Exam Questions JN0-351

Enterprise Routing and Switching - Specialist (JNCIS-ENT)

https://www.2passeasy.com/dumps/JN0-351/

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)

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

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

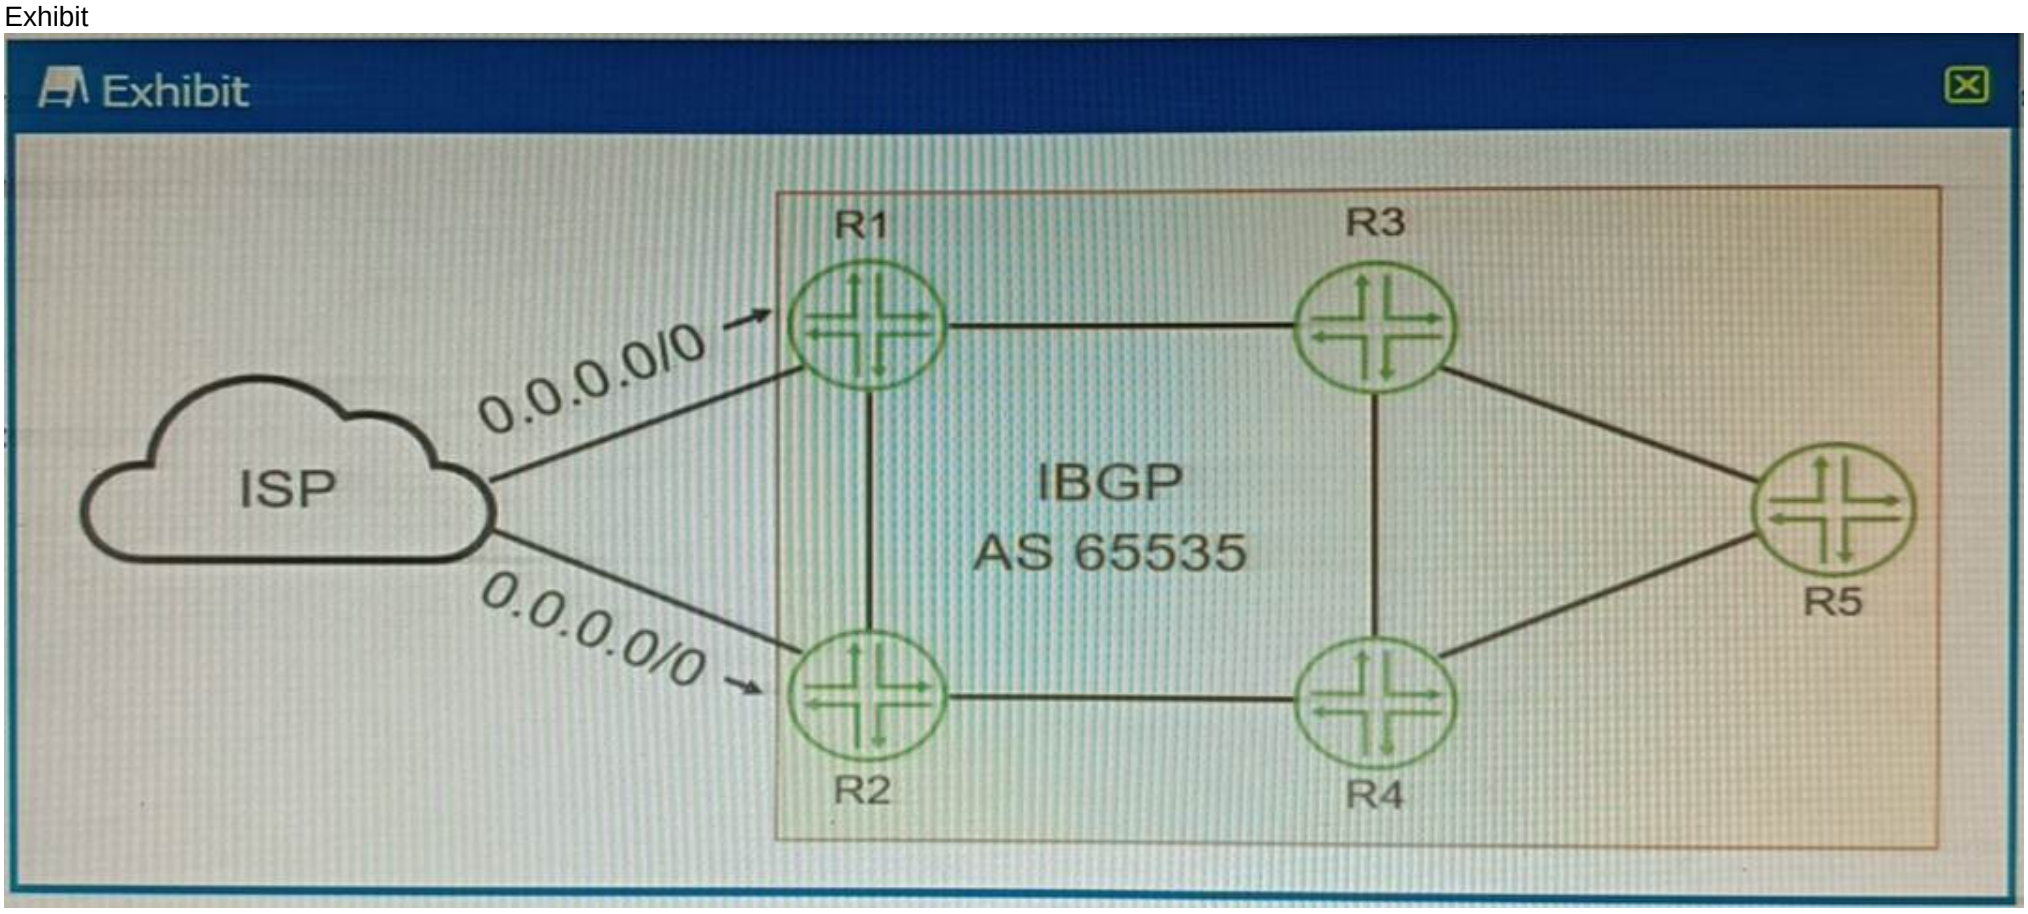

Your ISP is announcing a default route to both R1 and R2. You want your network routers to forward all Internet traffic through the R1 device
Which BGP attribute would you use?

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 3
Exhibit.

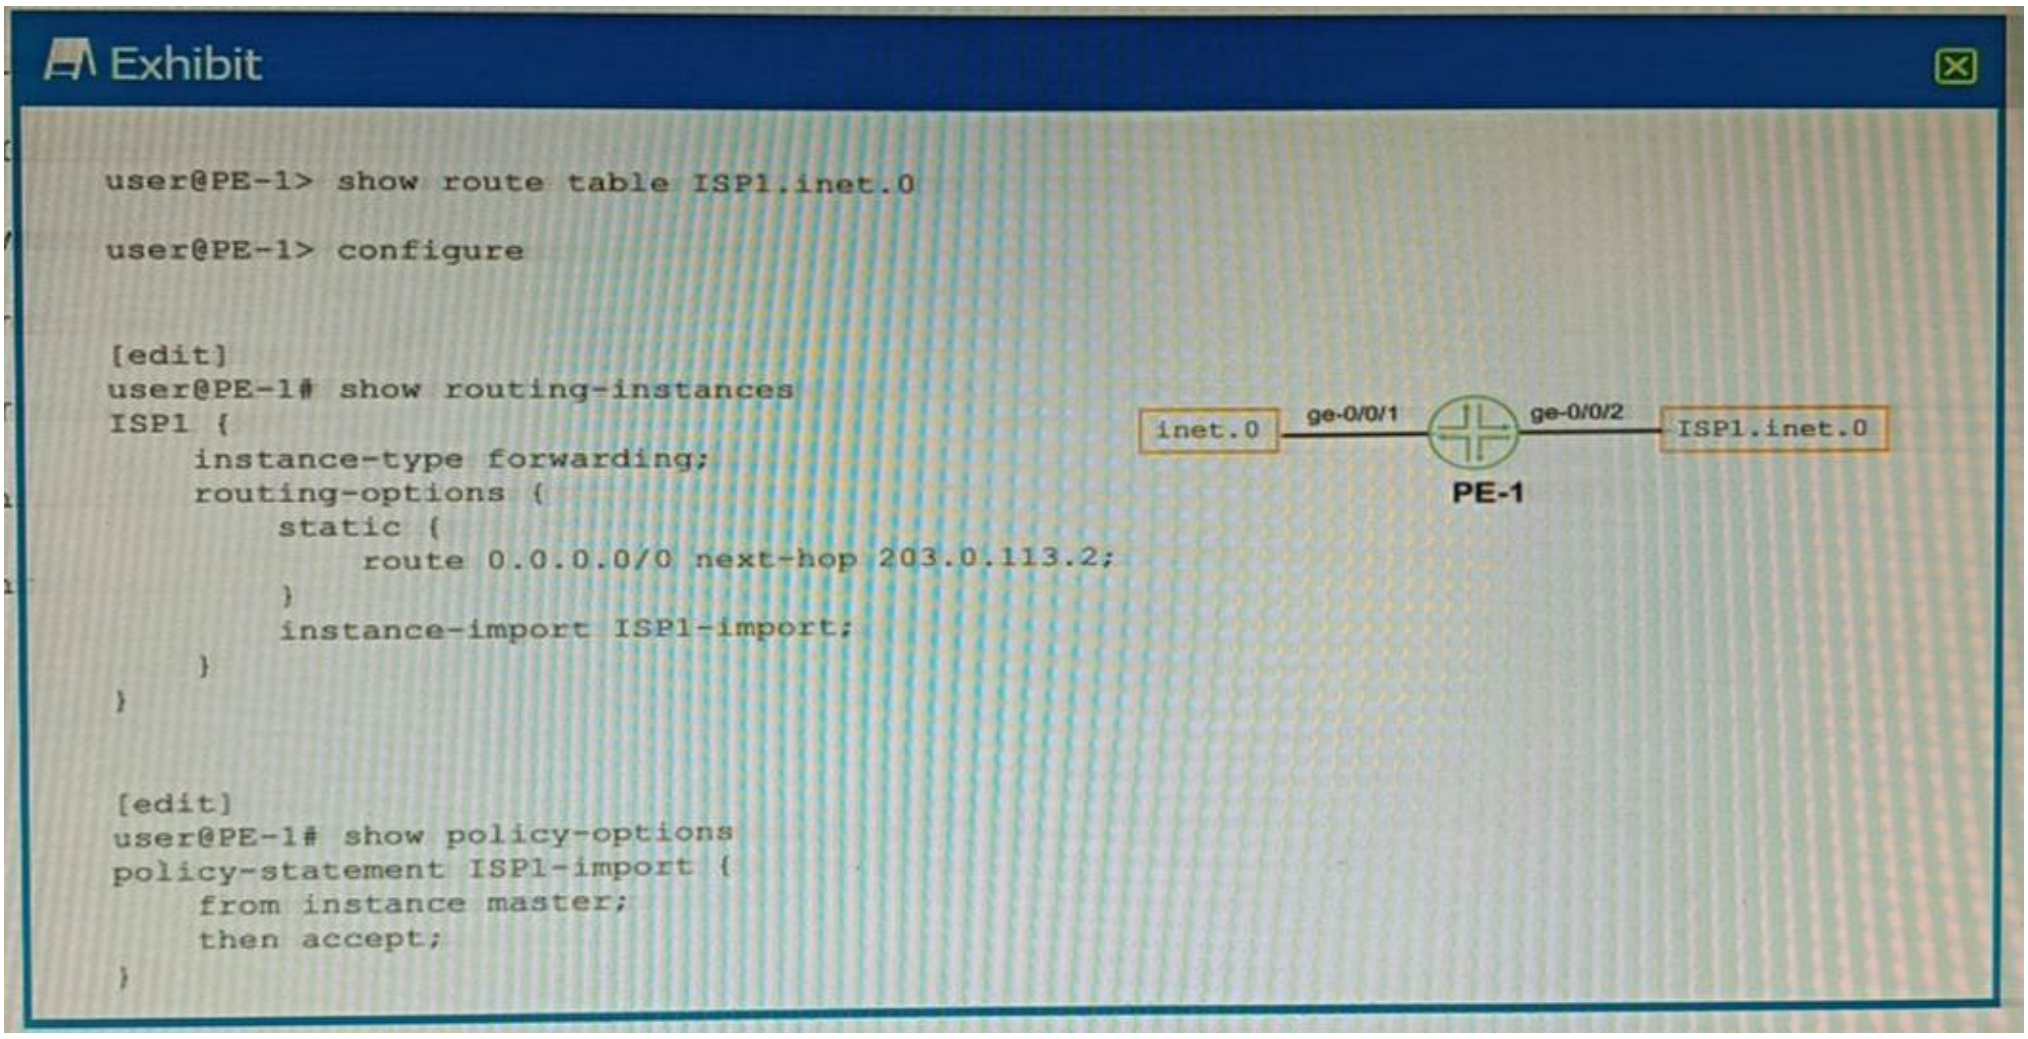

The ispi _ inet. 0 route table has currently no routes in it.
What will happen when you commit the configuration shown on the exhibit?

A. The ine
B. 0 route table will be completely overwritten by the ispi . ine
C. 0 route table.
D. The ine
E. 0 route table will be imported into the ispi . ine
F. 0 route table.
G. The ISPI . ine
H. 0 route table will be completely overwritten by the ine
I. o route table.
J. The ISPI . ine
K. 0 route table will be imported into the ine
L. 0 route table.

Answer: B

Explanation:
The configuration shown in the exhibit is an example of a routing instance of type virtual-router. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters that create a separate routing domain on a Juniper device1. A virtual-router routing instance allows administrators to divide a device into multiple independent virtual routers, each with its own routing table2.
The configuration also includes a rib-group statement, which is used to import routes from one routing table to another. A rib-group consists of an import-rib statement, which specifies the source routing table, and an export-rib statement, which specifies the destination routing table.
In this case, the rib-group name is inet-to-ispi, and the import-rib statement specifies inet.0 as the source routing table. The export-rib statement specifies ispi.inet.0 as the destination routing table. This means that the routes from inet.0 will be imported into ispi.inet.0. Therefore, the correct answer is B. The inet.0 route table will be imported into the ispi.inet.0 route table.
References:
1: Routing Instances Overview 2: Virtual Routing Instances : [rib-group (Routing Options)]

NEW QUESTION 4
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 5
You want to use filter-based forwarding (FBF) on your Internet peering router to load- balance traffic to two directly connected ISPs based on the source address.
Which two statements are correct in this scenario? (Choose two.)

A. FBF uses the no-forwarding routing instance type.
B. FBF uses the forwarding routing instance type.
C. RIB groups are used to copy routes from the ine
D. o routing table.
E. RIB groups are used to hide routes in the ine
F. 0 routing table.

Answer: BC

Explanation:
? Option B is correct. Filter-based forwarding (FBF), also known as Policy Based Routing (PBR), uses the forwarding routing instance type12.
? Option C is correct. Routing Information Base (RIB) groups are used to copy routes from one routing table to another34. In the context of FBF, RIB groups can be used to copy routes from the inet.0 routing table34.
? Option A is incorrect. FBF does not use the no-forwarding routing instance type15.
? Option D is incorrect. RIB groups are not used to hide routes in the inet.0 routing table34. They are used to share or copy routes between different routing tables34.

NEW QUESTION 6
Exhibit

Exhibit

```
user@R1> show bgp neighbor
Peer: 10.32.1.2+63645 AS 65401 Local: 10.32.1.1+179 AS 65400
  Description: EBGP peering to 10.32.1.2
  Group: IPCLOS_eBGP                Routing-Instance: master
  Forwarding routing-instance: master
  Type: External    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ IPCLOS_BGP_EXP ] Import: [ IPCLOS_BGP_IMP ]
  Options: <Preference PeerAS Multipath LocalAS Refresh>
  Options: <VpnApplyExport MtuDiscovery MultipathAs BfdEnabled>
  Holdtime: 90 Preference: 170 Local AS: 65400 Local System AS: 0
  Number of flaps: 0
  Peer ID: 10.52.100.2      Local ID: 10.52.100.1        Active Holdtime: 90
  Keepalive Interval: 30          Group index: 0     Peer index: 0     SNMP
index: 0
  I/O Session Thread: bgpio-0 State: Enabled
  BFD: enabled, up
  Local Interface: ge-0/0/1.0
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  Restart flag received from the peer: Notification
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer does not support LLGR Restarter functionality
  Peer supports 4 byte AS extension (peer-as 65401)
  Peer does not support Addpath
  Table inet.0 Bit: 20000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              6
    Received prefixes:            9
    Accepted prefixes:            9
    Suppressed due to damping:    0
    Advertised prefixes:          22
  Last traffic (seconds): Received 22   Sent 10   Checked 69617
  Input messages:  Total 2568 Updates 4 Refreshes 0 Octets 48991
  Output messages: Total 2572 Updates 8 Refreshes 0 Octets 49362
  Output Queue[1]: 0              (inet.0, inet-unicast)
```

You are a network operator troubleshooting BGP connectivity.
Which two statements are correct about the output shown in the exhibit? (Choose two.)

A. Peer 10.32.1.2 is configured for AS 63645.
B. The BGP session is not established.
C. The R1 is configured for AS 65400.
D. The routers are exchanging IPv4 routes.

Answer: BC

Explanation:
Option B suggests that the BGP session is not established. This is correct because in the output, the state of the BGP session is shown as ??Idle??. In BGP, an ??Idle?? state means that the BGP session is not currently established1.
Option C suggests that R1 is configured for AS 65400. This is also correct because in the output, it??s shown that the local AS number is 654001. The local AS number represents the Autonomous System (AS) number of the router on which you??re checking the BGP session1.

NEW QUESTION 7
Exhibit.

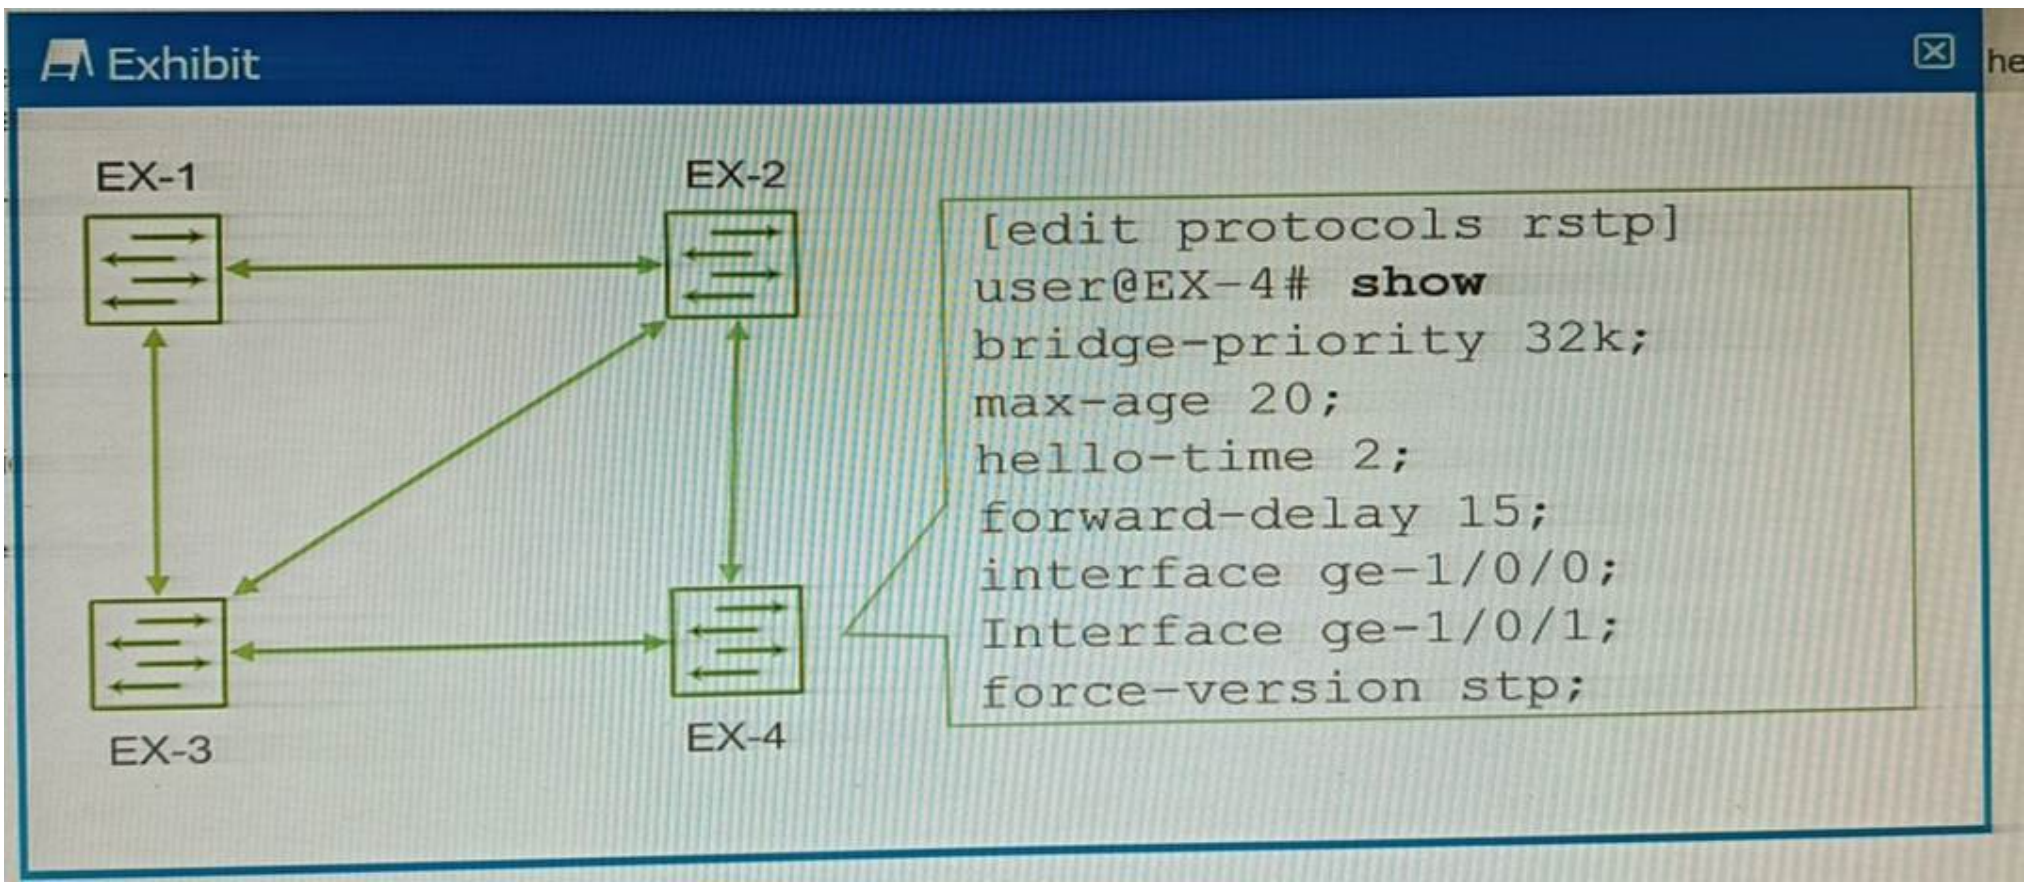

You have configured the four EX Series switches with RSTP, as shown in the exhibit. You discover that whenever a link between switches goes up or down, the switches take longer
than expected for RSTP to converge, using the default settings.
In this scenario, which action would solve the delay in RSTP 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 8
Exhibit.

Exhibit

```
{master:0}[edit]
user@switch# run show interfaces terse
Interface               Admin  Link  Proto      Local              Remote
ge-0/0/0                up     up
gr-0/0/0                up     up
pfe-0/0/0               up     up
ge-0/0/1                       up       up
ge-0/0/1.0                     up       up   inet     172.23.11.10/24
                                                     172.23.12.10/24
ge-0/0/2                       up       up
ge-0/0/2.0                     up       up   inet     172.23.11.100/24
ge-0/0/3                       up       up
ge-0/0/3.0                     up       up   inet     172.23.12.100/24
…
bme0                           up       up
bme0.0                         up       up   inet     128.0.0.1/2
                                                     128.0.0.4/2
                                                     128.0.0.16/2
                                                     128.0.0.63/2
…
jsrv.1                  up     up            inet     128.0.0.127/2
lo0                     up     up
lo0.16385               up     up            inet
lsi                     up     up
me0                     up     up
me0.0                   up     up            inet     10.210.20.233/29
mtun                    up     up
pimd                    up     up
pime                    up     up
tap                     up     up
vme                     up     down
```

What is the management IP address of the device shown in the exhibit?

A. 10.210.20.233
B. 172.23.12.100
C. 128.0.0.1
D. 172.23.11.10

Answer: B

Explanation:
The management IP address of a device is the IP address that is used to access the device for configuration and monitoring purposes. It is usually assigned to a dedicatedmanagement interface that is separate from the data interfaces. The management interface can be accessed via SSH, Telnet, HTTP, or other protocols. In the exhibit, the list of interfaces and their statuses shows that the management interface isme0. This interface has an admin status ofup, a protocol status ofinet, a local address of172.23.12.100/24, and a remote address ofunspecified. This means that the me0 interface is active, has an IPv4 address assigned, and is not connected to another device. Therefore, the management IP address of the device shown in the exhibit is172.23.12.100. References:
[Management Interfaces Overview] : [Displaying Interface Status Information]

NEW QUESTION 9
Which two statements are correct about tunnels? (Choose two.)

A. BFD cannot be used to monitor tunnels.
B. Tunnel endpoints must have a valid route to the remote tunnel endpoint.
C. IP-IP tunnels are stateful.
D. Tunnels add additional overhead to packet size.

Answer: BD

Explanation:
A tunnel is a connection between two computer networks, in which data is sent from one network to another through an encrypted link. Tunnels are commonly used to secure data communications between two networks or to connect two networks that use different protocols.
Option B is correct, because tunnel endpoints must have a valid route to the remote tunnel endpoint. A tunnel endpoint is the device that initiates or terminates a tunnel connection. For a tunnel to be established, both endpoints must be able to reach each other over the underlying network. This means that they must have a valid route to the IP address of the remote endpoint1.
Option D is correct, because tunnels add additional overhead to packet size. Tunnels work by encapsulating packets: wrapping packets inside of other packets. This means that the original packet becomes the payload of the surrounding packet, and the surrounding packet has its own header and trailer. The header and trailer of the surrounding packet add extra bytes to the packet size, which is called overhead. Overhead can reduce the efficiency and performance of a network, as it consumes more bandwidth and processing power2.
Option A is incorrect, because BFD can be used to monitor tunnels. 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. BFD can also be used to monitor the connectivity of tunnels, such as GRE, IPsec, or MPLS.
Option C is incorrect, because IP-IP tunnels are stateless. 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 are stateless, which means that they do not keep track of the state or status of the tunnel connection. Stateless tunnels do not require any signaling or negotiation between the endpoints, but they also do not provide any error detection or recovery mechanisms.
References:
1: What is Tunneling? | Tunneling in Networking 2: What Is Tunnel In Networking, Its Types, And Its Benefits? : [Configuring Bidirectional Forwarding Detection] : [IP-IP Tunneling]

NEW QUESTION 10
After receiving a BGP route, which two conditions are verified by the receiving router to ensure that the received route is valid? (Choose two)

A. The AS-path length is greater than 0.
B. The loops do not exist.
C. The next hop is reachable.
D. The local preference is greater than 0.

Answer: BC

Explanation:
? B is correct because the loops do not exist is one of the conditions that are verified by the receiving router to ensure that the received BGP route is valid. A loop in BGP means that a route has been advertised by the same AS more than once, which can cause routing instability and inefficiency1. To prevent loops, BGP uses the AS-path attribute, which lists the AS numbers that a route has traversed from the origin to the destination2. The receiving router checks the AS-path attribute of the received route and discards it if it finds its own AS number in the list2. This way, BGP avoids accepting routes that contain loops.
? C is correct because the next hop is reachable is one of the conditions that are verified by the receiving router to ensure that the received BGP route is valid. The next hop is the IP address of the next router that is used to forward packets to the destination network3. The receiving router checks the next hop attribute of the received route and verifies that it has a valid route to reach it3. If the next hop is not reachable, the received route is not usable and is rejected by the receiving router3. This way, BGP ensures that only feasible routes are accepted.

NEW QUESTION 10
Exhibit

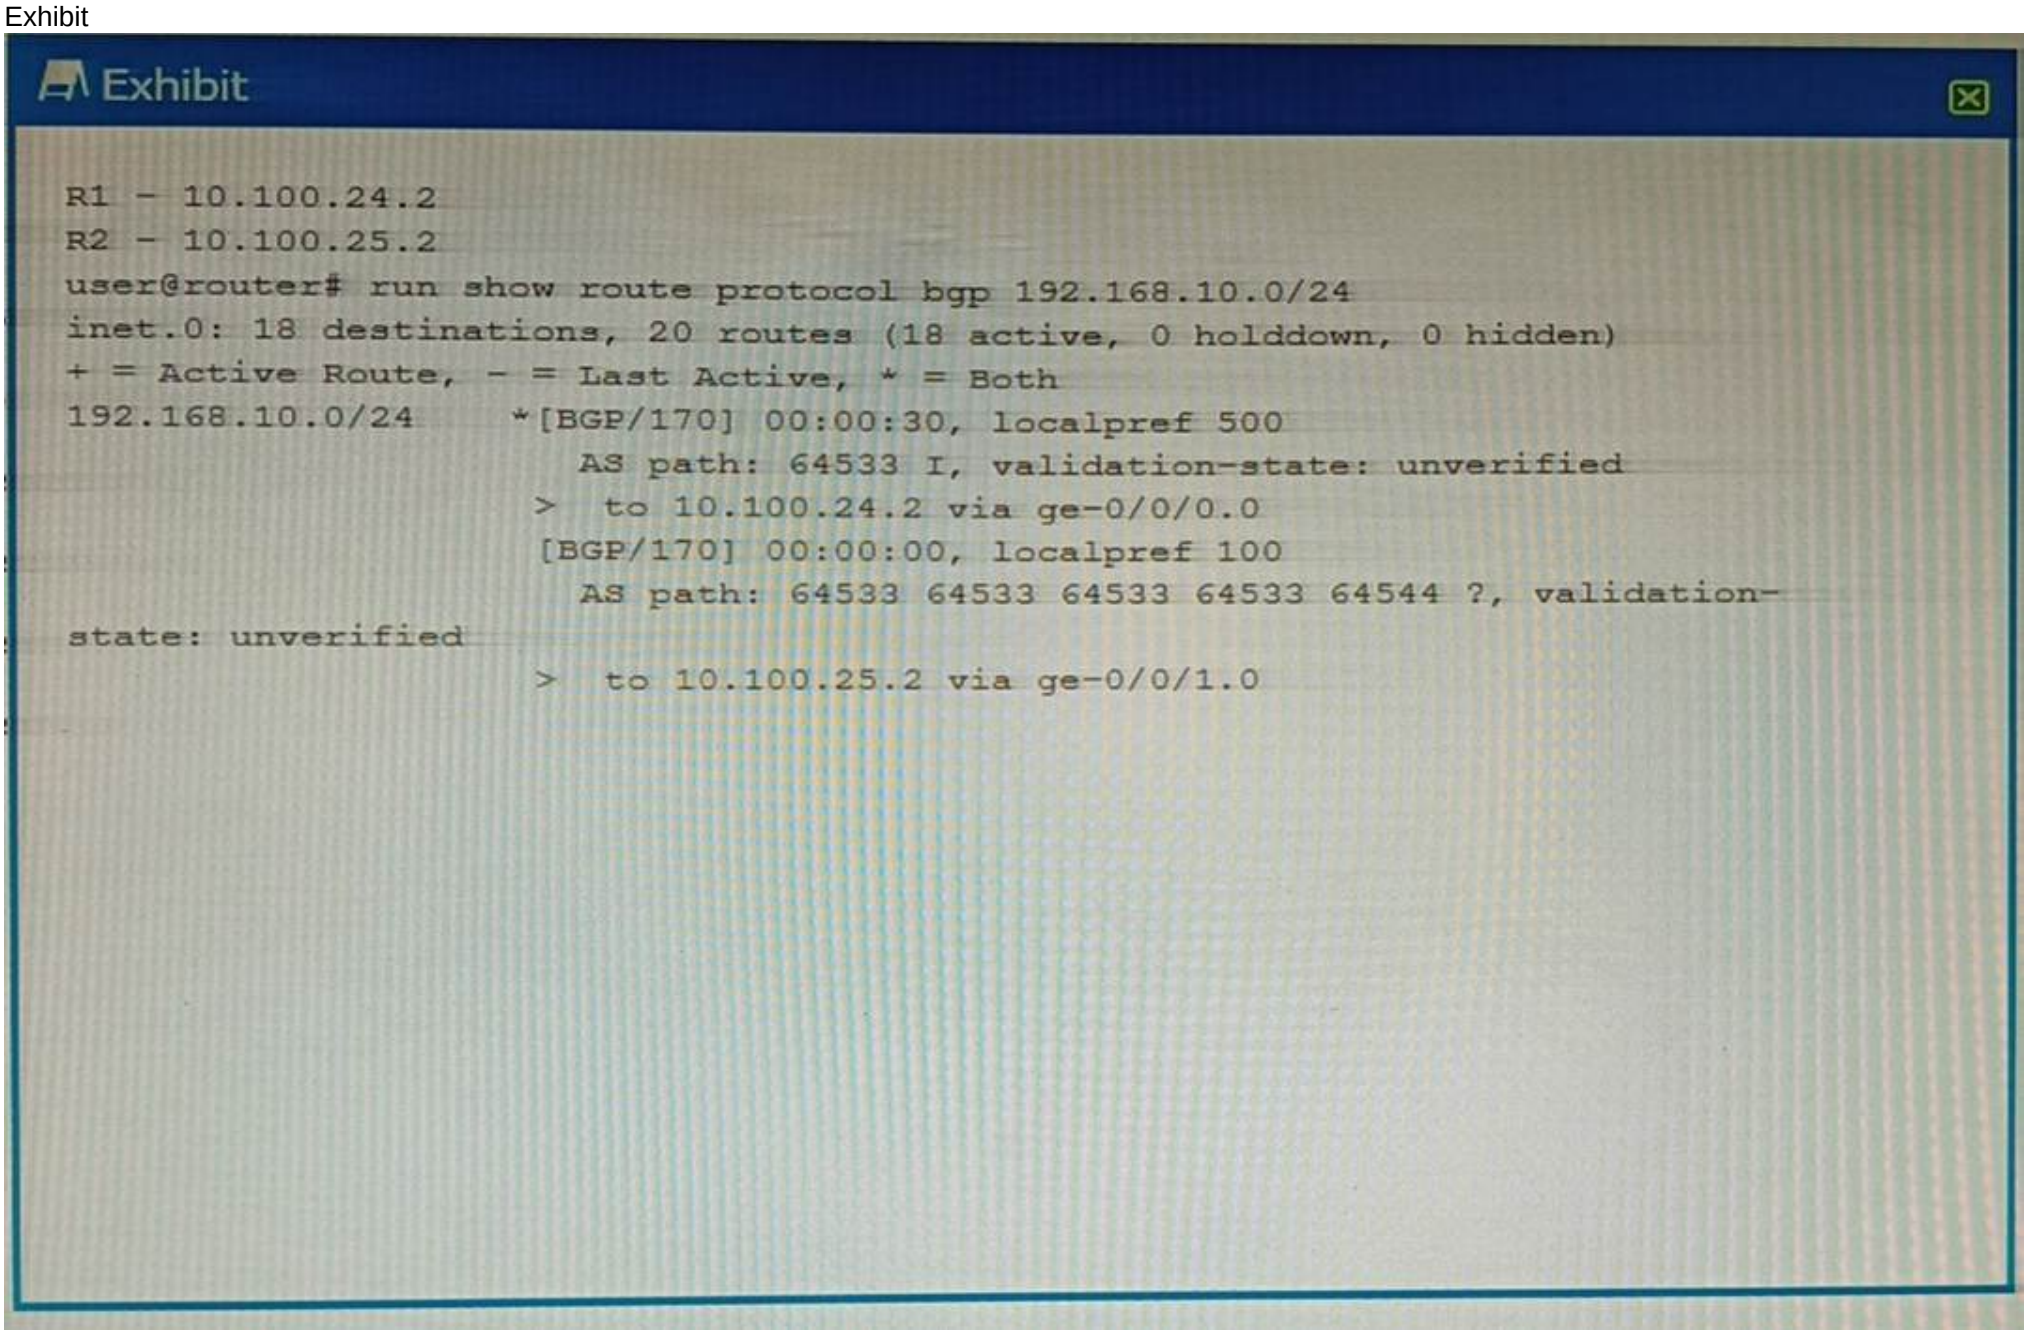

You are troubleshooting an issue where traffic to 192.168.10.0/24 is being sent to R1 instead of your desired path through R2.
Referring to the exhibit, what is the reason for the problem?

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 11
Exhibit.

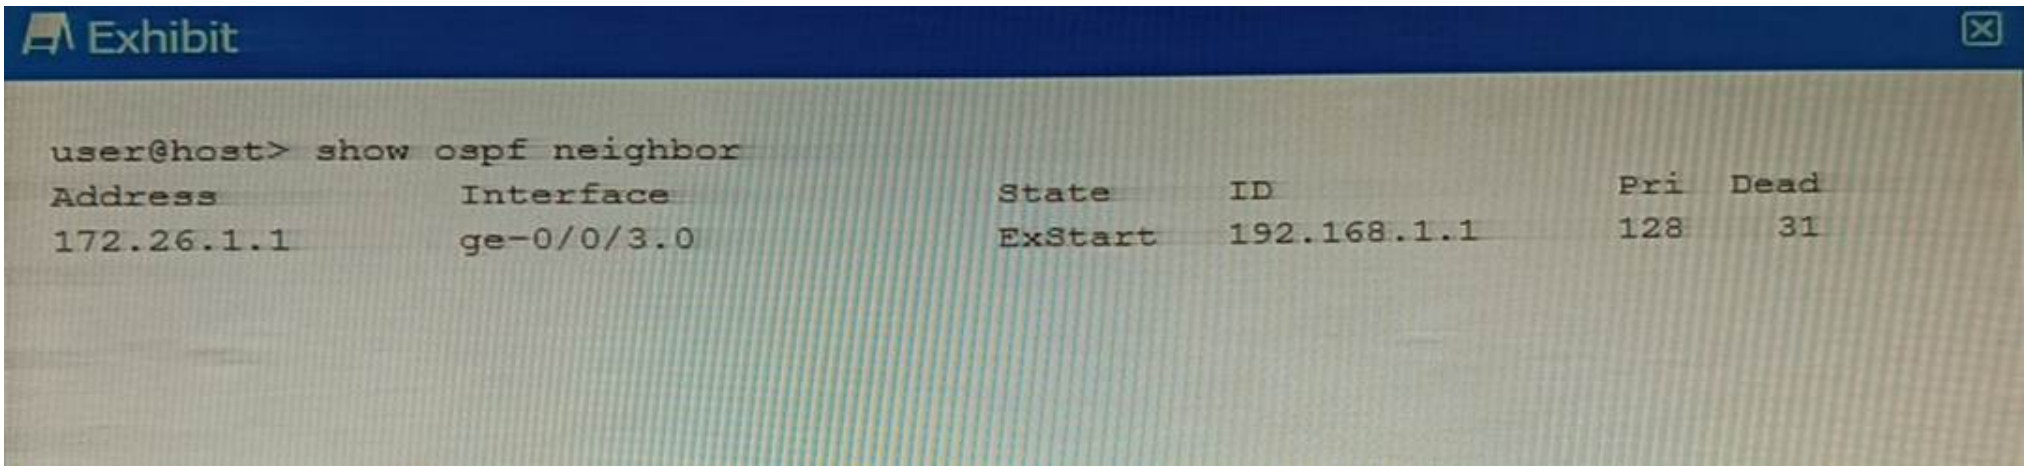

Why is this OSPF adjacency remaining in this state?

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 12
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 15
Refer to the exhibit.

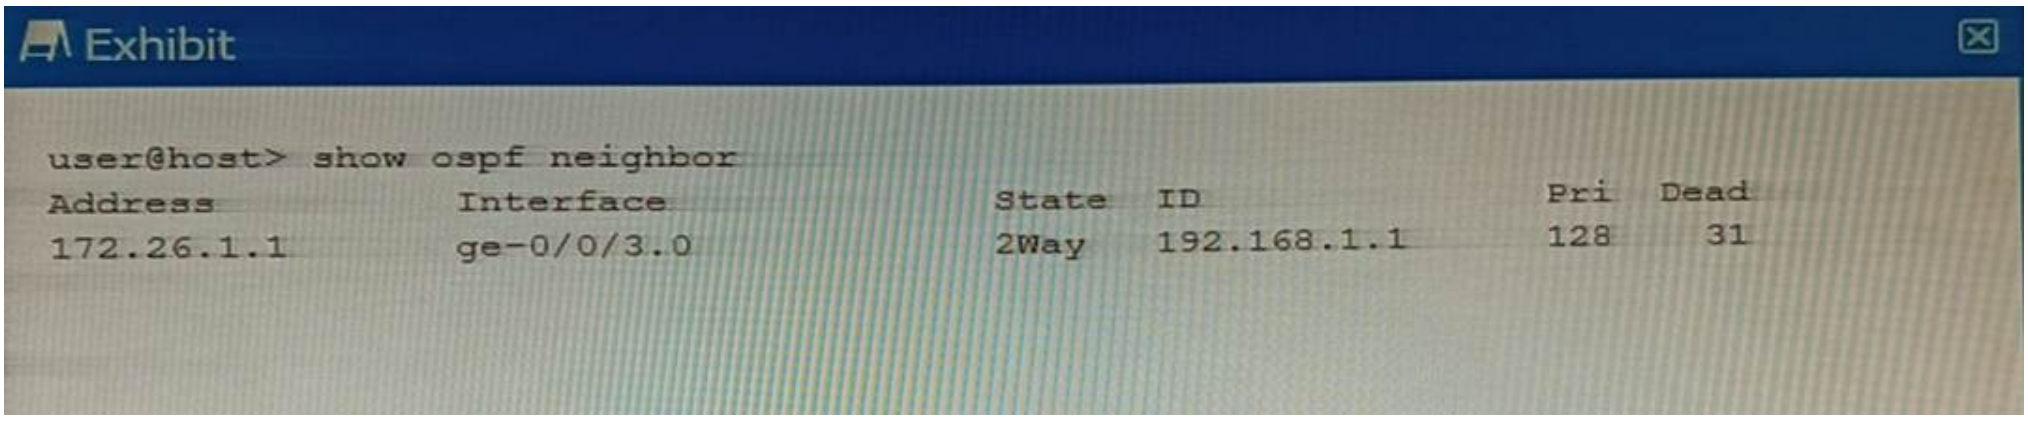

Referring to the output shown in the exhibit, which statement is correct?

A. The state is normal for a DR neighbor.
B. The state is normal for a DRother neighbor
C. An MTU mismatch exists between the OSPF neighbors.
D. An area ID mismatch exists between the OSPF neighbors

Answer: B

Explanation:
In OSPF, the state of the neighbor relationship is determined by the exchange of OSPF packets between routers1. The state ??2Way?? as shown in the exhibit indicates that bi-directional communication has been established between the two OSPF routers1. This is the normal state for a neighbor that is not the Designated Router (DR) or Backup Designated Router (BDR) on a broadcast, non-broadcast multi-access (NBMA), or point-to-multipoint network1. These neighbors are often referred to as "DRothers"1. Therefore, option B is correct.

NEW QUESTION 20
What is the maximum allowable MTU size for a default GRE tunnel without IPv4 traffic fragmentation?

A. 1496 bytes
B. 1480 bytes
C. 1500 bytes
D. 1476 bytes

Answer: D

Explanation:
The maximum allowable MTU size for a default GRE tunnel without IPv4 traffic fragmentation is 1476 bytes1. This is because GRE packets are formed by the addition of the original packets and the required GRE headers1. These headers are 24- bytes in length and since these headers are added to the original frame, depending on the original size of the packet we may run into IP MTU problems1. The most common IP MTU is 1500-bytes in length (Ethernet)1. When the tunnel is created, it deducts the 24-bytes it needs to encapsulate the passenger protocols and that is the IP MTU it will use1. For
example, if we are forming a tunnel over FastEthernet (IP MTU 1500)the IOS calculates the IP MTU on the tunnel as: 1500-bytes from Ethernet - 24-bytes for the GRE encapsulation = 1476-Bytes1.

NEW QUESTION 22
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 27
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 28
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 29
Which two mechanisms are part of building and maintaining a Layer 2 bridge table? (Choose two.)

A. blocking
B. flooding
C. learning
D. listening

Answer: BC

Explanation:
? Option B is correct. Flooding is a mechanism used in Layer 2 bridging where the switch sends incoming packets to all its ports except for the port where the packet originated1. This is done when the switch doesn??t know the destination MAC address or when the packet is a broadcast or multicast1.
? Option C is correct. Learning is another mechanism used in Layer 2 bridging where the switch learns the source MAC addresses of incoming packets and associates them with the port on which they were received23. This information is stored in a MAC address table, also known as a bridge table23.
? Option A is incorrect. Blocking is a state in Spanning Tree Protocol (STP) used to prevent loops in a network2. It??s not a mechanism used in building and maintaining a Layer 2 bridge table2.
? Option D is incorrect. Listening is also a state in Spanning Tree Protocol (STP) where the switch listens for BPDUs to make sure no loops occur in the network before transitioning to the learning state2. It??s not a mechanism used in building and maintaining a Layer 2 bridge table2.

NEW QUESTION 32
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 34
A new network requires multiple topology support. You decide to use IS-IS in this situation. Which three protocol topologies are supported in this scenario? (Choose three.)

A. IPsec
B. anycast
C. IPv6
D. multicast
E. IPv4

Answer: CDE

Explanation:
IS-IS (Intermediate System to Intermediate System) is a routing protocol that is designed to move information efficiently within a computer network12. It supports multiple protocol topologies, including IPv4, IPv6, and multicast12. Therefore, options C, E, and D are correct.

NEW QUESTION 38
You are concerned about spoofed MAC addresses on your LAN.
Which two Layer 2 security features should you enable to minimize this concern? (Choose two.)

A. dynamic ARP inspection
B. IP source guard
C. DHCP snooping
D. static ARP

Answer: AC

Explanation:
? A is correct because dynamic ARP inspection (DAI) is a Layer 2 security feature that prevents ARP spoofing attacks. ARP spoofing is a technique that allows an attacker to send fake ARP messages to associate a spoofed MAC address with a legitimate IP address. This can result in traffic redirection, man-in-the-middle attacks, or denial-of-service attacks. DAI validates ARP packets by checking the source MAC address and IP address against a trusted database, which is usually built by DHCP snooping1. DAI discards any ARP packets that do not match the database or have invalid formats1.
? C is correct because DHCP snooping is a Layer 2 security feature that prevents DHCP spoofing attacks. DHCP spoofing is a technique that allows an attacker to act as a rogue DHCP server and offer fake IP addresses and other network parameters to unsuspecting clients. This can result in traffic redirection, man-in- the-middle attacks, or denial-of-service attacks. DHCP snooping filters DHCP messages by classifying switch ports as trusted or untrusted. Trusted ports are allowed to send and receive any DHCP messages, while untrusted ports are allowed to send only DHCP requests and receive only valid DHCP replies from trusted ports2. DHCP snooping also builds a database of MAC addresses, IP addresses, lease times, and binding types for each client2.

NEW QUESTION 43
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 44
Exhibit

Exhibit

```
{master:0}
user@switch> show vlans brief
Routing instance        VLAN name               Tag         Interfaces
default-switch          default                 1
                                                            ge-0/0/0.0*
                                                            ge-0/0/1.0*
                                                            ge-0/0/2.0*
                                                            ge-0/0/3.0*
                                                            ge-0/0/4.0*
                                                            ge-0/0/5.0*
```

What does the * indicate in the output shown in the exhibit?

A. The switch ports have a router attached.
B. The interface is down.
C. The interface is active.
D. All interfaces have elected a root bridge.

Answer: C

Explanation:
? The exhibit shows the output of the command show vlans brief, which displays brief information about VLANs and their associated interfaces1.
? The output has four columns: Routing instance, VLAN name, Interfaces, and Tagging.
? The * symbol indicates that the interface is active, meaning that it is up and forwarding traffic1. This can be verified by the command show interfaces terse, which displays the status of the interfaces2.

NEW QUESTION 46
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 50
What are two reasons for creating multiple areas in OSPF? (Choose two.)

A. to reduce the convergence time
B. to increase the number of adjacencies in the backbone
C. to increase the size of the LSDB
D. to reduce LSA flooding across the network

Answer: AD

Explanation:
Option A is correct. Creating multiple areas in OSPF can help to reduce the convergence time . This is because changes in one area do not affect other areas, so fewer routers need to run the SPF algorithm in response to a change.
Option D is correct. Creating multiple areas in OSPF can help to reduce Link State Advertisement (LSA) flooding across the network. This is because LSAs are not flooded out of their area of origin.

NEW QUESTION 55

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 56
Exhibit

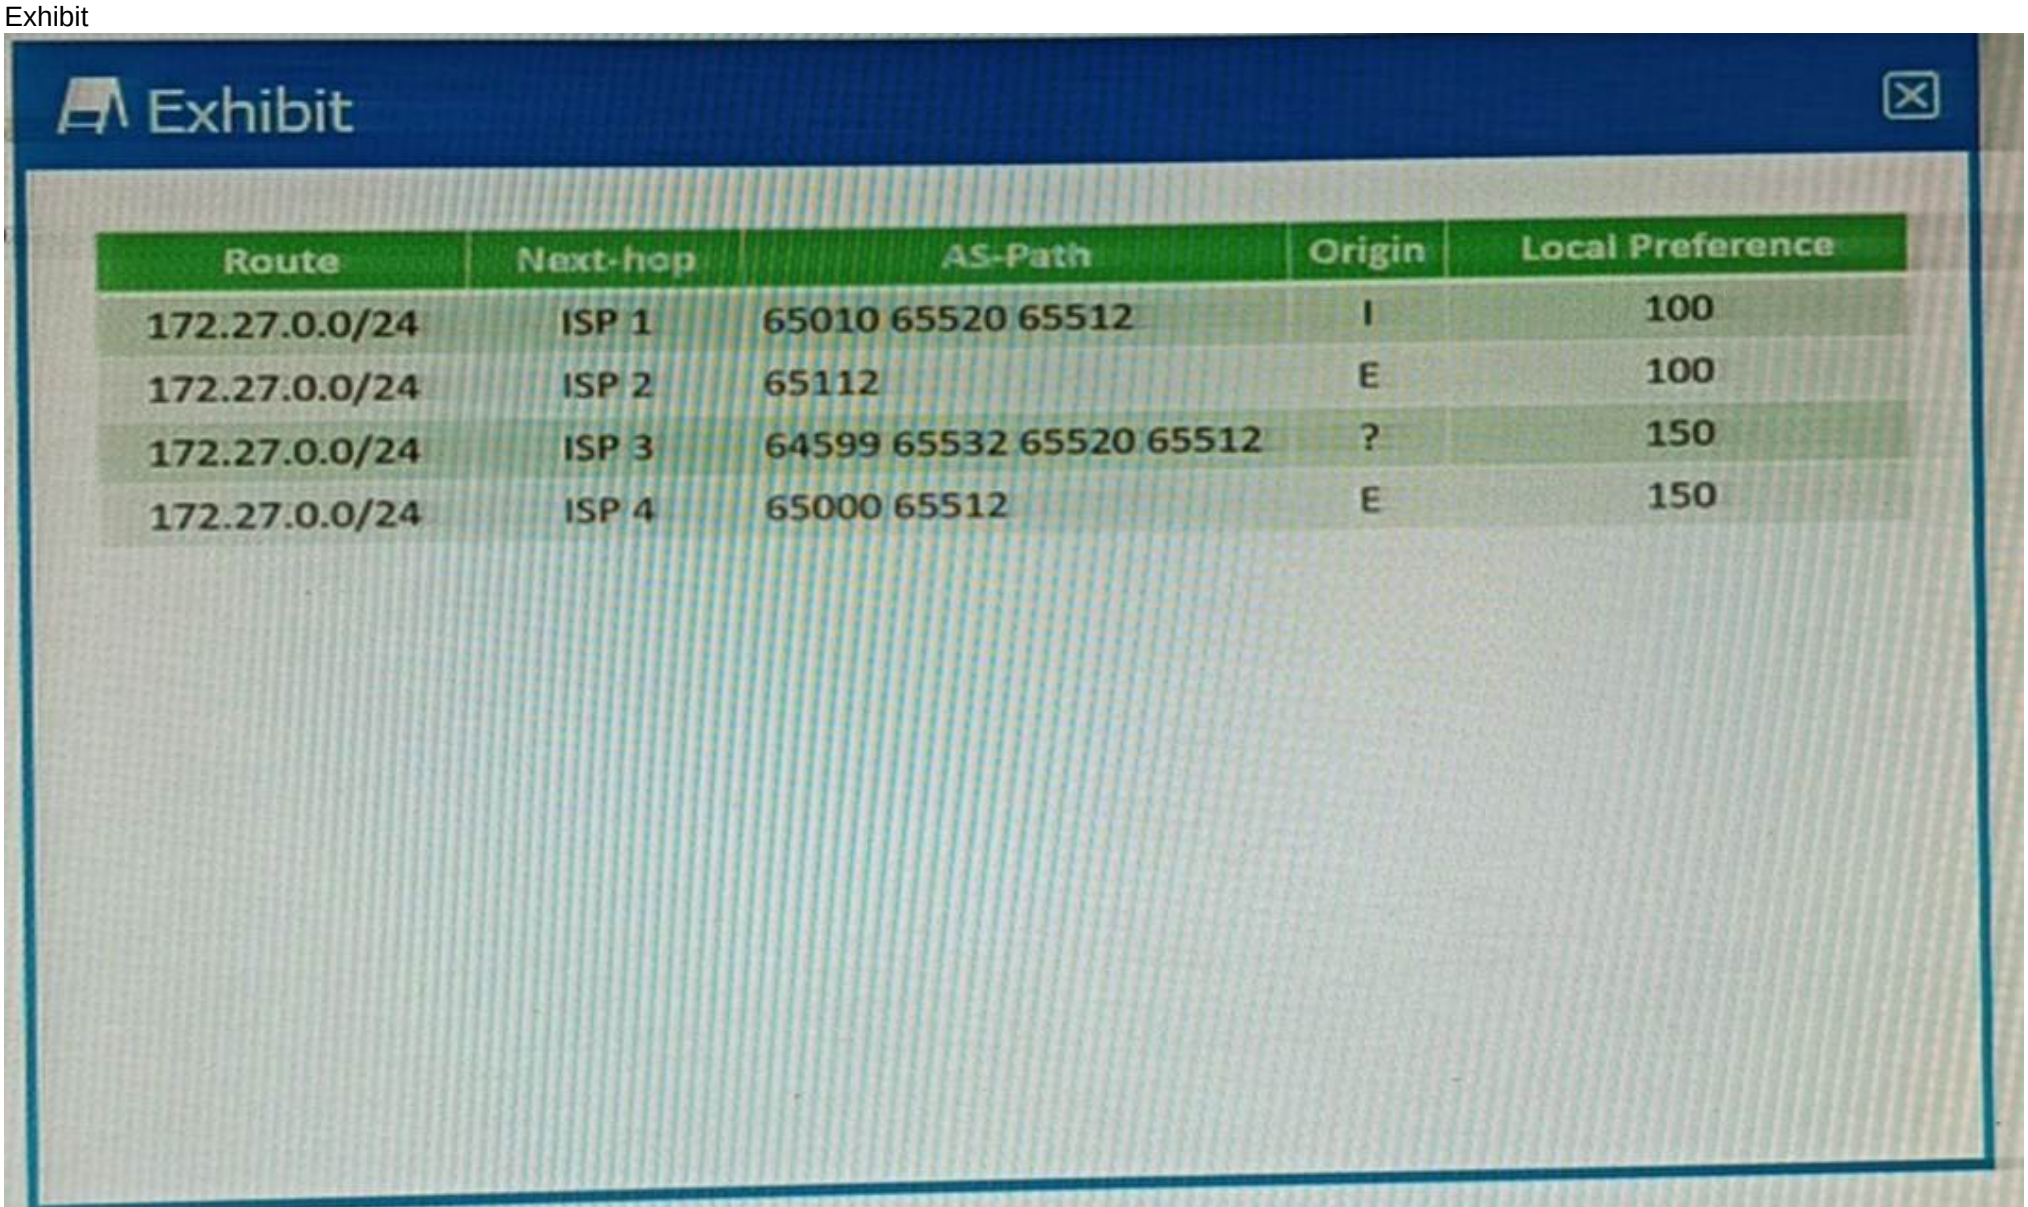

Route	Next-hop	AS-Path	Origin	Local Preference
172.27.0.0/24	ISP 1	65010 65520 65512	I	100
172.27.0.0/24	ISP 2	65112	E	100
172.27.0.0/24	ISP 3	64599 65532 65520 65512	?	150
172.27.0.0/24	ISP 4	65000 65512	E	150

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 60
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 61
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 62
Exhibit

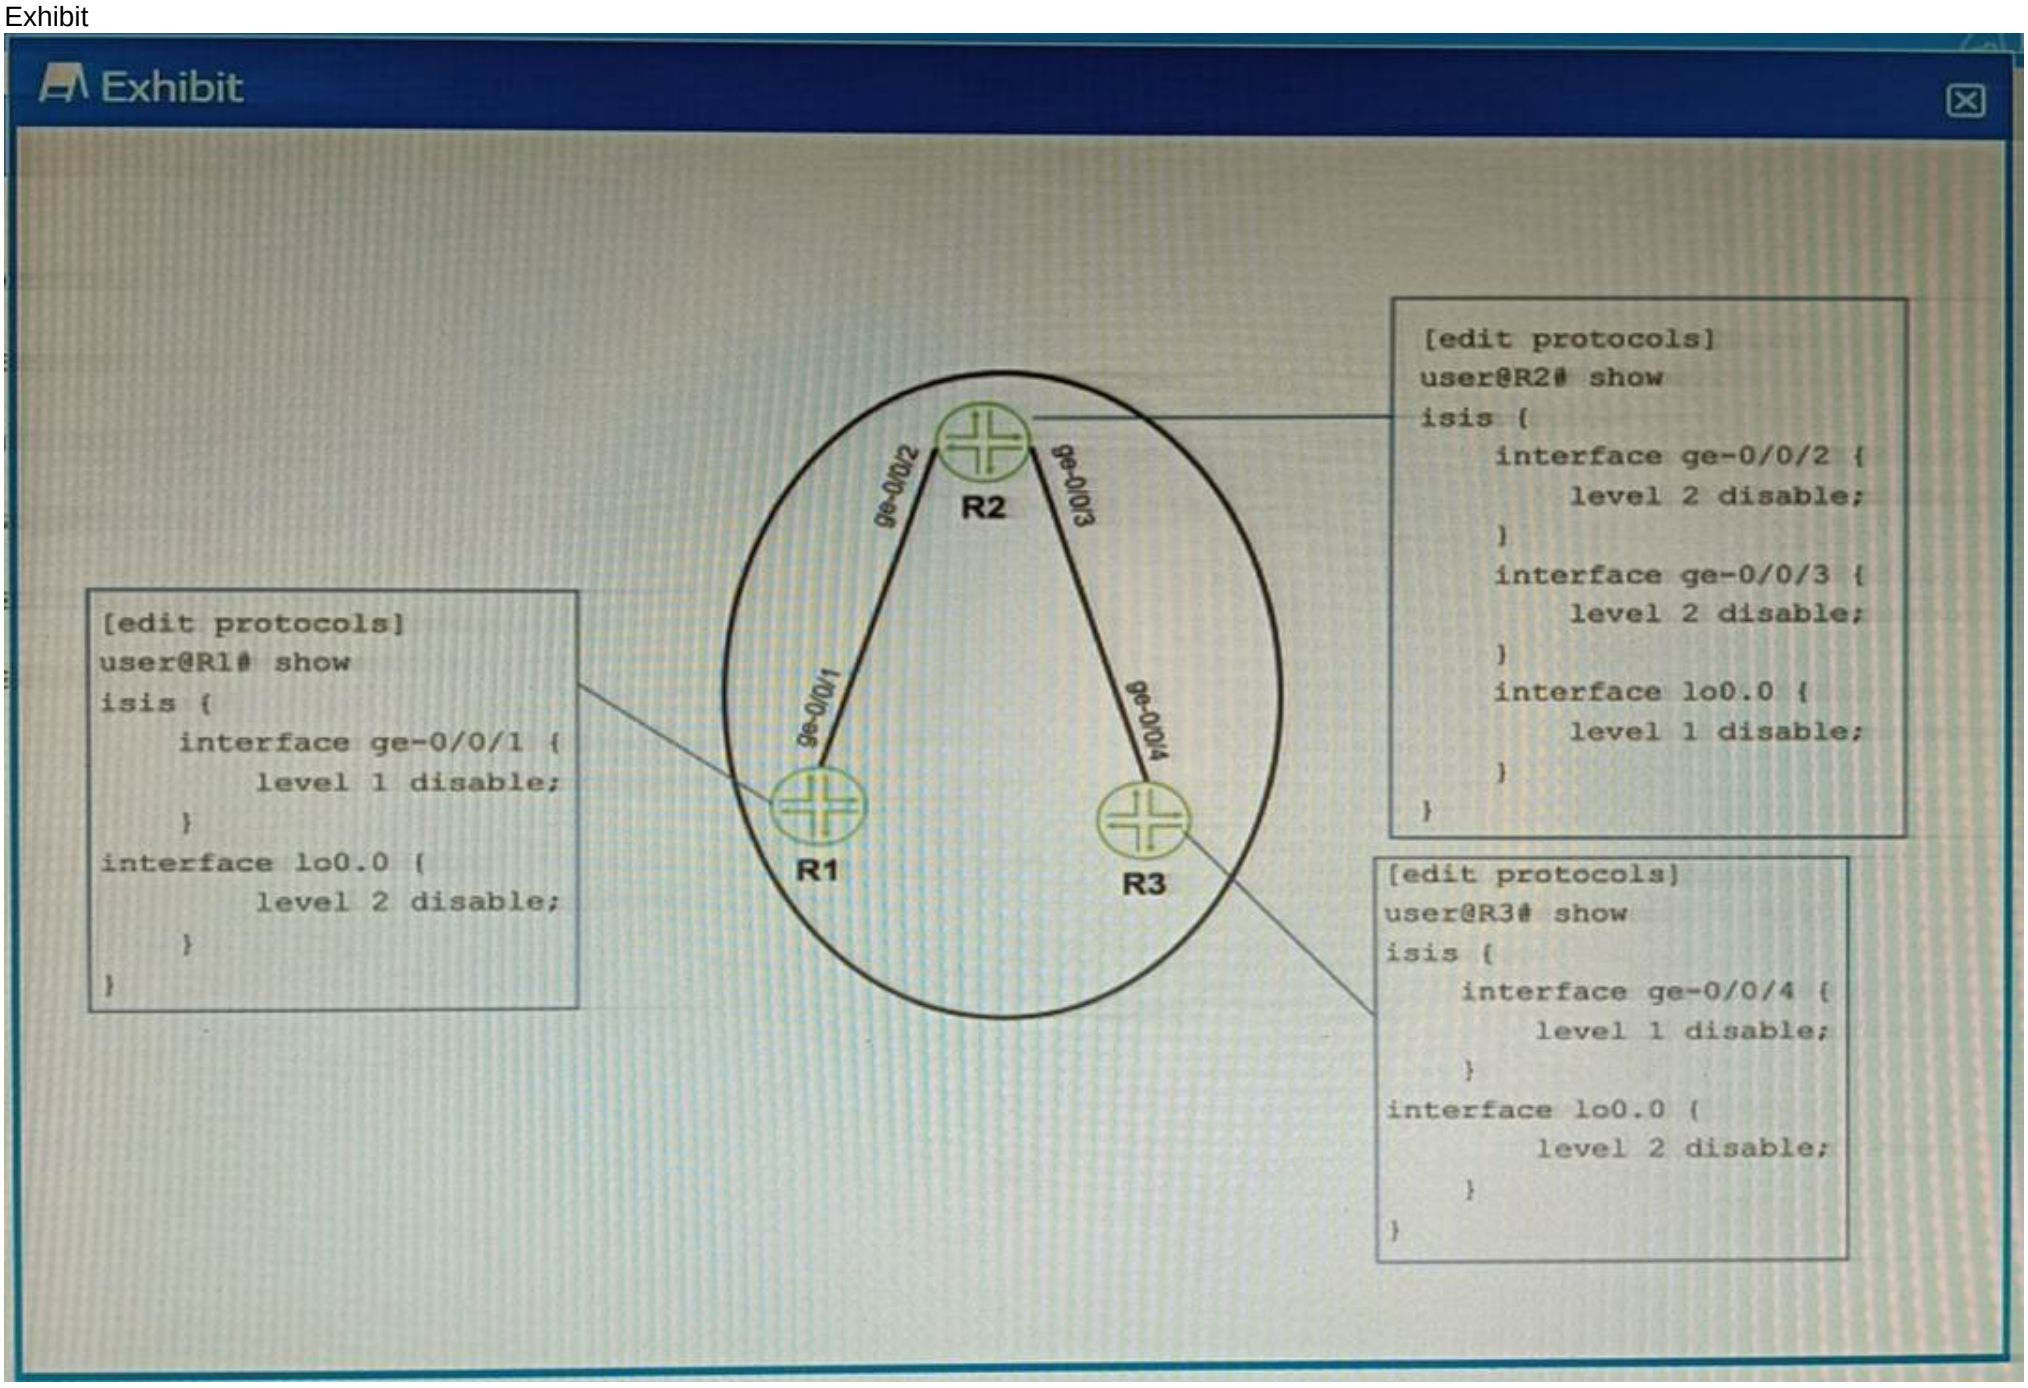

Referring to the exhibit, which two configuration changes must you apply for packets to reach from R1 to R3 using IS-IS? (Choose two.)

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 66
.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual JN0-351 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the JN0-351 Product From:

https://www.2passeasy.com/dumps/JN0-351/

Money Back Guarantee

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