Juniper

Exam Questions JN0-351

Enterprise Routing and Switching - Specialist (JNCIS-ENT)

NEW QUESTION 1
Exhibit.

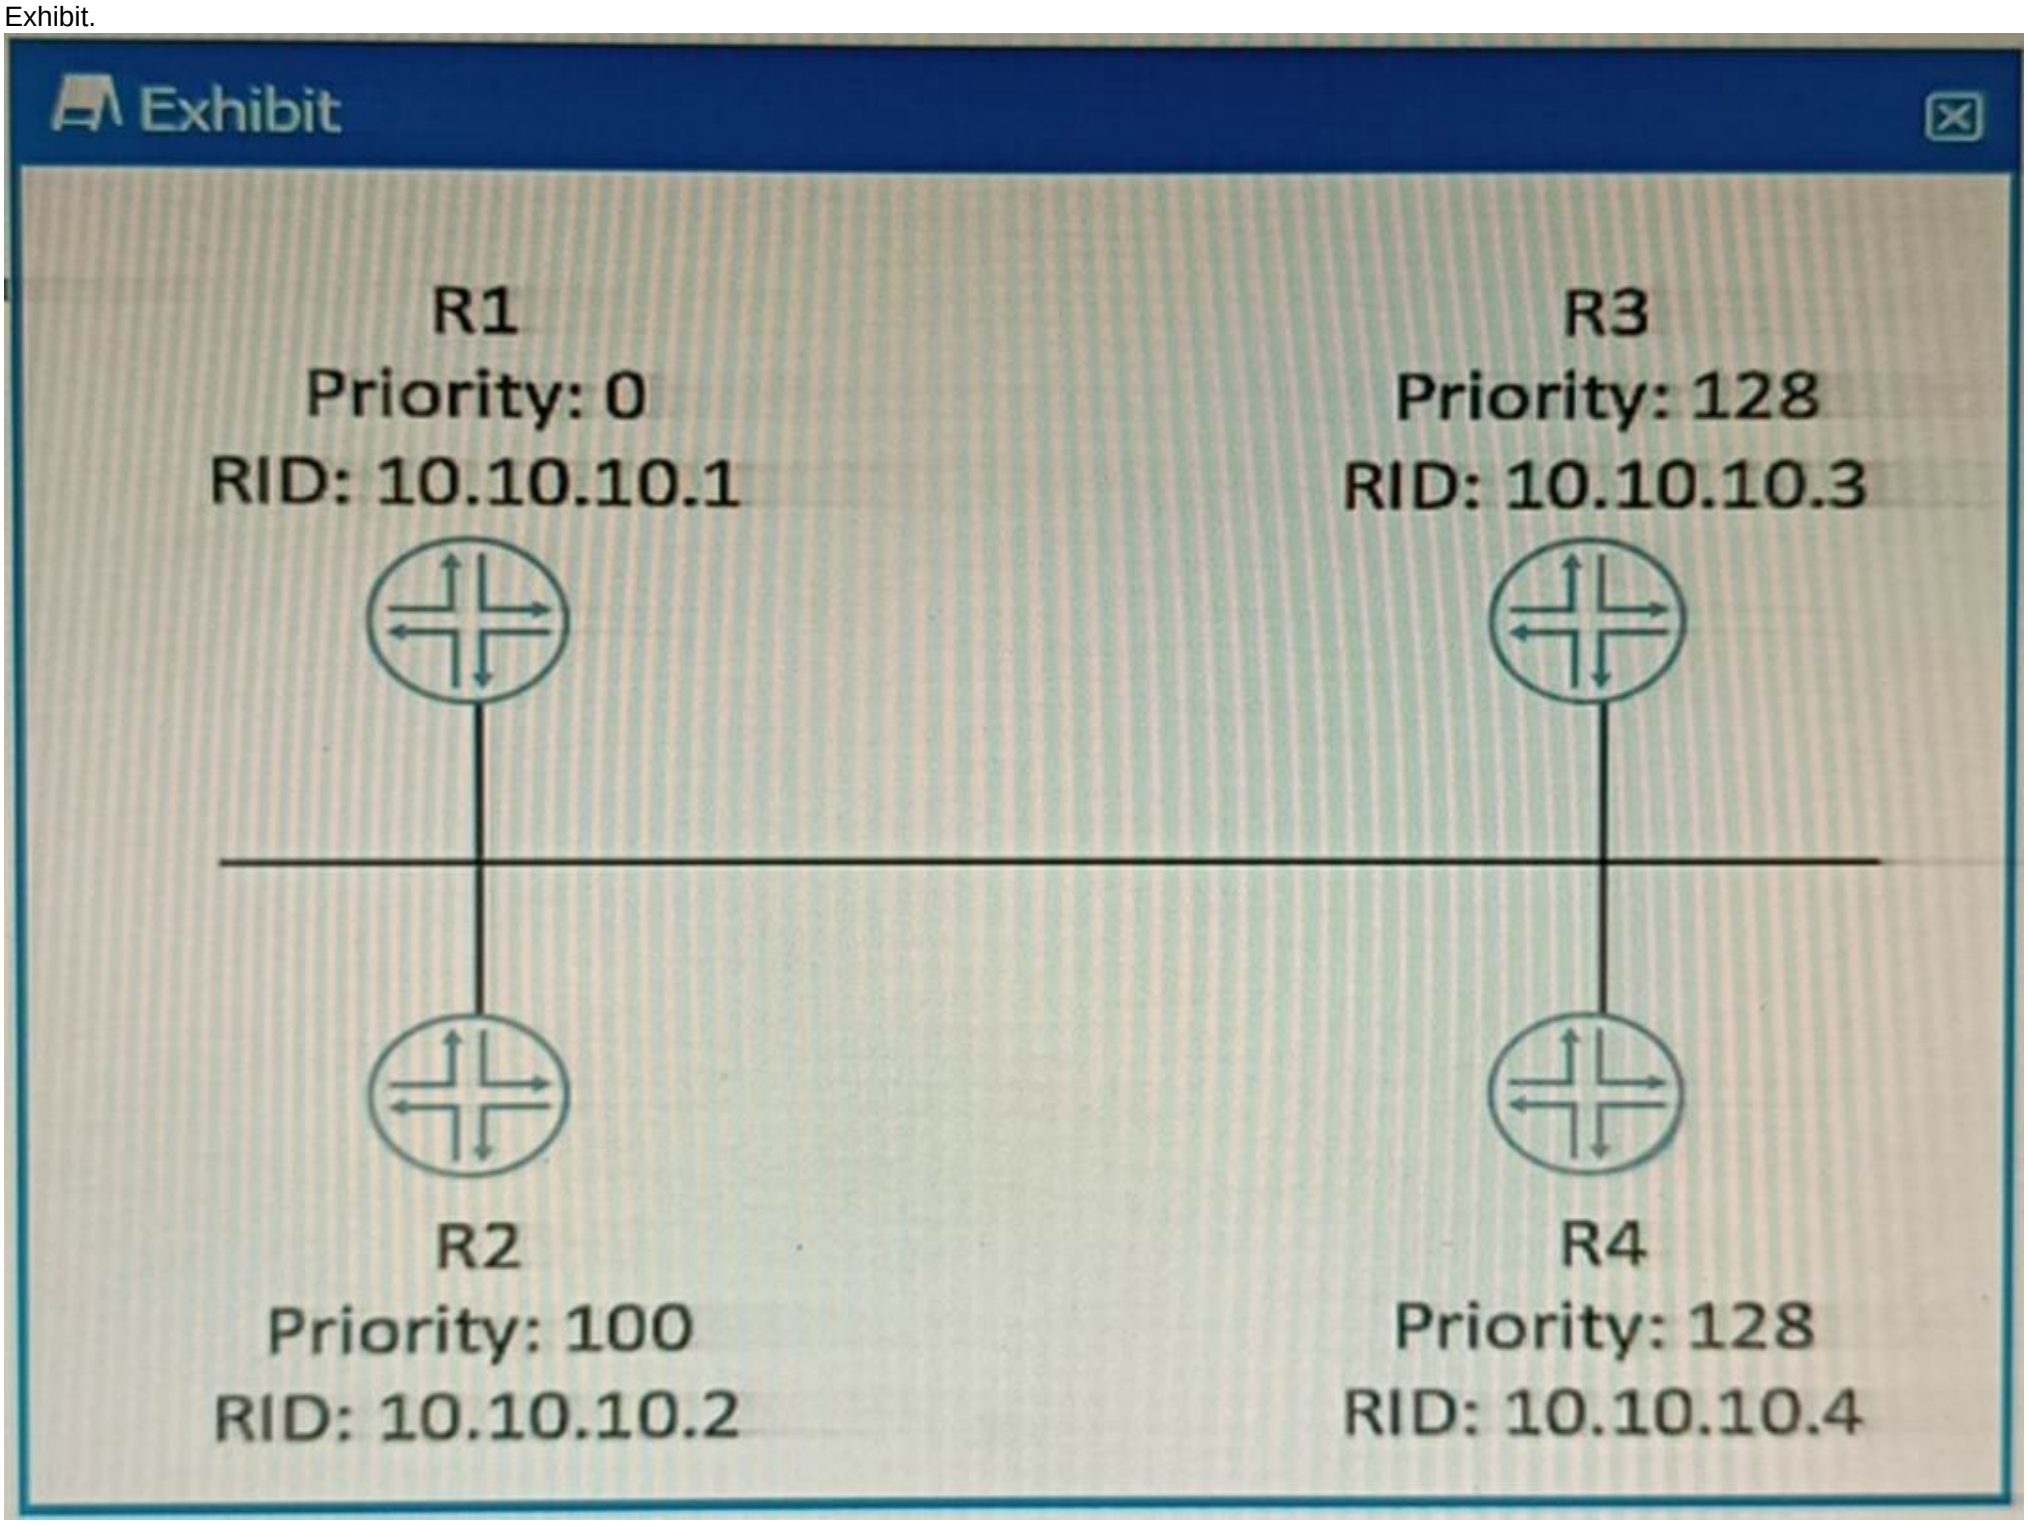

Which router will become the OSPF BDR if all routers are powered on at the same time?

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 2
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 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

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 5
Exhibit.

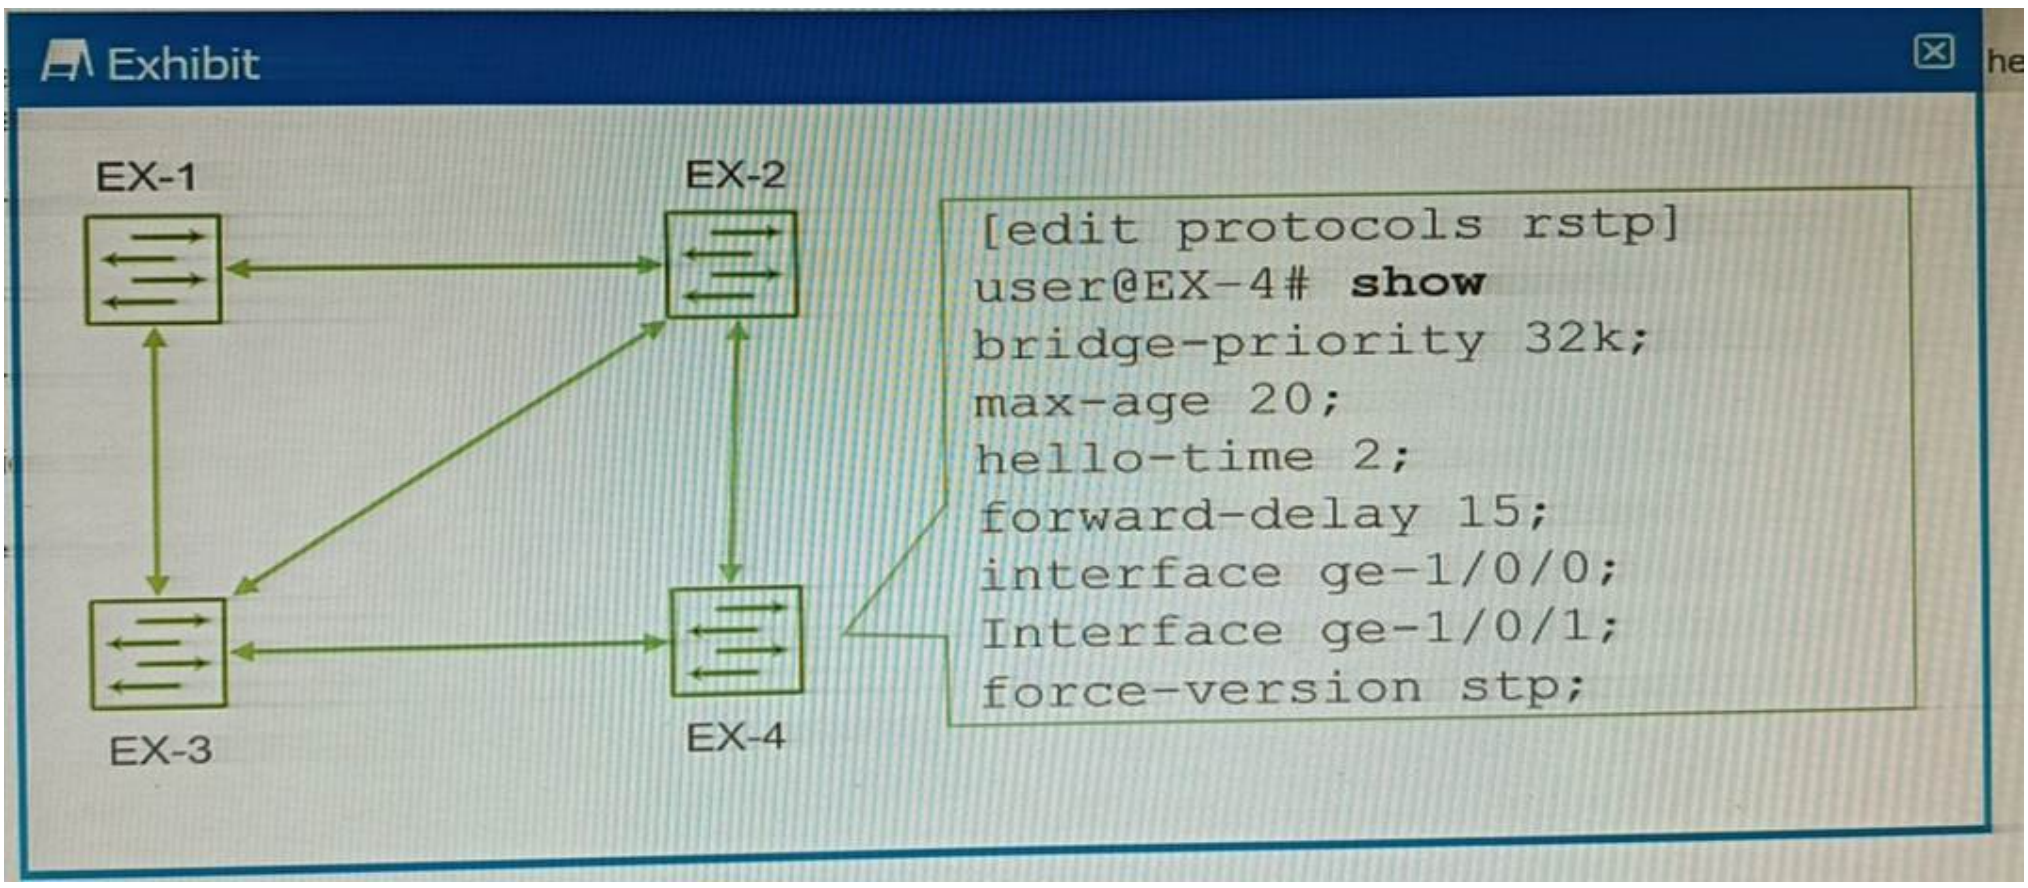

You have configured the four EX Series switches with RSTP, as shown in the exhibit. You discover that whenever a link between switches goes up or down, the switches take longer
than expected for RSTP to converge, using the default settings.
In this scenario, which action would solve the delay in RSTP 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 6
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 7
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 8
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 9
Exhibit

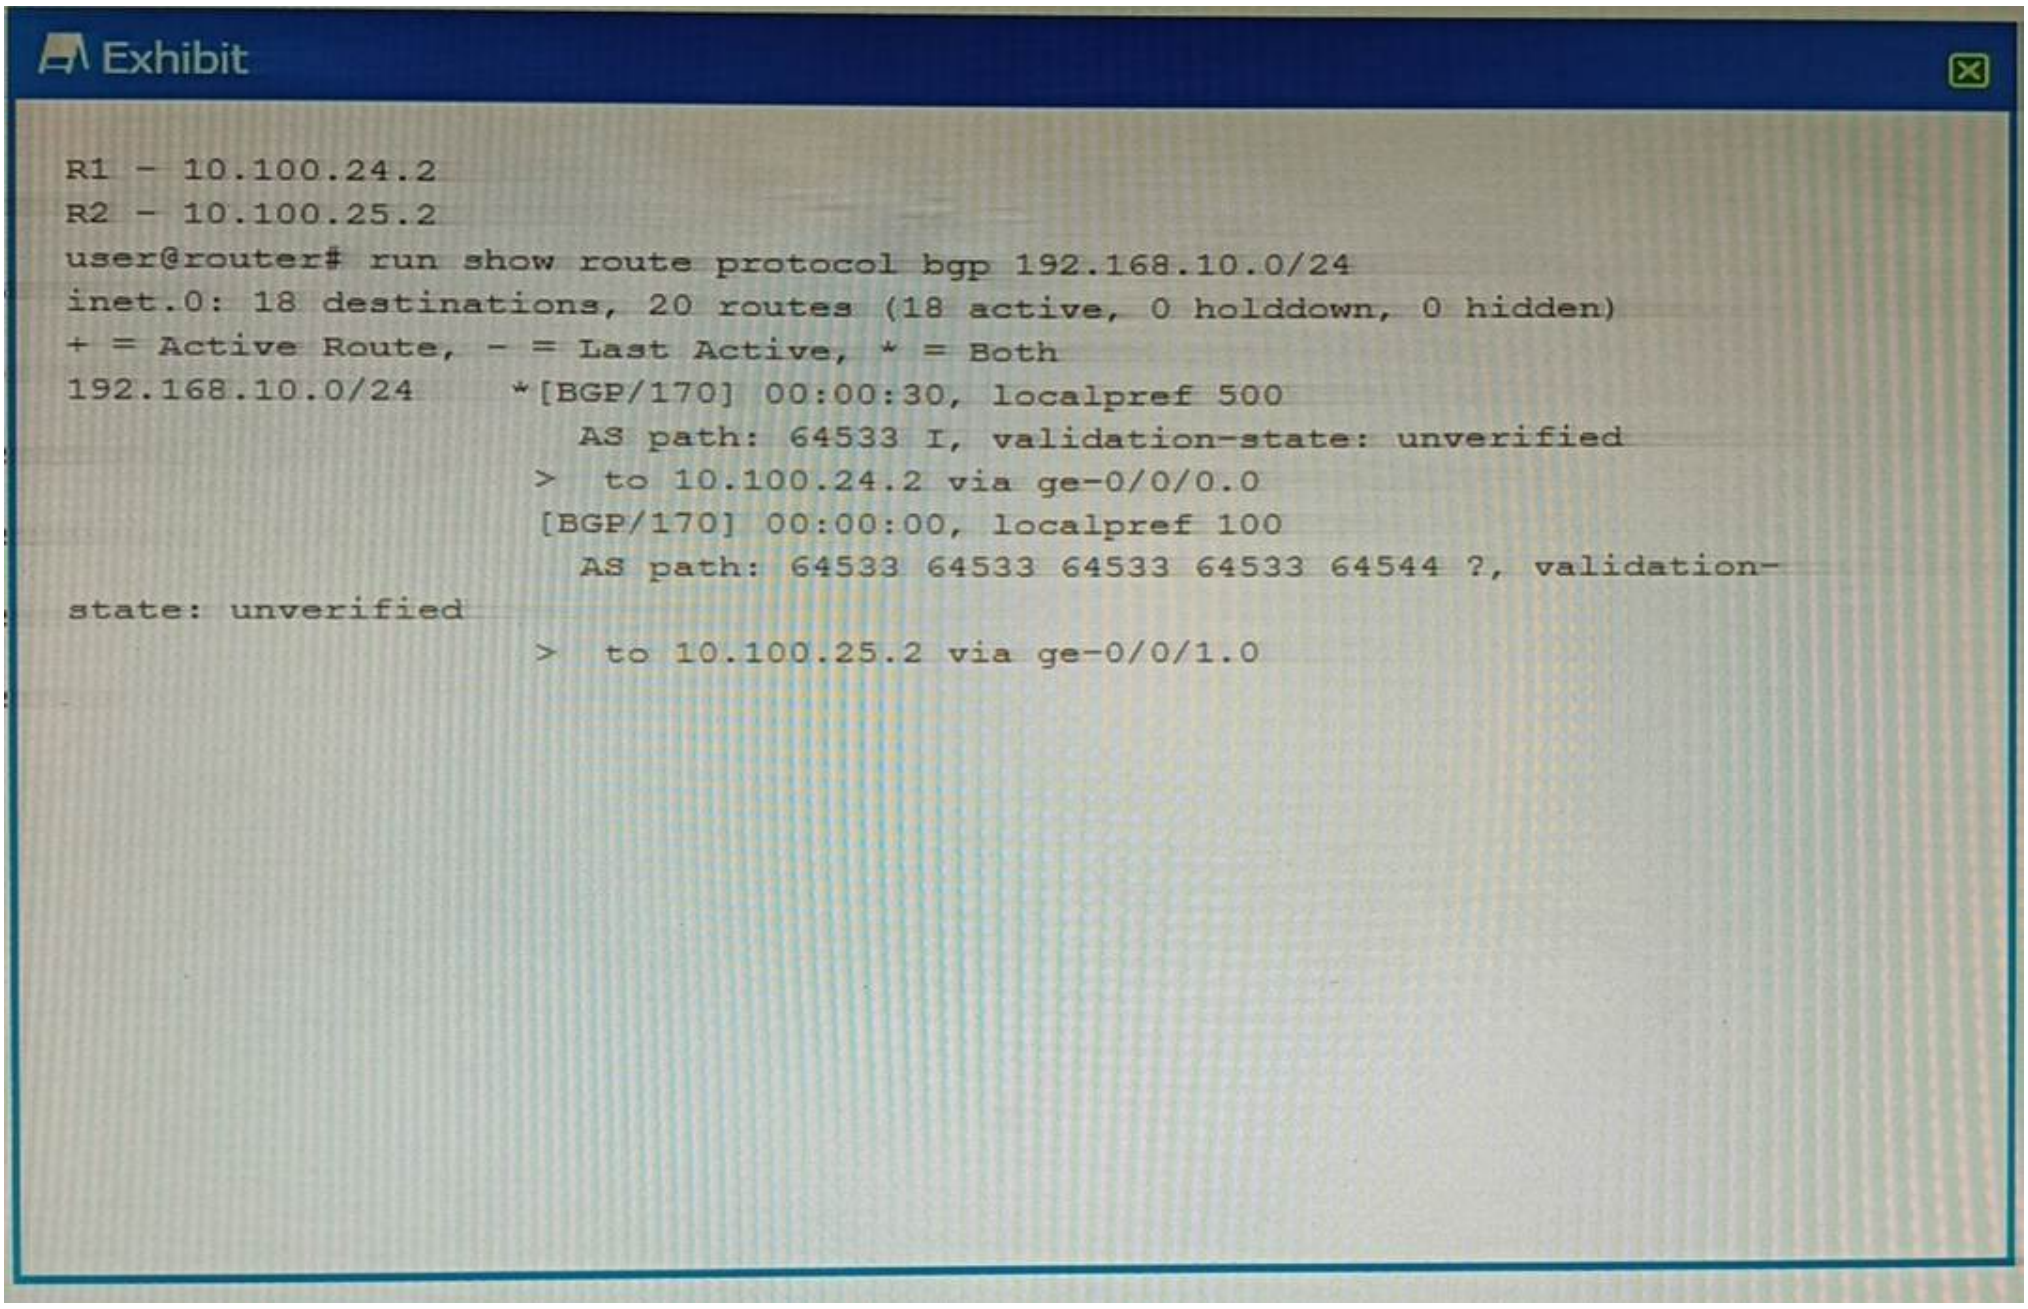

You are troubleshooting an issue where traffic to 192.168.10.0/24 is being sent to R1 instead of your desired path through R2.
Referring to the exhibit, what is the reason for the problem?

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 10
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
Exhibit

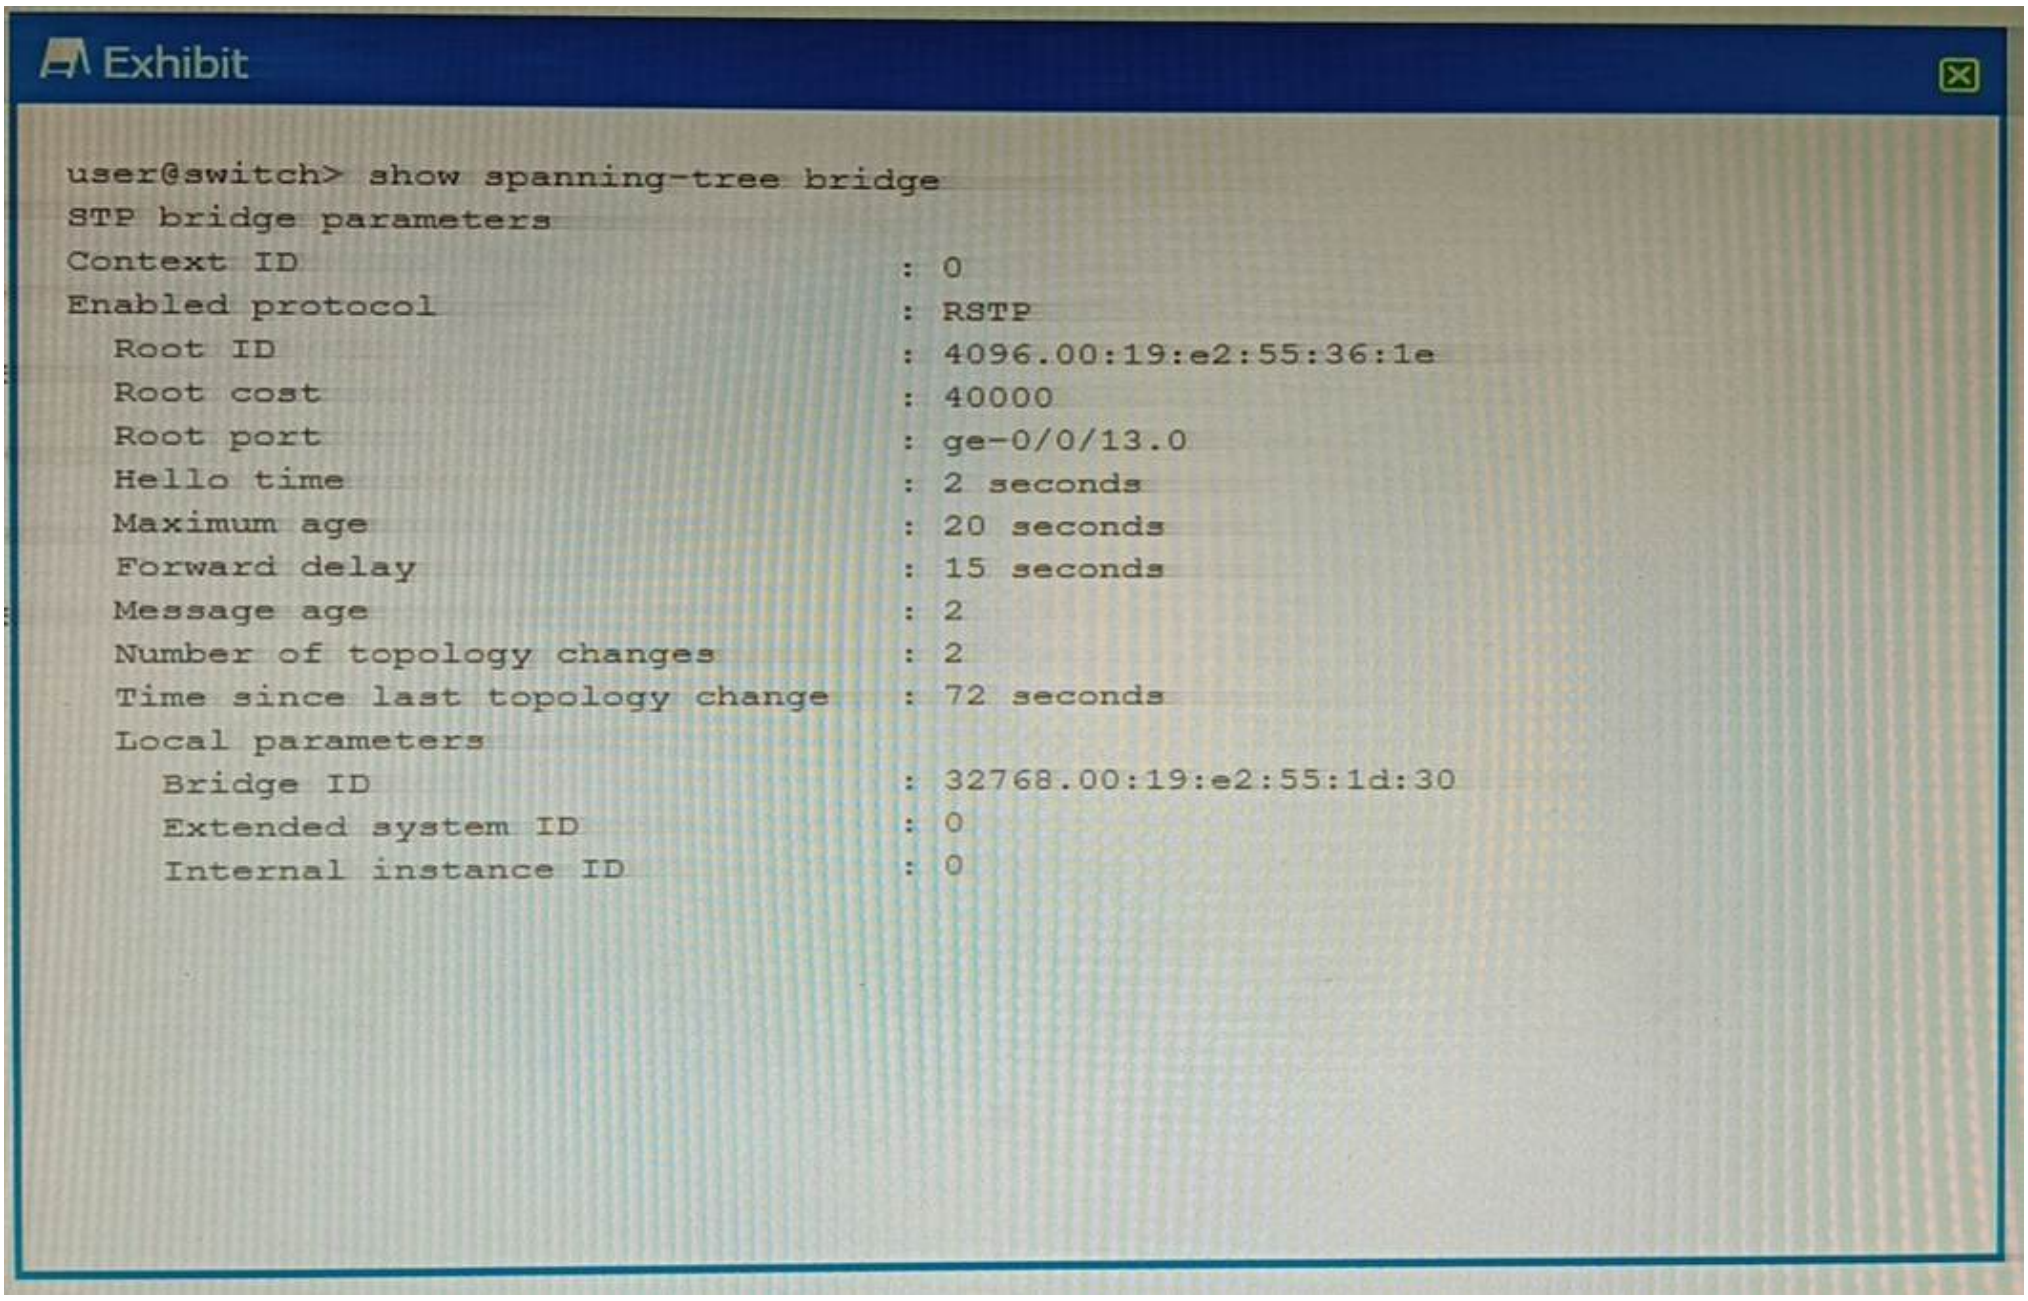

Referring to the exhibit, which statement is correct?

A. The local device is using a bridge priority of 4k.
B. The root bridge is using a bridge priority of 4k.
C. The root bridge has not been elected for this RSTP topology.
D. The local device is the root bridge for this RSTP topology.

Answer: D

Explanation:
In a Rapid Spanning Tree Protocol (RSTP) topology, the root bridge is determined by the switch with the lowest bridge priority value12. If all switches have the same priority, then the root bridge is assigned to the switch whose MAC address??s hex value is the lowest2. The default bridge priority value is 3276832. However, without the actual exhibit, it??s difficult to definitively determine which device is the root bridge. But based on the options provided, if we assume that the local device has a lower bridge priority or a lower MAC address than other devices in the network, then it could be considered as the root bridge for this RSTP topology45.

NEW QUESTION 11
Refer to the exhibit.

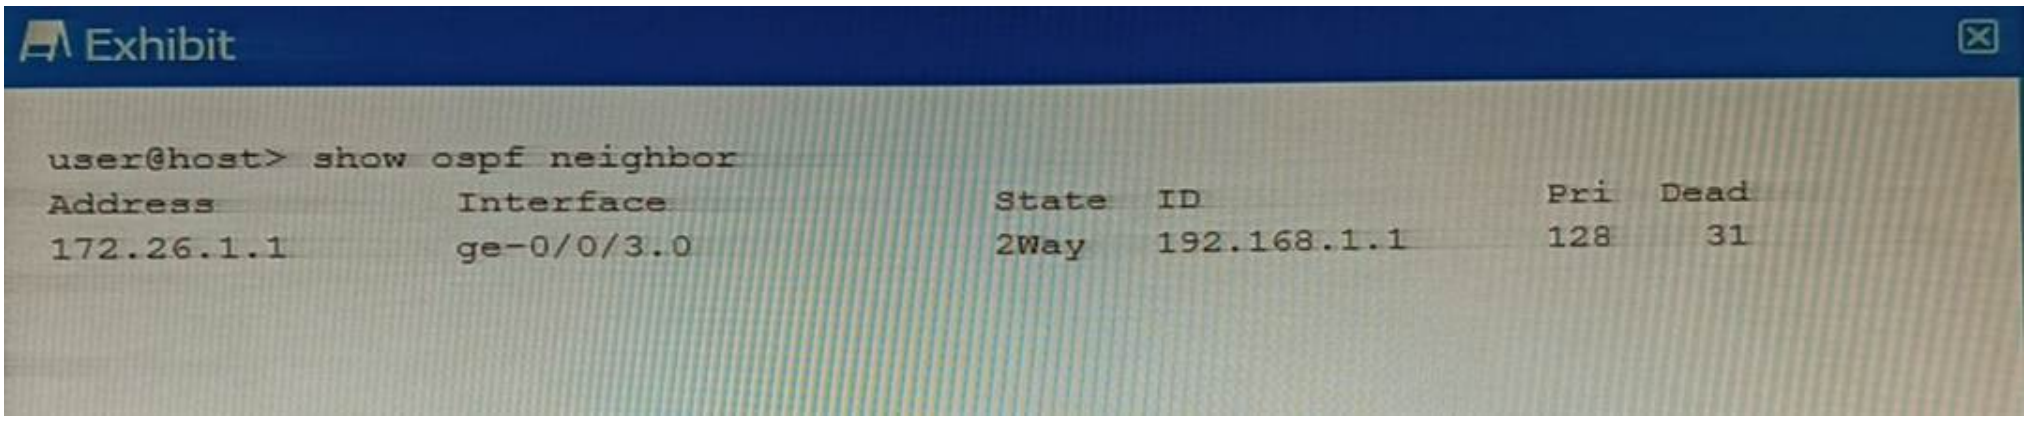

Referring to the output shown in the exhibit, which statement is correct?

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 13
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 16
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 18
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 22
You implemented the MAC address limit feature with the shutdown action on all interfaces on your switch.
In this scenario, which statement is correct when a violation occurs?

A. By default, you must manually clear the violation for the interface to send and receive traffic again.
B. By default, the violation will automatically be cleared after 300 seconds and the interface will resume sending and receiving traffic for all learned devices.
C. By default, devices that are learned before the violation occurs are still allowed to send and receive traffic through the specific interface.
D. By default, the interface will continue to send and receive traffic for all connected devices after a violation has occurred.

Answer: A

Explanation:
When the MAC address limit feature with the shutdown action is implemented on a switch, if a violation occurs, the interface is disabled and a system log entry is generated1. If the switch has been configured with the port-error-disable statement, the disabled interface recovers automatically upon expiration of the specified disable timeout1. However, if the switch has not been configured for auto-recovery from port error disabled conditions, you must manually clearthe violation by running the clear ethernet- switching port-error command for the interface to send and receive traffic again1. This explanation is based on the Enterprise Routing and Switching Specialist (JNCIS-ENT) documents and learning resources available at Juniper Networks1.

NEW QUESTION 23
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 25
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 29
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 34
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 37
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 39
Which two statements correctly describe RSTP port roles? (Choose two.)

A. The designated port forwards data to the downstream network segment or device.
B. The backup port is used as a backup for the root port.
C. The alternate port is a standby port for an edge port.
D. The root port is responsible for forwarding data to the root bridge.

Answer: AD

Explanation:
In Rapid Spanning Tree Protocol (RSTP), there are several port roles that determine the behavior of the port in the spanning tree1.
Option A suggests that the designated port forwards data to the downstream network segment or device. This is correct because the designated port is the port on a network segment that has the best path to the root bridge1. It??s responsible for forwarding frames towards the root bridge and sending configuration messages into its segment1.
Option D suggests that the root port is responsible for forwarding data to the root
bridge. This is also correct because the root port is always the link directly connected to the root bridge, or the shortest path to the root bridge1. It??s used to forward traffic towards the root bridge1.
Therefore, options A and D are correct.

NEW QUESTION 41
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 45
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 49
Which two statements are true about the default VLAN on Juniper switches? (Choose two.)

A. The default VLAN is set to a VLAN ID of 1 by default
B. The default VLAN ID is not assigned to any interface.
C. The default VLAN ID is not visible.
D. The default VLAN ID can be changed.

Answer: AD

Explanation:
On Juniper switches, the default VLAN is set to a VLAN ID of 1 by default12. This means that all interfaces on the switch are members of VLAN 1 until they are specifically assigned to another VLAN12. Therefore, option A is correct.
The default VLAN ID can be changed12. This allows network administrators to configure the switch to use a different VLAN as the default, if necessary12. Therefore, option D is correct.

NEW QUESTION 54
Exhibit

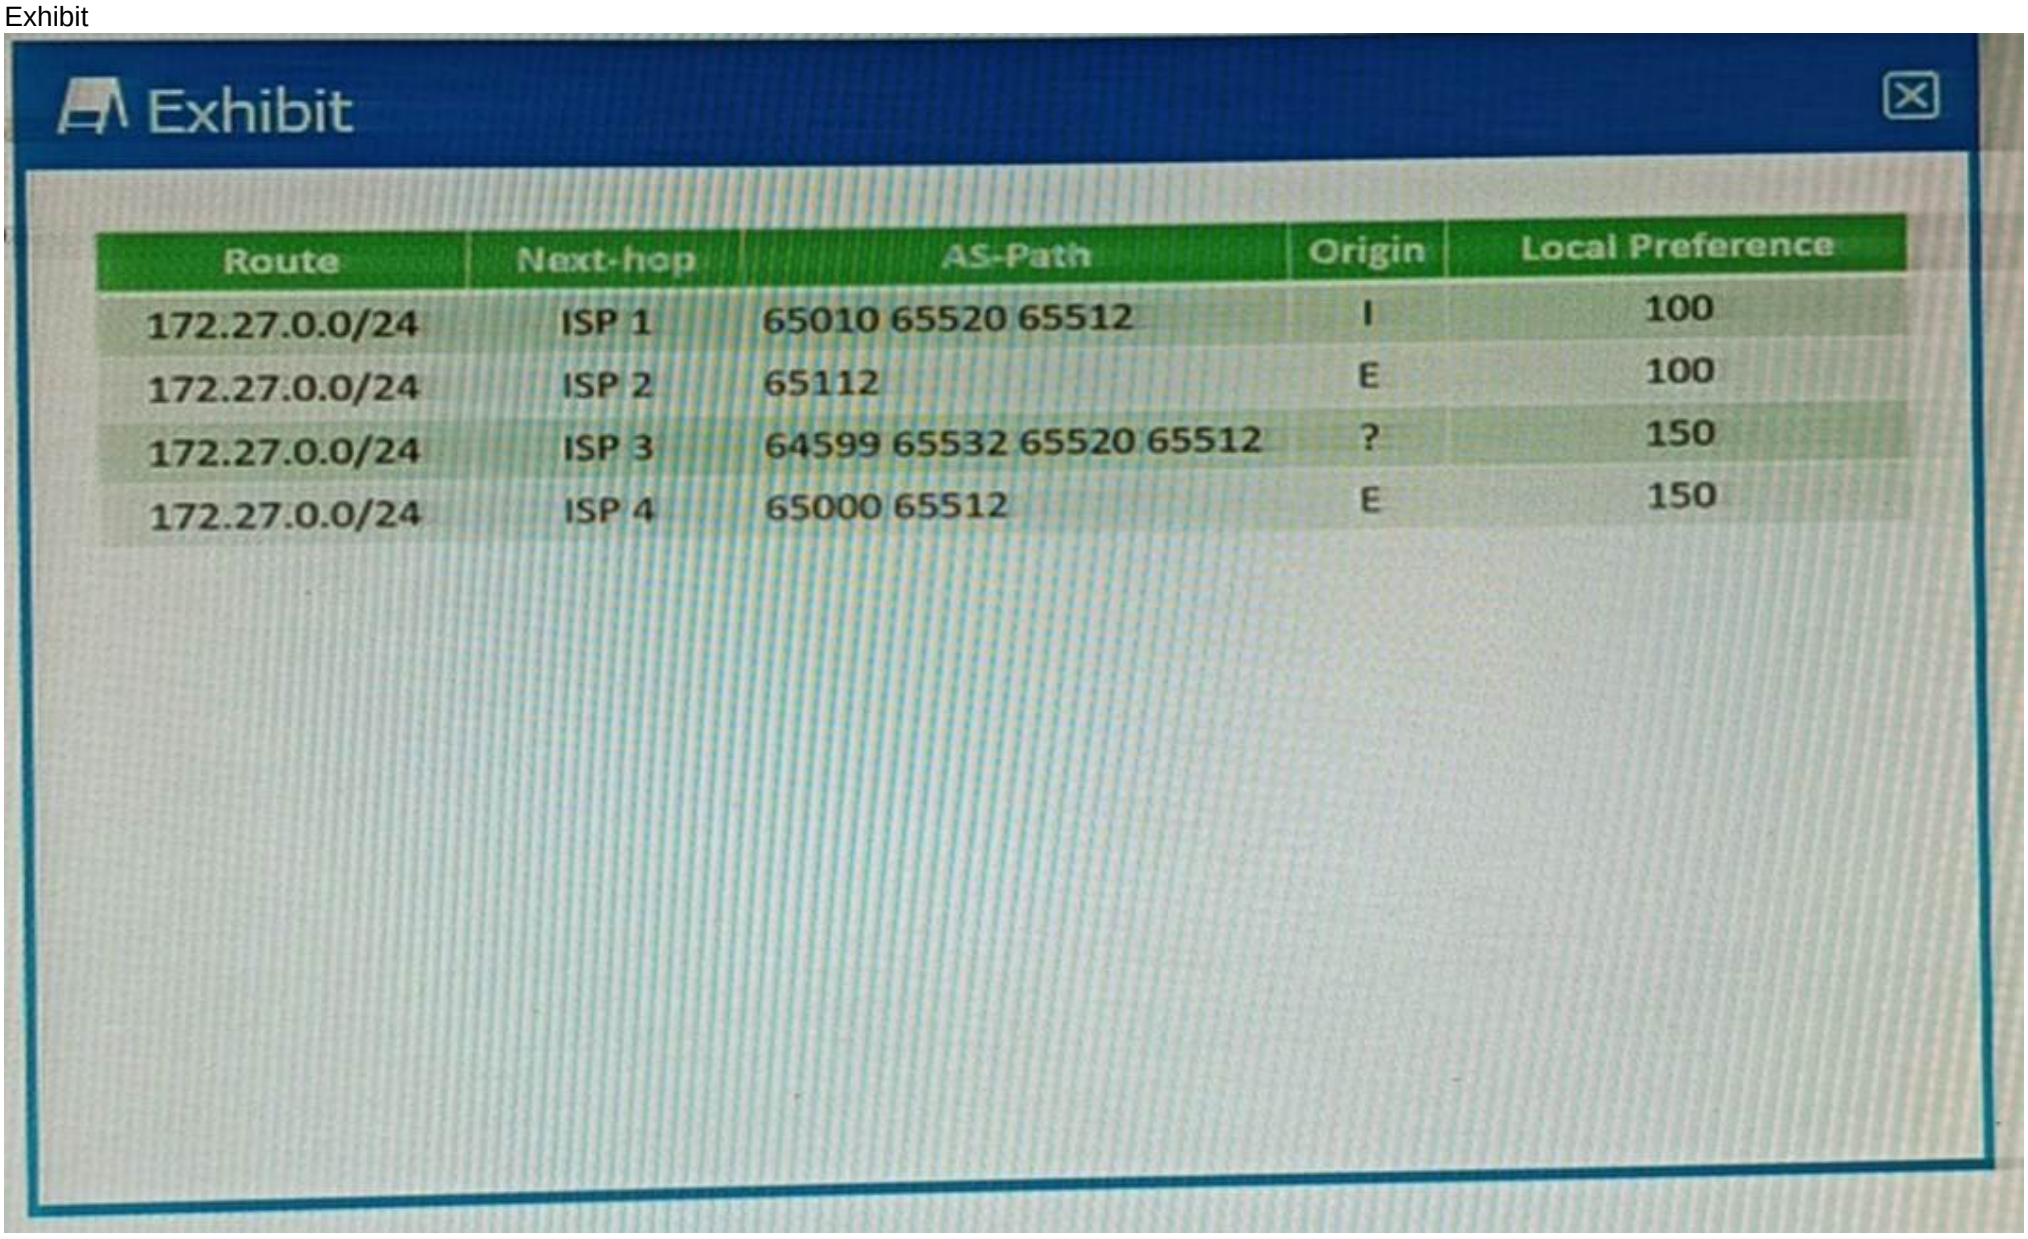

You are receiving the BGP route shown in the exhibit from four different upstream ISPs. Referring to the exhibit, which ISP will be selected as the active path?

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 55
Which statement is correct about controlling the routes installed by a RIB group?

A. An import policy is applied to the RIB group.
B. Only routes in the last table are installed.
C. A firewall filter must be configured to install routes in the RIB groups.
D. An export policy is applied to the RIB group.

Answer: A

Explanation:
A RIB group is a configuration that allows a routing protocol to install routes into multiple routing tables in Junos OS. 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 or group. A RIB group can also include an import- policy statement, which specifies one or more policies to control which routes are imported into the destination routing table or group1.
An import policy is a policy statement that defines the criteria for accepting or rejecting routes from the source routing table. An import policy can also modify the attributes of the imported routes, such as preference, metric, or community. An import policy can be applied to a RIB group by using the import-policy statement under the [edit routing-options rib- groups] hierarchy level1.
Therefore, option A is correct, because an import policy is applied to the RIB group to control which routes are installed in the destination routing table or group.
Option B is incorrect, because all routes in the source routing table are imported into the destination routing table or group, unless filtered by an import policy.
Option C is incorrect, because a firewall filter is not used to install routes in the RIB groups; a firewall filter is used to filter packets based on various criteria. Option D is incorrect, because an export policy is not applied to the RIB group; an export policy is applied to a routing protocol to control which routes are advertised to other devices.
References:
1: rib-groups | Junos OS | Juniper Networks

NEW QUESTION 59
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 60
You need to configure a LAG between your switches. In this scenario, which two statements are correct? (Choose two.)

A. Duplex and speed settings are not required to match on both participating devices.
B. Duplex and speed settings are required to match on both participating devices.
C. Member links are not required to be contiguous ports.
D. Member links are required to be contiguous ports.

Answer: BC

Explanation:
? B is correct because duplex and speed settings are required to match on both participating devices. According to the Juniper Networks documentation1, all the interfaces in a LAG must have the same speed and be in full-duplex mode. This ensures that the LAG can operate as a single logical link without any performance or compatibility issues.
? C is correct because member links are not required to be contiguous ports. According to the Juniper Networks documentation2, you can group any Ethernet interfaces on a switch into a LAG, regardless of their physical location or slot number. This provides flexibility and scalability for configuring LAGs on switches.

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 63
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 64
Exhibit

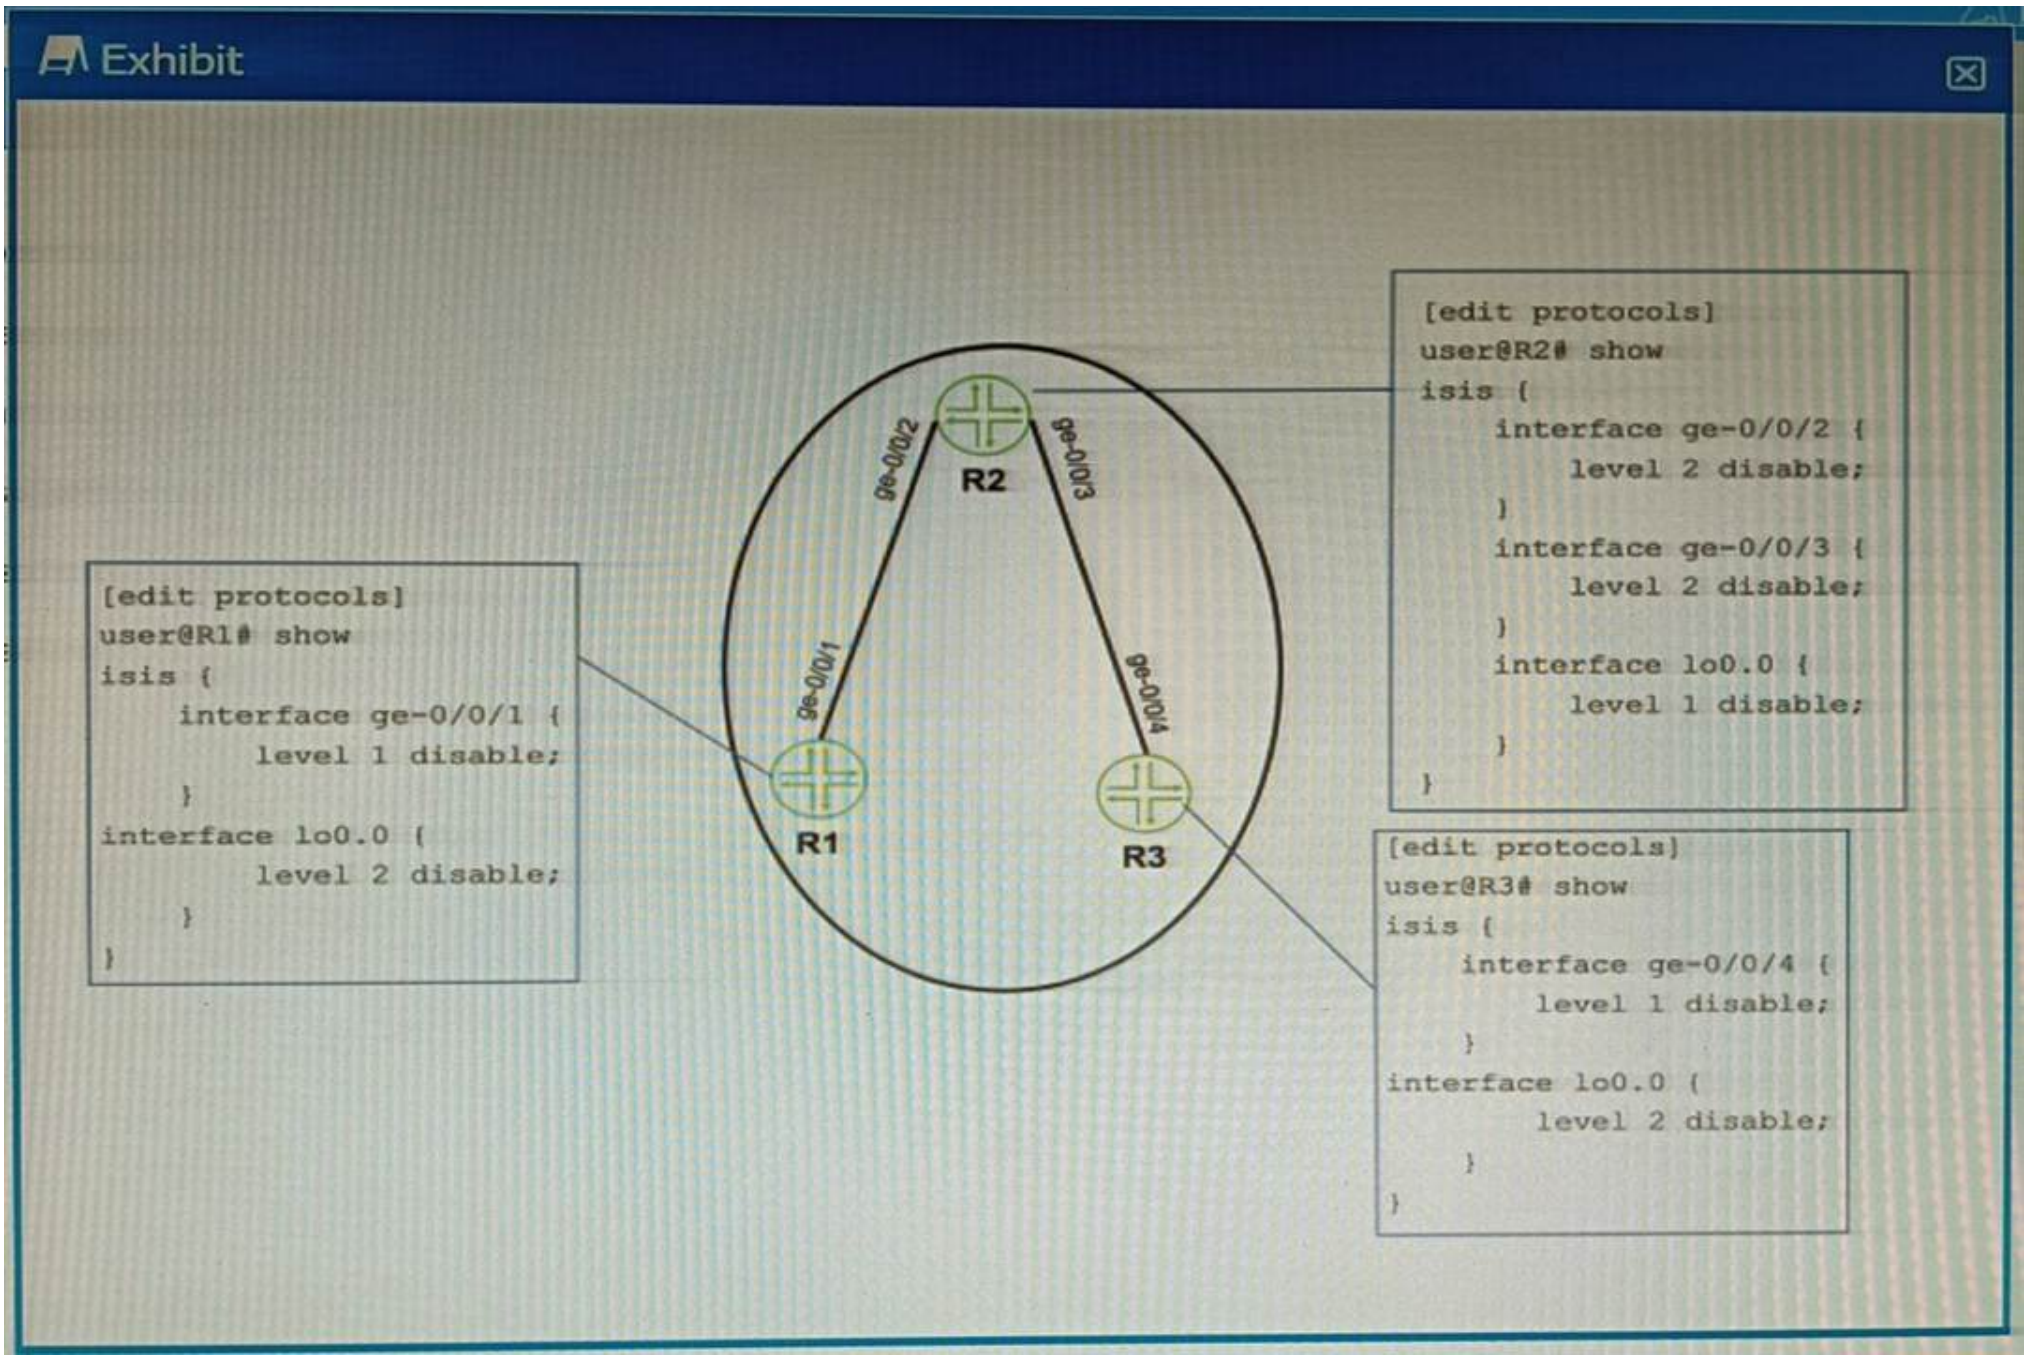

Referring to the exhibit, which two configuration changes must you apply for packets to reach from R1 to R3 using IS-IS? (Choose two.)

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 67

.....

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