physical interfaces2. Loopback interfaces are virtual interfaces that are always up andreachable as long as the router is operational2. By using loopback interface peering, you can avoid the dependency on a single physical interface or link for the BGP session, and use multiple paths to reach the loopback address of the peer2. This can provide redundancy and load balancing for the EBGP peering.
? C is correct because you need to configure routes for the peer loopback interface
IP addresses to enable redundancy for the EBGP peering between the two routers. Routes for the peer loopback interface IP addresses are necessary to ensure that the routers can reach each other??s loopback addresses over multiple hops2. You can use static routes or dynamic routing protocols to advertise and learn the routes for the peer loopback interface IP addresses2. Without these routes, the routers will not be able to establish or maintain the BGP session using their loopback interfaces.

NEW QUESTION 3
Exhibit

Exhibit

```
user# show protocols bgp


group ext-64501 {
    type external;
    peer-as 64501;
    neighbor 172.30.1.2;
}
group int-64503 {
    type internal;
    local-address 192.168.100.1;
    neighbor 192.168.100.2;
}
bfd-liveness-detection {
    minimum-interval 10;
}
```

Your BGP neighbors, one in the USA and one in France, are not establishing a connection with each other.
Referring to the exhibit, which statement is correct?

A. The BFD liveness is set too low.
B. The BFD liveness must be configured on the BGP neighbor.
C. The BFD liveness must be configured on the BGP group.
D. The BFD liveness is set too high.

Answer: B

Explanation:
? The exhibit shows the configuration of BFD liveness detection for BGP at the global level, which applies to all BGP neighbors by default1. However, this configuration does not specify the session mode, which determines whether BFD uses single-hop or multihop mode to communicate with a neighbor2.
? For single-hop BGP neighbors, which are directly connected on the same subnet, the session mode can be either automatic or single-hop. For multihop BGPneighbors, which are not directly connected and require multiple hops to reach, the session mode must be multihop2.
? Since your BGP neighbors are in different countries, they are likely to be multihop neighbors. Therefore, you need to configure the session mode as multihop for each neighbor individually at the [edit protocols bgp group group-name neighbor address bfd-liveness-detection] hierarchy level2. For example:
protocols { bgp { group usa { neighbor 192.0.2.1 { bfd-liveness-detection { session-mode multihop; } } } group france { neighbor 198.51.100.1 { bfd-liveness-

detection { session-mode multihop; } } } } }
? If you do not configure the session mode for multihop neighbors, BFD will use the
default mode of automatic, which will try to use single-hop mode and fail to establish a BFD session with the remote neighbor2. This will prevent BGP from using BFD to detect liveliness and failover.
? Therefore, the answer B is correct, as you need to configure the BFD liveness
detection on the BGP neighbor level with the appropriate session mode for multihop neighbors.

NEW QUESTION 4
Exhibit.

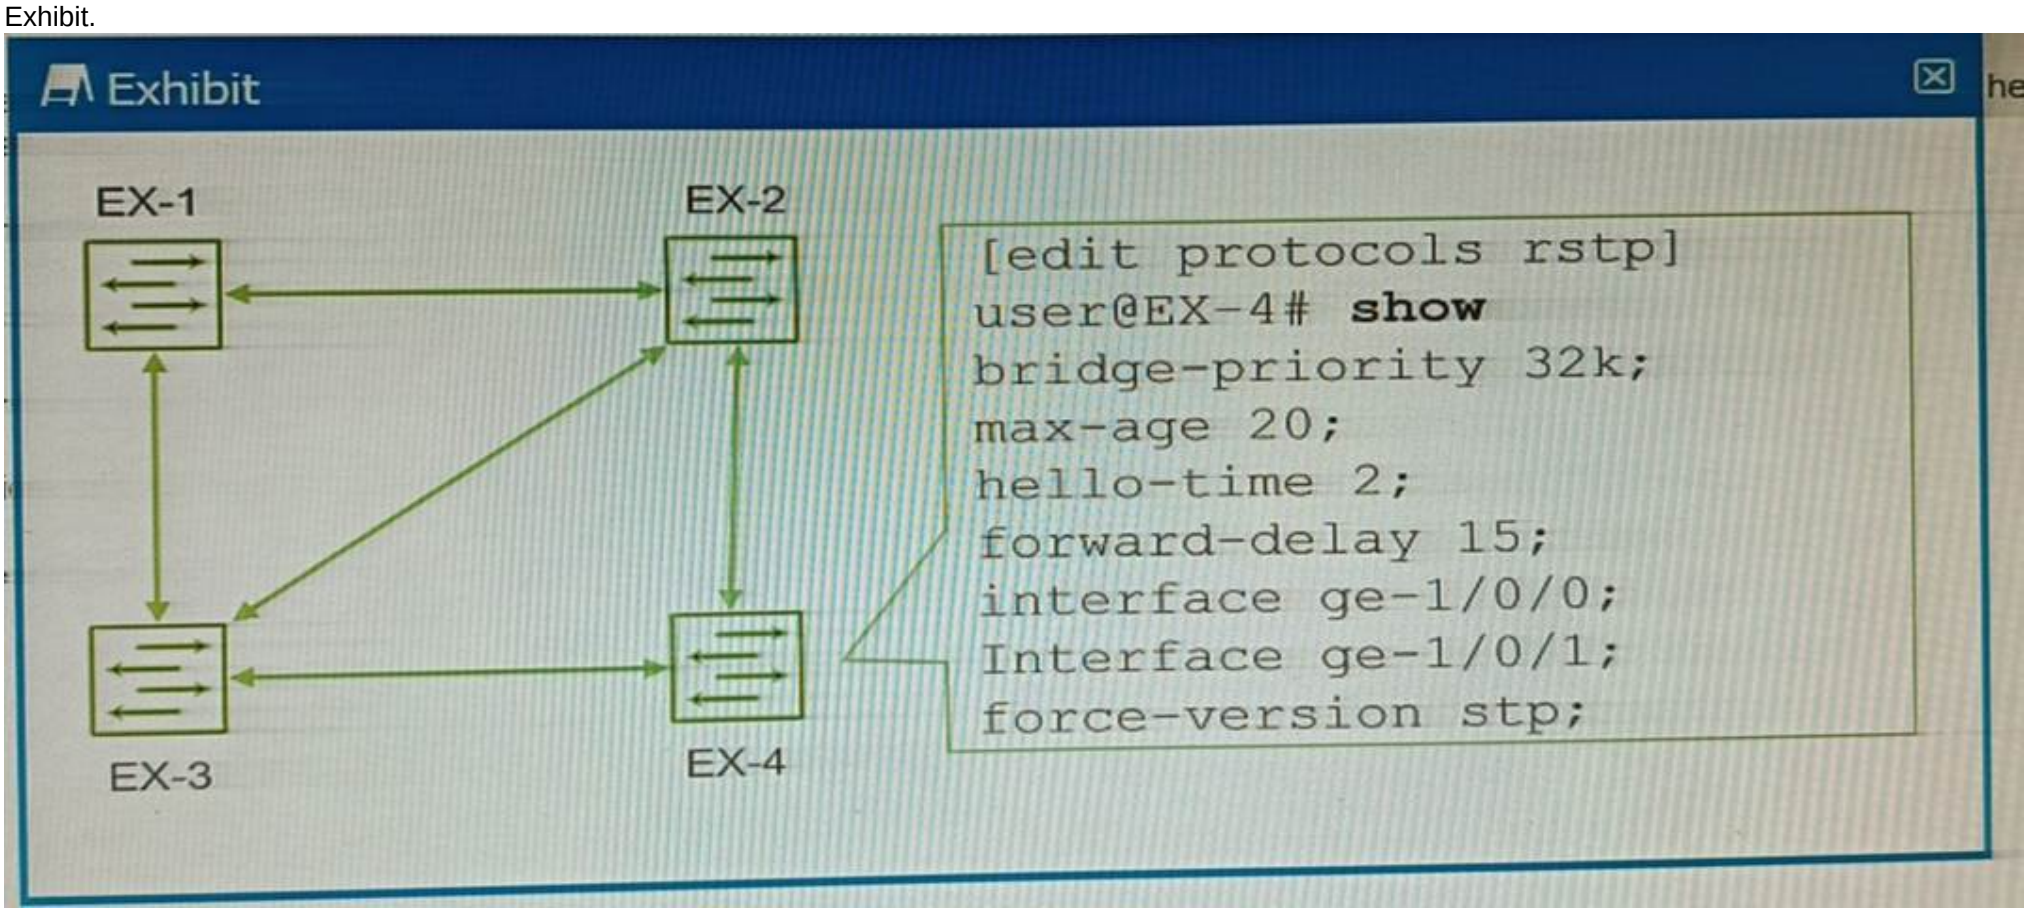

You have configured the four EX Series switches with RSTP, as shown in the exhibit. You discover that whenever a link between switches goes up or down, the switches take longer
than expected for RSTP to converge, using the default settings.
In this scenario, which action would solve the delay in RSTP convergence?

A. The hello-time must be increased.
B. The force-version must be removed.
C. The bridge priority for EX-4 must be set at 4000.
D. The max-age must be increased to 20

Answer: B

Explanation:
? The exhibit shows the configuration of RSTP on EX-4, which has the command force-version stp. This command forces the switch to use the legacy STP protocol instead of RSTP, even though the switch supports RSTP1. This means that EX-4 will not be able to take advantage of the faster convergence and enhanced features of RSTP, such as edge ports, link type, and proposal/agreement sequence2.
? The other switches in the network are likely to be running RSTP, as it is the default protocol for EX Series switches3. Therefore, there will be a compatibility issue between EX-4 and the other switches, which will result in longer convergence times and suboptimal performance. The switch will also generate a warning message that says ??Warning: STP version mismatch with neighbor?? when it receives a BPDU from a RSTP neighbor1.
? To solve this problem, the force-version command must be removed from EX-4, so that it can run RSTP natively and interoperate with the other switches in the network. This will enable faster convergence and better stability for the network topology. To remove the command, you can use the delete protocols rstp force-version command in configuration mode1.

NEW QUESTION 5
Which two BGP attributes must be supported by all BGP implementations and must be included in every update? (Choose two.)

A. AS path
B. MED
C. next hop
D. community

Answer: AC

Explanation:
BGP attributes are properties that BGP uses for route advertisement, path selection, and loop prevention1. There are four categories of BGP attributes123:
? Well-known mandatory: Must be recognized by all BGP routers, present in all BGP updates, and passed on to other BGP routers123.
? Well-known discretionary: Supported by all BGP implementations, and are optionally included in BGP updates1.
? Optional transitive: May not be supported by all implementations of BGP1.
? Optional non-transitive: May not be supported by all implementations of BGP1. The well-known mandatory attributes must be supported by all BGP implementations and must be included in every update123. These include the AS path and next hop attributes23. Therefore, options A and C are correct.

NEW QUESTION 6
Which two statements about redundant trunk groups on EX Series switches are correct? (Choose two.)

A. Redundant trunk groups load-balance traffic across two designated uplink interfaces.

B. If the active link fails, then the secondary link automatically takes over.
C. Layer 2 control traffic is permitted on the secondary link
D. Redundant trunk groups must be connected to the same aggregation switch.

Answer: BD

Explanation:
Redundant Trunk Groups (RTGs) on EX Series switches provide a simple solution for network recovery when a trunk port on a switch goes down1. They are configured on the access switch and contain two links: a primary or active link, and a secondary link1. Therefore, option B is correct because if the active link fails, the secondary link automatically starts forwarding data traffic without waiting for normal spanning-tree protocol convergence1.
Option D is also correct. In a typical enterprise network composed of distribution and access layers, RTGs are used where one Access switch is connected to two different uplink switches2. This implies that RTGs must be connected to the same aggregation switch2.

NEW QUESTION 7
Exhibit

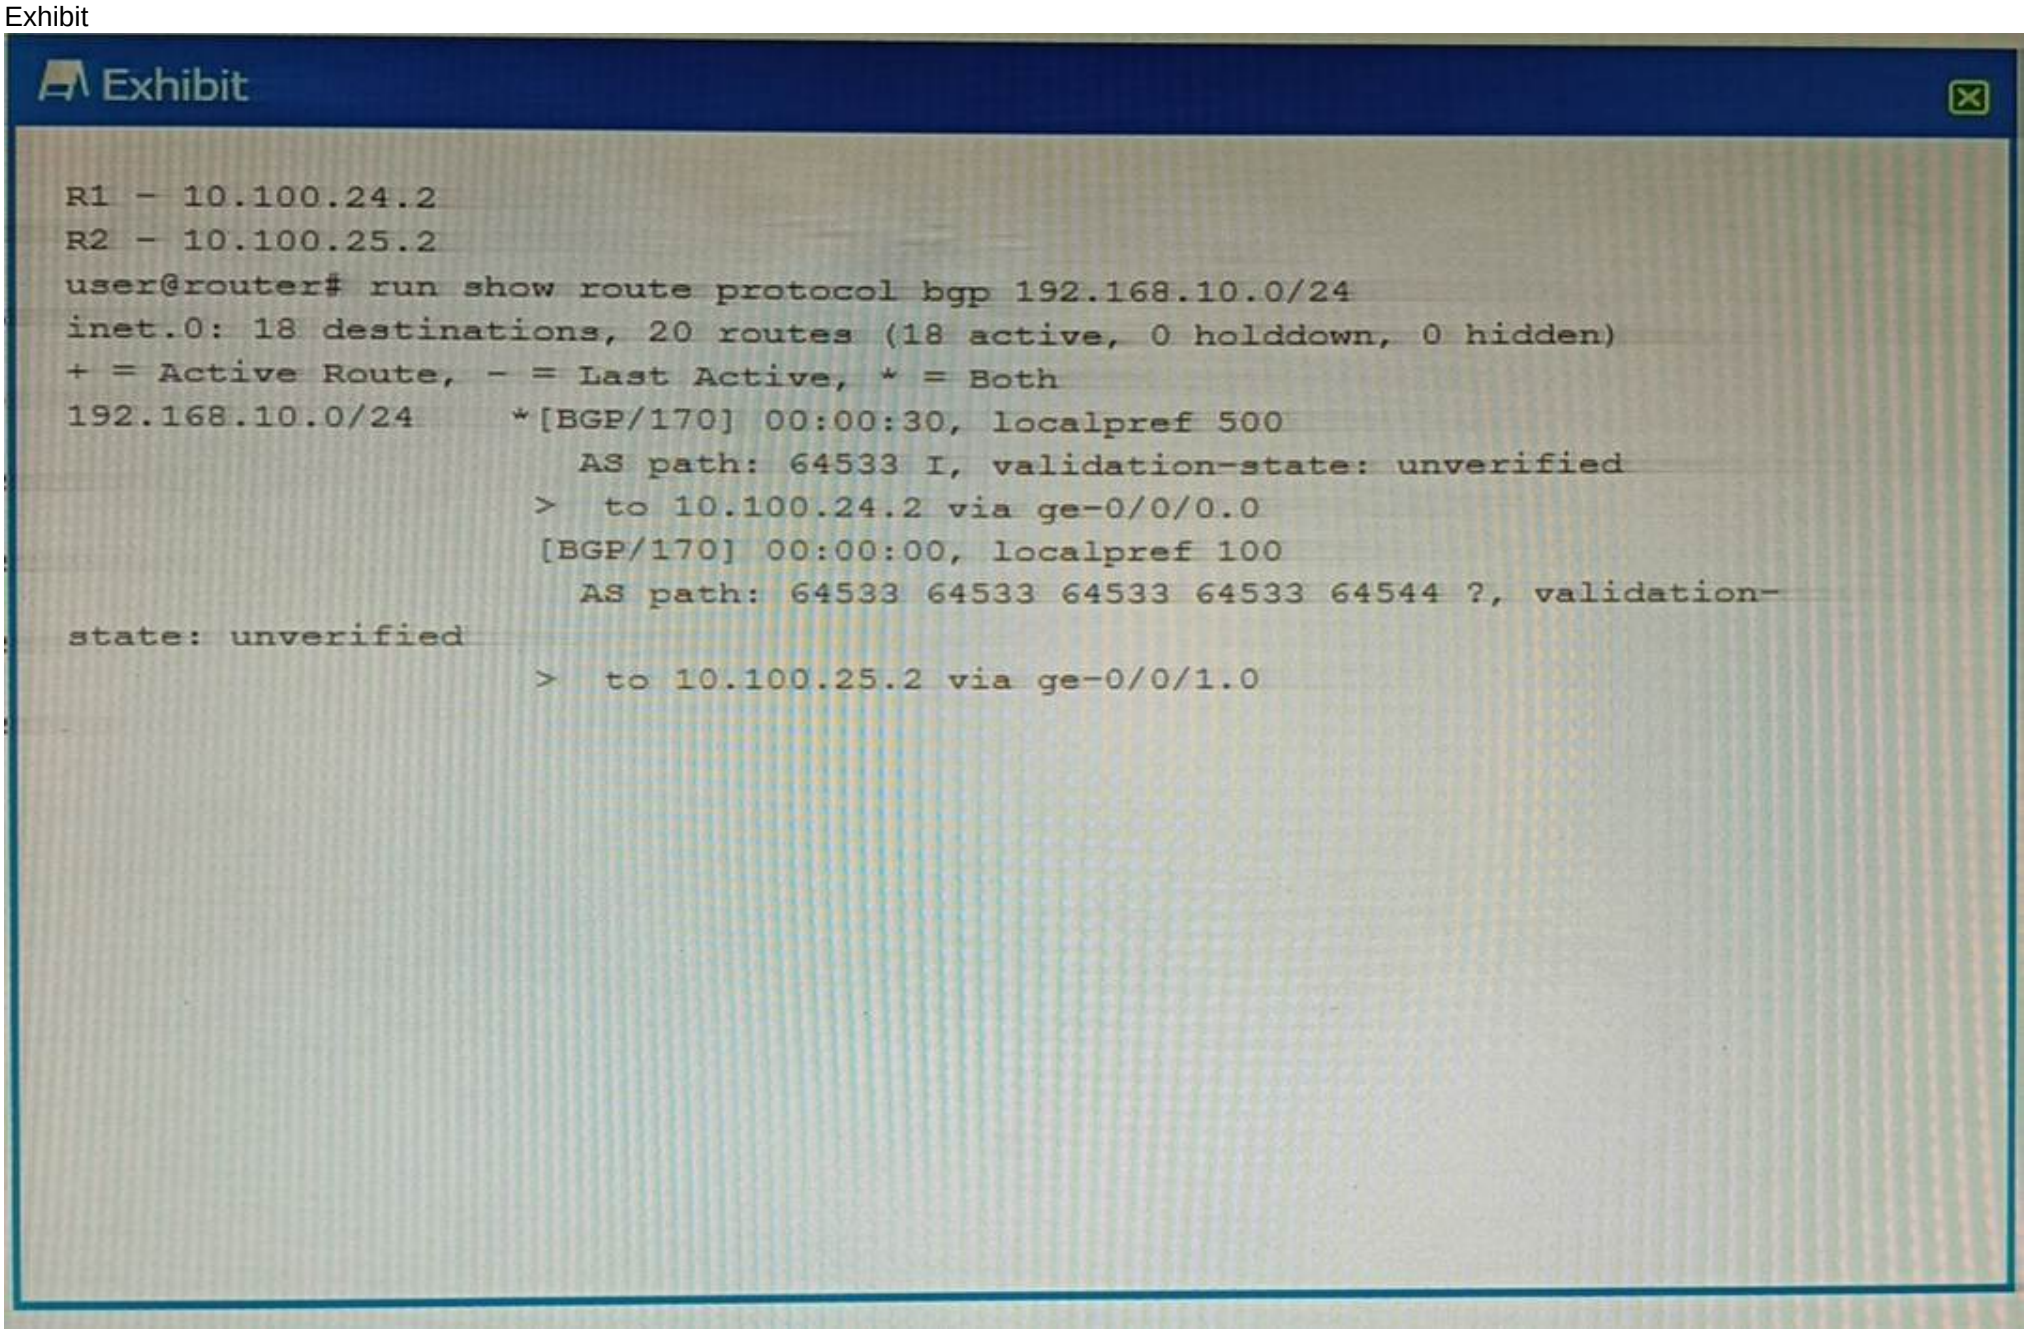

You are troubleshooting an issue where traffic to 192.168.10.0/24 is being sent to R1 instead of your desired path through R2.
Referring to the exhibit, what is the reason for the problem?

A. R2's route is not the best path due to loop prevention.
B. R2's route is not the best path due to a lower origin code.
C. R1's route is the best path due to a higher local preference
D. R1's route is the best path due to the shorter AS path.

Answer: C

Explanation:
? The exhibit shows the output of the command show ip bgp, which displays information about the BGP routes in the routing table1. The output shows two routes for the destination 192.168.10.0/24, one from R1 and one from R2.
? The route from R1 has a local preference of 200, while the route from R2 has a local preference of 100. Local preference is a BGP attribute that indicates the degree of preference for a route within an autonomous system (AS)2. A higher local preference means a more preferred route2.
? BGP uses a best path selection algorithm to choose the best route for each destination among multiple paths. The algorithm compares different attributes of the routes in a specific order of precedence3. The first attribute that is compared is weight, which is a Cisco-specific attribute that is local to the router3. If the weight is equal or not set, the next attribute that is compared is local preference3.
? In this case, both routes have the same weight of 0, which means that they are learned from external BGP (eBGP) peers3. Therefore, the next attribute that is compared is local preference. Since R1??s route has a higher local preference than R2??s route, it is chosen as the best path and installed in the routing table3. The other attributes, such as origin code and AS path, are not considered in this case.

NEW QUESTION 8
Exhibit.

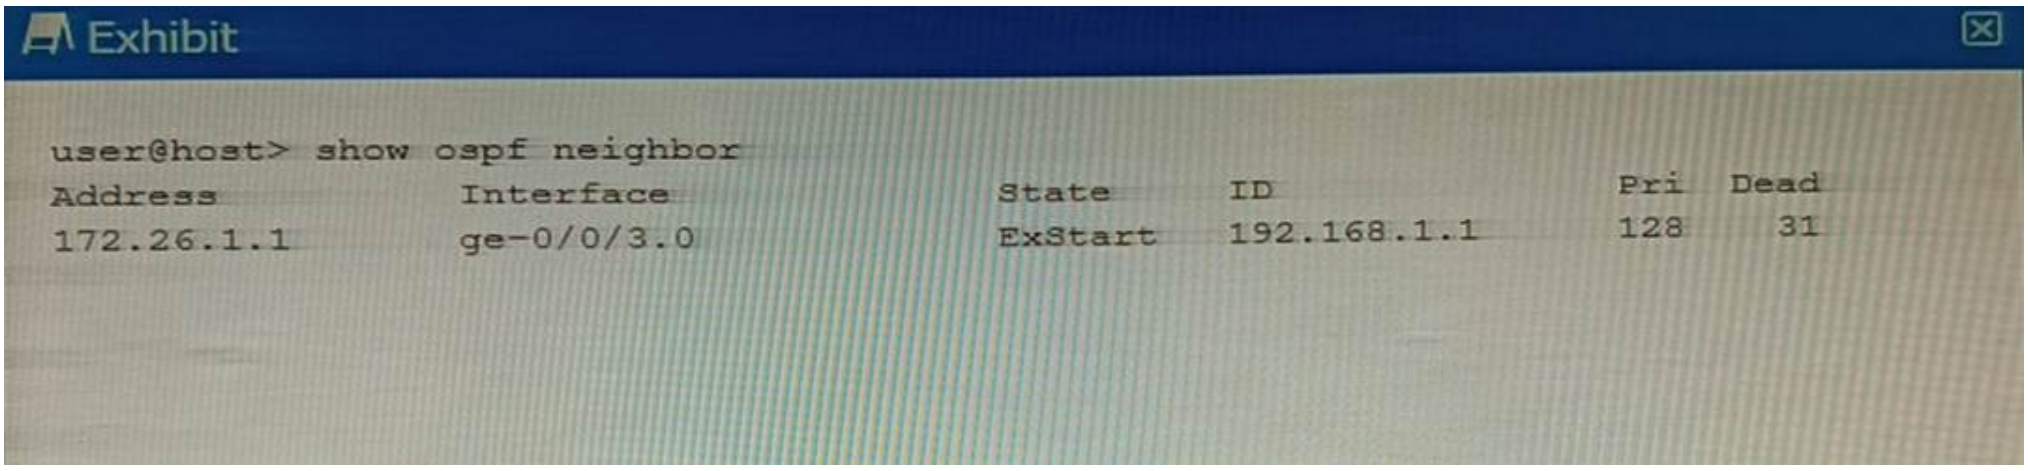

Why is this OSPF adjacency remaining in this state?

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 9
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 10
You are an operator for a network running 1S-IS. Two routers are failing to form an adjacency. What are two reasons for this problem? (Choose two.)

A. There are mismatched router IDs on the L2 routers.
B. There is no configured ISO address on any IS-IS interface.
C. There is a mismatched area ID between the L2 routers.
D. The family iso configuration is missing from the adjacency interface.

Answer: BD

Explanation:
The two reasons for the failure to form an adjacency in a network running IS- IS could be:
* B. There is no configured ISO address on any IS-IS interface. IS-IS requires each router interface to have an ISO address configured. Without this address, the routers cannot form an adjacency1.
* D. The family iso configuration is missing from the adjacency interface. The ??family iso?? configuration is essential for IS-IS to function correctly. If this configuration is missing from the adjacency interface, it could prevent the formation of an adjacency1.
These explanations are based on the Enterprise Routing and Switching Specialist (JNCIS- ENT) documents and learning resources available at Juniper Networks23.

NEW QUESTION 10
An update to your organization's network security requirements document requires management traffic to be isolated in a non-default routing-instance. You want to implement
this requirement on your Junos-based devices.
Which two commands enable this behavior? (Choose two.)

A. set routing—instances mgmtjunoa interface ge-0/0/0.0

B. set routing—instances mgmt_junos interface em1
C. set system management—instance
D. set routing—instances mgmt_junos

Answer: CD

Explanation:
To isolate management traffic in a non-default routing-instance on Junos- based devices, you can use the set system management-instance and set routing-instances mgmt_junos commands12.
? set system management-instance: This command associates the management
interface (usually named fxp0 or em0 for Junos OS, or re0:mgmt-* or re1:mgmt-* for Junos OS Evolved) with the non-default virtual routing and forwarding (VRF) instance1. After you configure the non-default management VRF instance, management traffic no longer has to share a routing table with other control traffic or protocol traffic1.
? set routing-instances mgmt_junos: This command creates a new routing instance
named mgmt_junos. The name of the dedicated management VRF instance is reserved and hardcoded as mgmt_junos; you cannot configure any other routing instance by the name mgmt_junos1.
Therefore, options C and D are correct. Options A and B are not correct because they attempt to assign an interface to the mgmt_junos routing instance, which is not necessary for isolating management traffic1.

NEW QUESTION 14
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 18
Which statement is correct about IP-IP tunnels?

A. IP-IP tunnels only support encapsulating IP traffic.
B. IP-IP tunnels only support encapsulating non-IP traffic.
C. The TTL in the inner packet is decremented during transit to the tunnel endpoint.
D. There are 24 bytes of overhead with IP-IP encapsulation.

Answer: A

Explanation:
IP-IP tunnels are a type of tunnels that use IP as both the encapsulating and encapsulated protocol. IP-IP tunnels are simple and easy to configure, but they do not provide any security or authentication features. IP-IP tunnels only support encapsulating IP traffic, which means that the payload of the inner packet must be an IP packet. IP-IP tunnels cannot encapsulate non-IP traffic, such as Ethernet frames or MPLS labels1.
Option A is correct, because IP-IP tunnels only support encapsulating IP traffic. Option B is incorrect, because IP-IP tunnels only support encapsulating non-IP traffic. Option C is incorrect, because the TTL in the inner packet is not decremented during transit to the tunnel endpoint. The TTL in the outer packet is decremented by each router along the path, but the TTL in the inner packet is preserved until it reaches the tunnel endpoint2. Option D is incorrect, because there are 20 bytes of overhead with IP-IP encapsulation. The overhead consists of the header of the outer packet, which has a fixed size of 20 bytes for IPv43.
References:
1: IP-IP Tunneling 2: What is tunneling? | Tunneling in networking 3: IPv4 - Header

NEW QUESTION 22
You are configuring an IS-IS IGP network and do not see the IS-IS adjacencies established. In this scenario, what are two reasons for this problem? (Choose two.)

A. MTU is not at least 1492 bytes.
B. IP subnets are not a /30 address.
C. The Level 2 routers have mismatched areas.
D. The lo0 interface is not included as an IS-IS interface.

Answer: AD

Explanation:
Option A suggests that the MTU is not at least 1492 bytes. This is correct because IS-IS requires a minimum MTU of 1492 bytes to establish adjacencies1. If the MTU is less than this, IS-IS adjacencies will not be established1.
Option D suggests that the lo0 interface is not included as an IS-IS interface. This is also correct because the loopback interface (lo0) is typically used as the router ID in IS-IS1. If the loopback interface is not included in IS-IS, it could prevent IS-IS adjacencies from being established1.
Therefore, options A and D are correct.

NEW QUESTION 24
Exhibit.

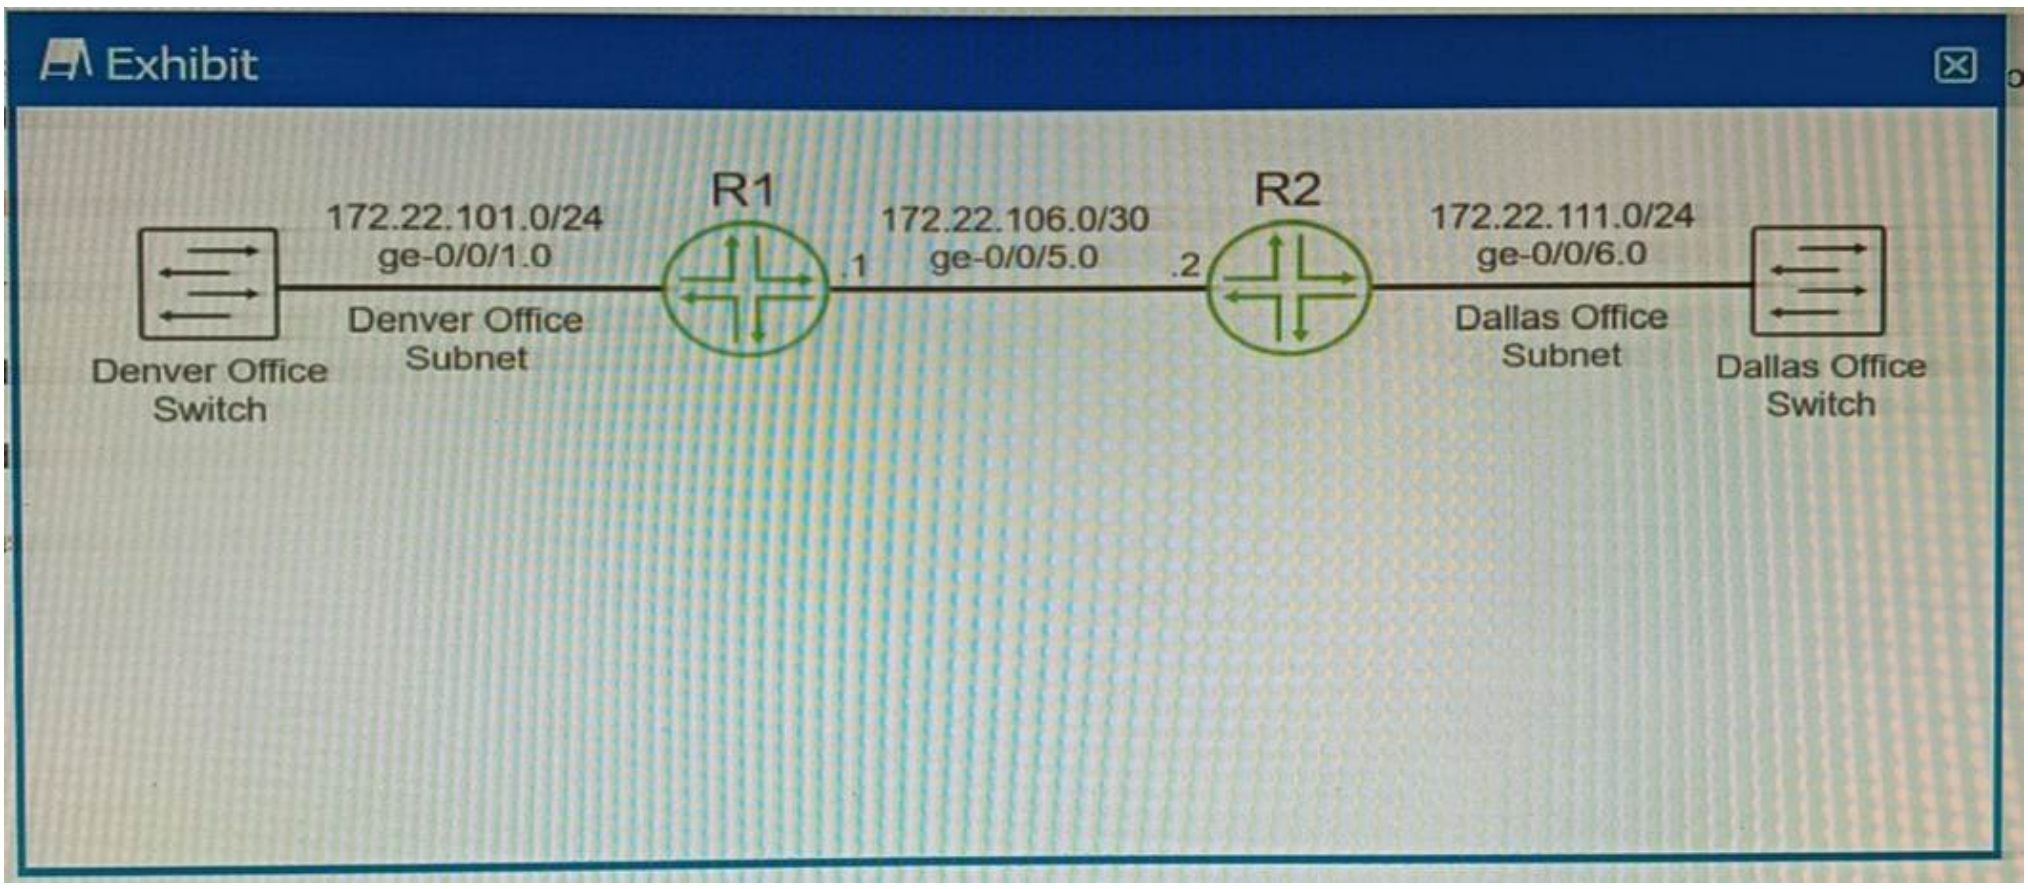

You are using OSPF to advertise the subnets that are used by the Denver and Dallas offices. The routers that are directly connected to the Dallas and Denver subnets are not advertising the connected subnets.
Referring to the exhibit, which two statements are correct? (Choose two.)

A. Create static routes on the switches using the local vMX router's loopback interface for the next hop.
B. Configure and apply a routing policy that redistributes the Dallas and Denver subnets using Type 5 LSAs.
C. Configure and apply a routing policy that redistributes the connected Dallas and Denver subnets.
D. Enable the passive option on the OSPF interfaces that are connected to the Dallas and Denver subnets.

Answer: CD

Explanation:
The routers that are directly connected to the Dallas and Denver subnets are not advertising the connected subnets. This can be resolved by redistributing the connected subnets into OSPF1.
Option C suggests to configure and apply a routing policy that redistributes the connected Dallas and Denver subnets. This is correct because redistribution allows routes from one routing protocol to be communicated to another, and in this case, it allows the connected subnets to be advertised through OSPF1.
Option D suggests enabling the passive option on the OSPF interfaces that are connected to the Dallas and Denver subnets. This is also correct because in OSPF, a passive interface is an interface that belongs to the OSPF router, but does not send OSPF Hello packets1. It??s typically used on an interface that you don??t want to use for OSPF adjacencies, but you still want to advertise its IP address1. Therefore, enabling passive interface can help in advertising the Dallas and Denver subnets.

NEW QUESTION 29
Which three protocols support BFD? (Choose three.)

A. RSTP
B. BGP
C. OSPF
D. LACP
E. FTP

Answer: BCD

Explanation:
BFD is a protocol that can be used to quickly detect failures in the forwarding path between two adjacent routers or switches. BFD can be integrated with various routing protocols and link aggregation protocols to provide faster convergence and fault recovery. According to the Juniper Networks documentation, the following protocols support BFD on Junos OS devices1:
? BGP: BFD can be used to monitor the connectivity between BGP peers and trigger a session reset if a failure is detected. BFD can be configured for both internal and external BGP sessions, as well as for IPv4 and IPv6 address families2.
? OSPF: BFD can be used to monitor the connectivity between OSPF neighbors and trigger a state change if a failure is detected. BFD can be configured for both OSPFv2 and OSPFv3 protocols, as well as for point-to-point and broadcast network types3.
? LACP: BFD can be used to monitor the connectivity between LACP members and trigger a link state change if a failure is detected. BFD can be configured for both active and passive LACP modes, as well as for static and dynamic LAGs4.
Other protocols that support BFD on Junos OS devices are:
? IS-IS: BFD can be used to monitor the connectivity between IS-IS neighbors and trigger a state change if a failure is detected. BFD can be configured for both level 1 and level 2 IS-IS adjacencies, as well as for point-to-point and broadcast network types.
? RIP: BFD can be used to monitor the connectivity between RIP neighbors and trigger a route update if a failure is detected. BFD can be configured for both RIP version 1 and version 2 protocols, as well as for IPv4 and IPv6 address families.
? VRRP: BFD can be used to monitor the connectivity between VRRP routers and trigger a priority change if a failure is detected. BFD can be configured for both VRRP version 2 and version 3 protocols, as well as for IPv4 and IPv6 address families.
The protocols that do not support BFD on Junos OS devices are:
? RSTP: RSTP is a spanning tree protocol that provides loop prevention and rapid convergence in layer 2 networks. RSTP does not use BFD to detect link failures, but relies on its own hello mechanism that sends BPDU packets every 2 seconds by default.
? FTP: FTP is an application layer protocol that is used to transfer files between hosts over a TCP connection. FTP does not use BFD to detect connection failures, but relies on TCP??s own retransmission and timeout mechanisms.
References:

1: [Configuring Bidirectional Forwarding Detection] 2: [Configuring Bidirectional Forwarding Detection for BGP] 3: [Configuring Bidirectional Forwarding Detection for OSPF] 4: [Configuring Bidirectional Forwarding Detection for Link Aggregation Control Protocol] : [Configuring Bidirectional Forwarding Detection for IS-IS] : [Configuring Bidirectional Forwarding Detection for RIP] : [Configuring Bidirectional Forwarding Detection for VRRP] : [Understanding Rapid Spanning Tree Protocol] : [Understanding FTP]

NEW QUESTION 31
What are two characteristics of RSTP alternate ports? (Choose two.)

A. RSTP alternate ports block traffic while receiving superior BPDUs from a neighboring switch.
B. RSTP alternate ports provide an alternate lower cost path to the root bridge.
C. RSTP alternate ports provide an alternate higher cost path to the root bridge.
D. RSTP alternate ports are active ports used to forward frames toward the root bridge.

Answer: AC

Explanation:
? A is correct because RSTP alternate ports block traffic while receiving superior BPDUs from a neighboring switch. An alternate port is a backup port for a root port, which means it receives better BPDUs from another bridge than the current root port1. However, an alternate port does not forward any traffic, as it is in a discarding state2. It only listens to BPDUs and waits for the root port to fail. If the root port fails, the alternate port can immediately transition to a forwarding state and become the new root port1.
? C is correct because RSTP alternate ports provide an alternate higher cost path to the root bridge. An alternate port is selected based on the same criteria as the root port, which are the lowest bridge ID, the lowest path cost, the lowest sender port ID, and the lowest receiver port ID3. However, an alternate port receives a higher cost BPDU than the root port, otherwise it would be the root port itself1. Therefore, an alternate port provides an alternate higher cost path to the root bridge than the root port.

NEW QUESTION 35
Exhibit

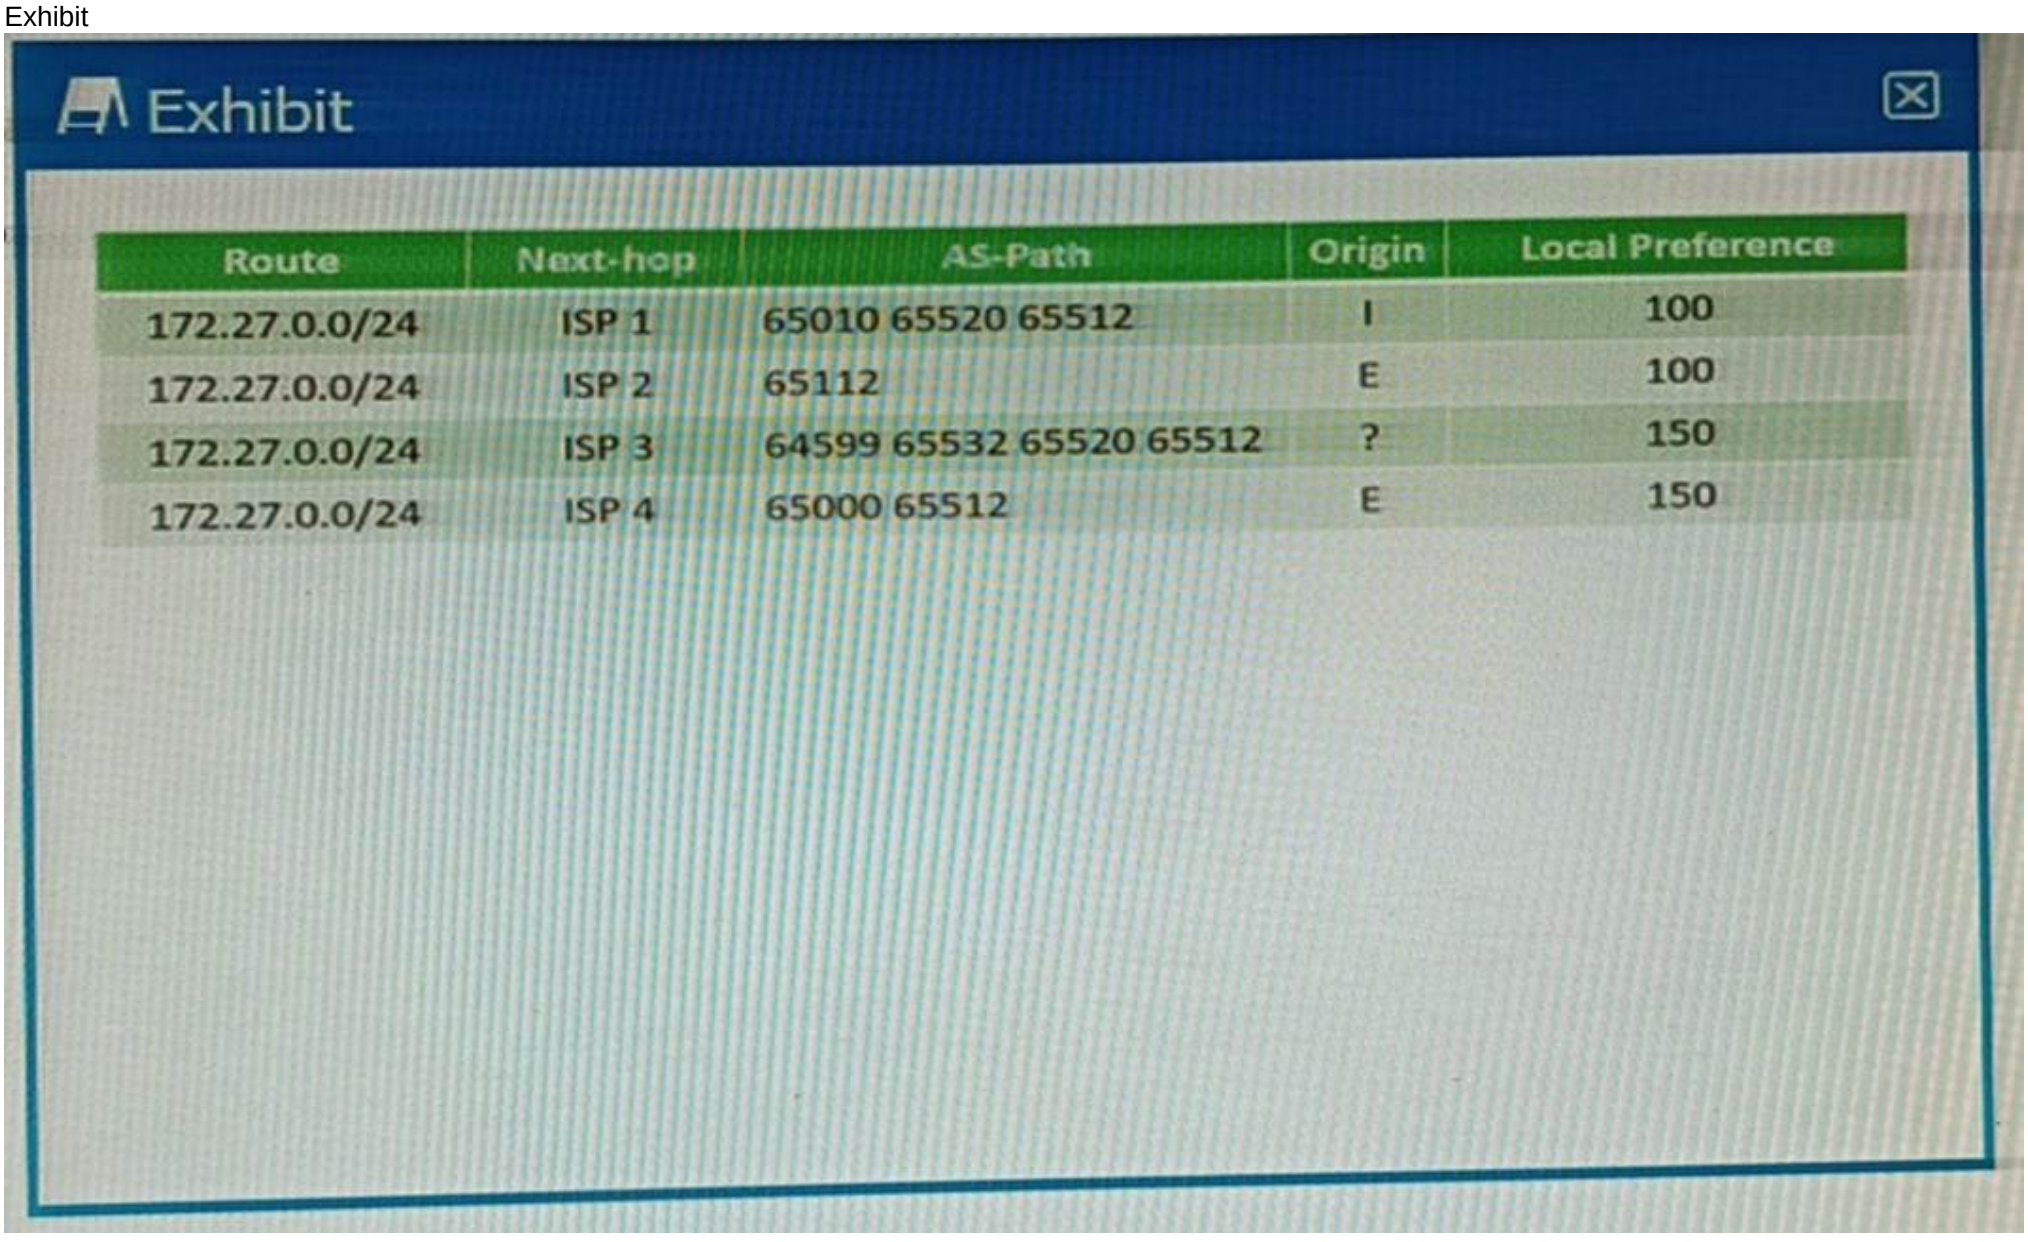

You are receiving the BGP route shown in the exhibit from four different upstream ISPs. Referring to the exhibit, which ISP will be selected as the active path?

A. ISP1
B. ISP 3
C. ISP 4
D. ISP 2

Answer: C

Explanation:
In BGP, the path selection process is based on a set of attributes1. The process starts by preferring the path with the highest weight, then the highest local preference, then the locally originated routes, and so on1. If all these attributes are the same, then it prefers the path with the shortest AS path1.
Referring to the exhibit, all four ISPs have the same weight, local preference, and origin1. However, ISP 4 has the shortest AS path1. Therefore, ISP 4 will be selected as the active path. So, option C is correct.

NEW QUESTION 38
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 39
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 44
.....

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