Juniper

Exam Questions JN0-351

Enterprise Routing and Switching - Specialist (JNCIS-ENT)

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

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

B)

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

C)

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

D)

```
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

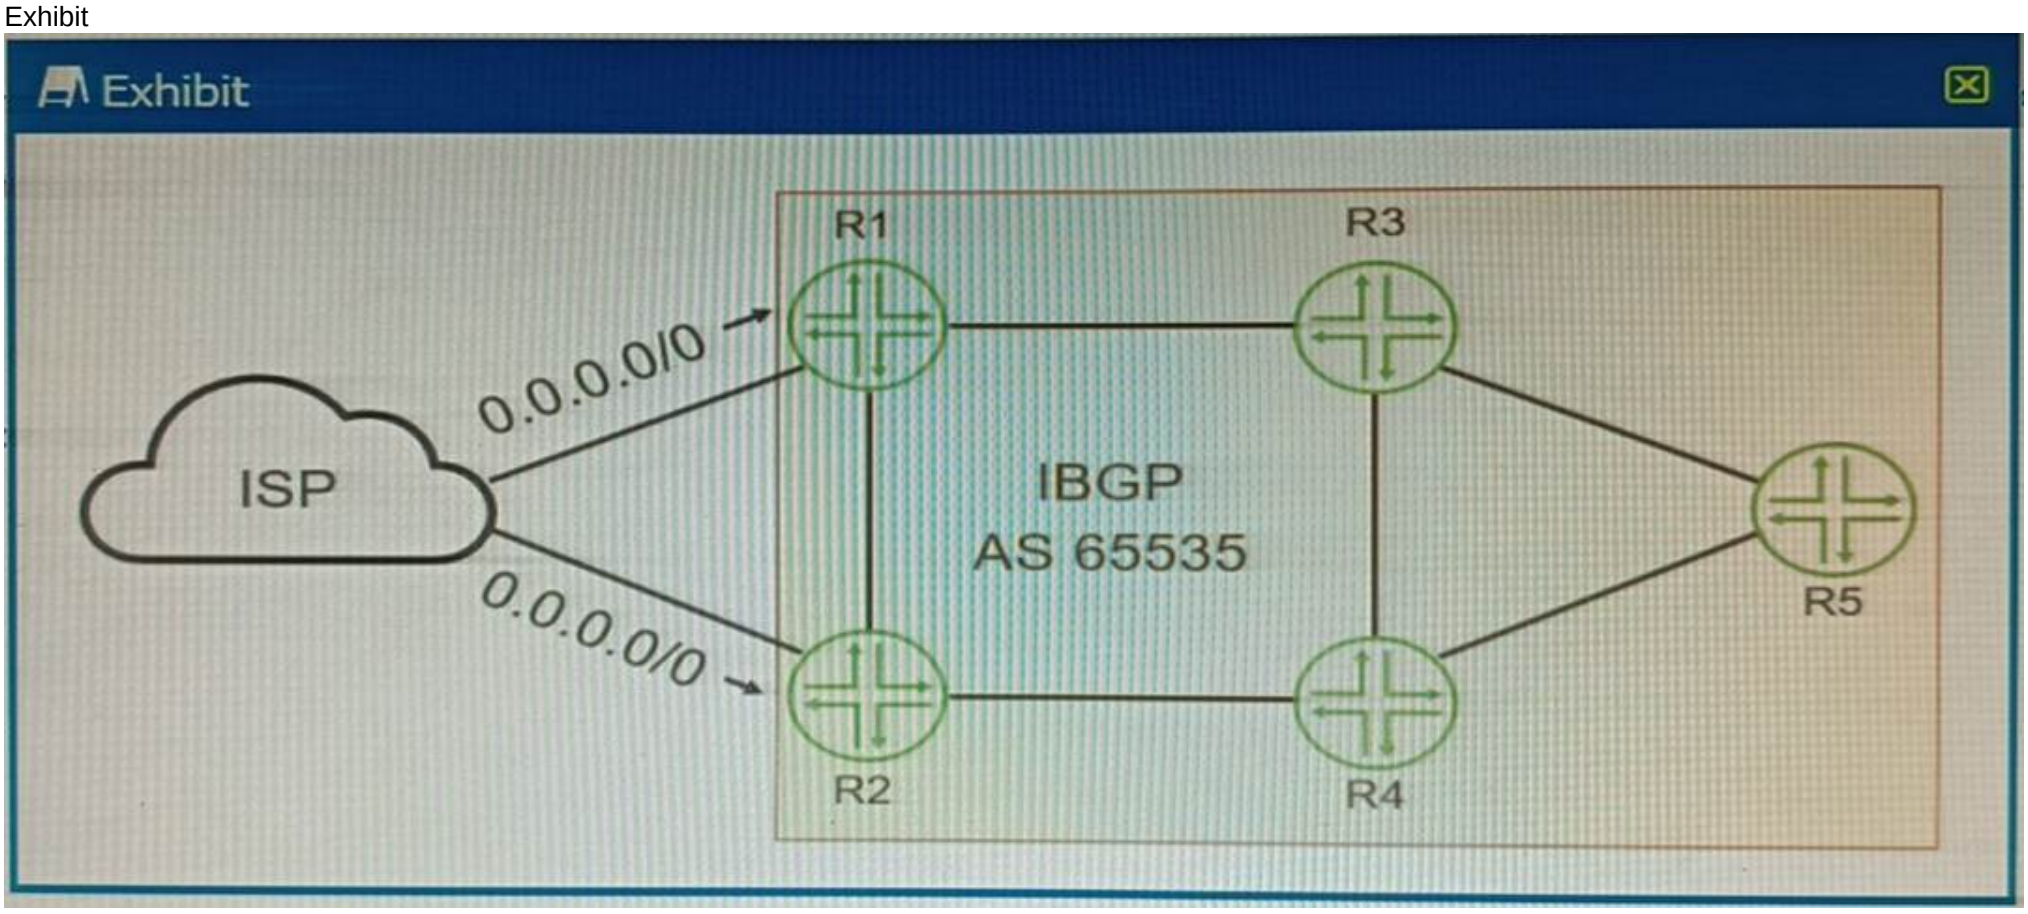

Your ISP is announcing a default route to both R1 and R2. You want your network routers to forward all Internet traffic through the R1 device
Which BGP attribute would you use?

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.

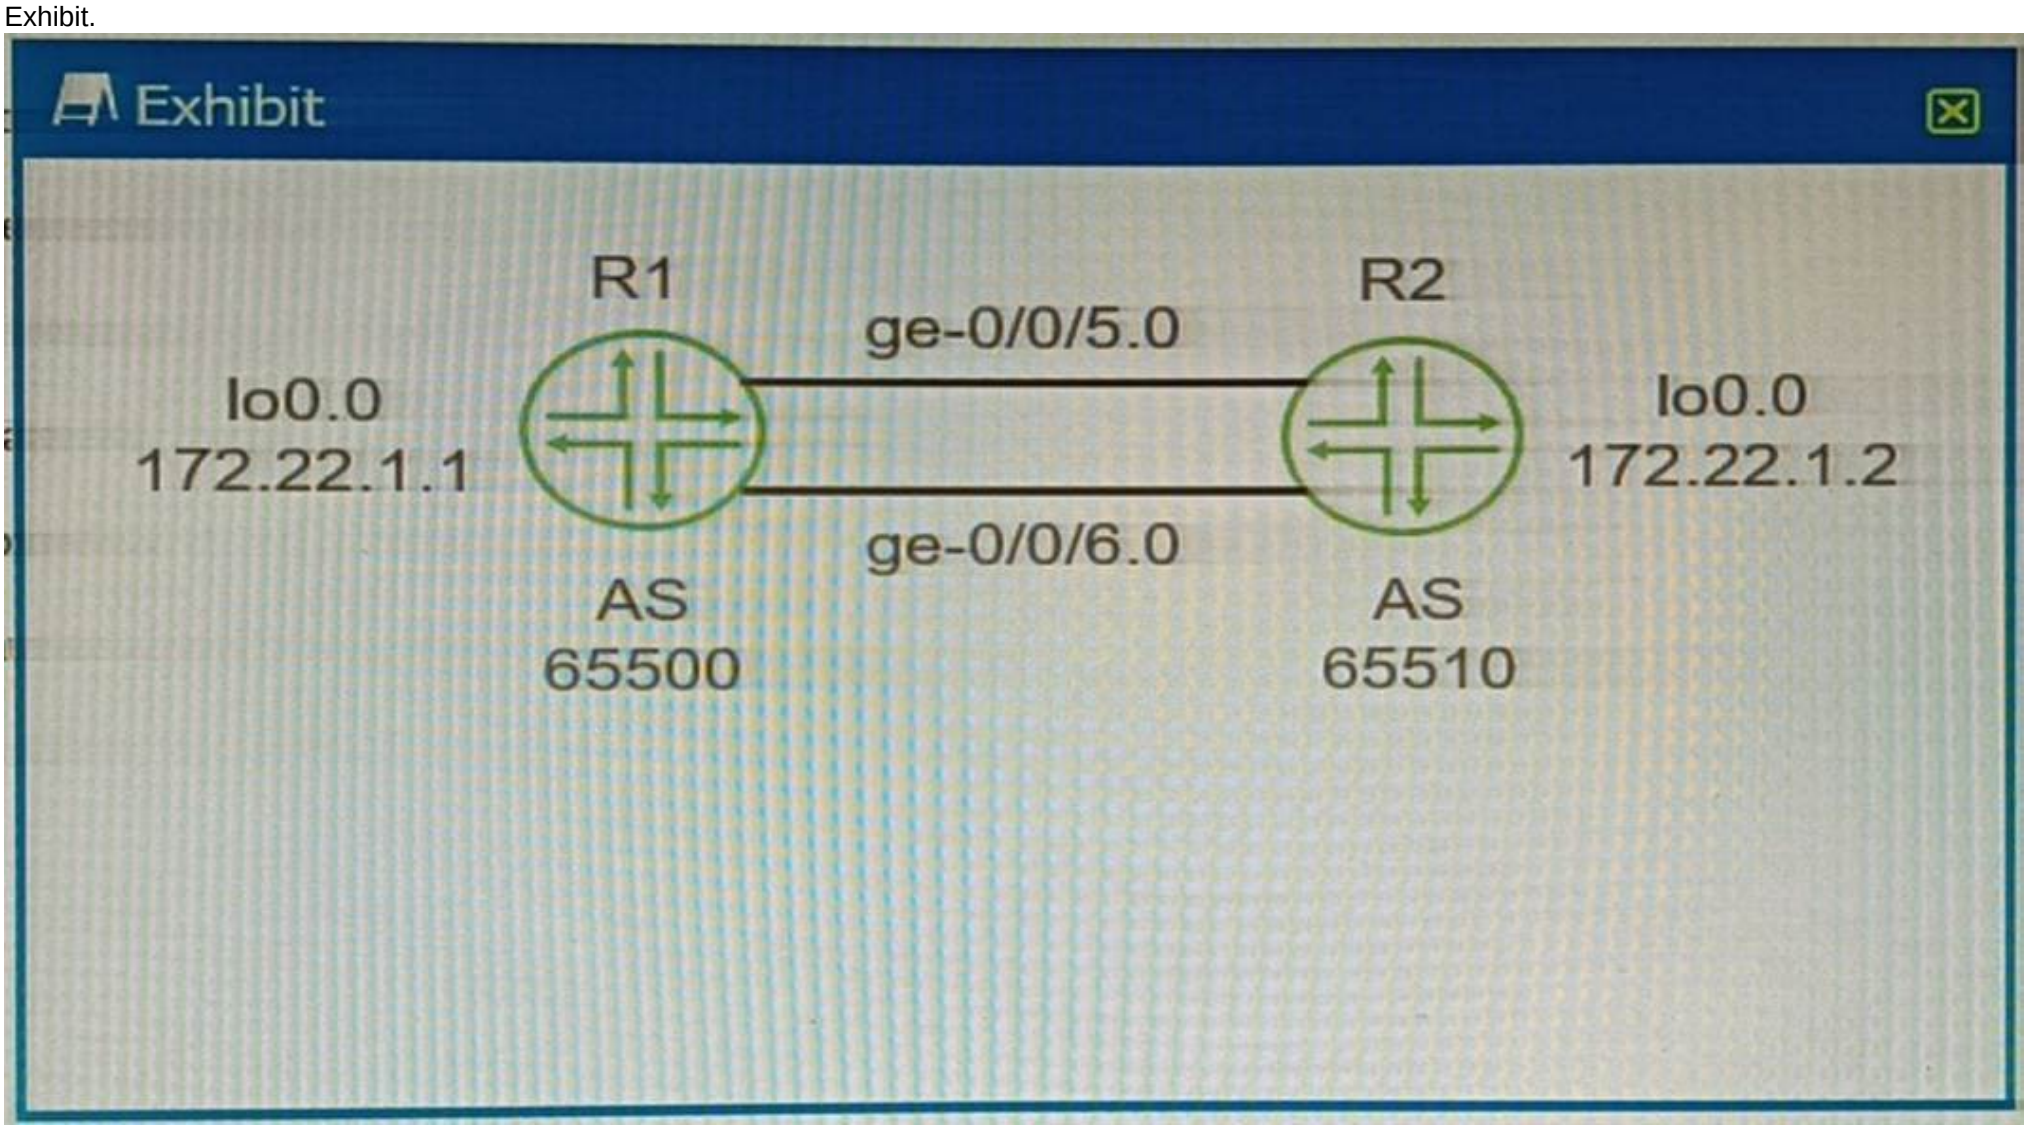

You want to enable redundancy for the EBGP peering between the two routers shown in the exhibit. Which three actions will you perform in this scenario? (Choose three.)

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

Answer: ABC

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

NEW QUESTION 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
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 6
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 7
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 8
What is the default keepalive time for BGP?

A. 10 seconds
B. 60 seconds
C. 30 seconds
D. 90 seconds

Answer: B

Explanation:
The default keepalive time for BGP is 60 seconds1. The keepalive time is the interval at which BGP sends keepalive messages to maintain the connection with its peer1. If the keepalive message is not received within the hold time, the connection is considered lost1. By default, the hold time is three times the keepalive time, which is 180 seconds1.

NEW QUESTION 9
Exhibit.

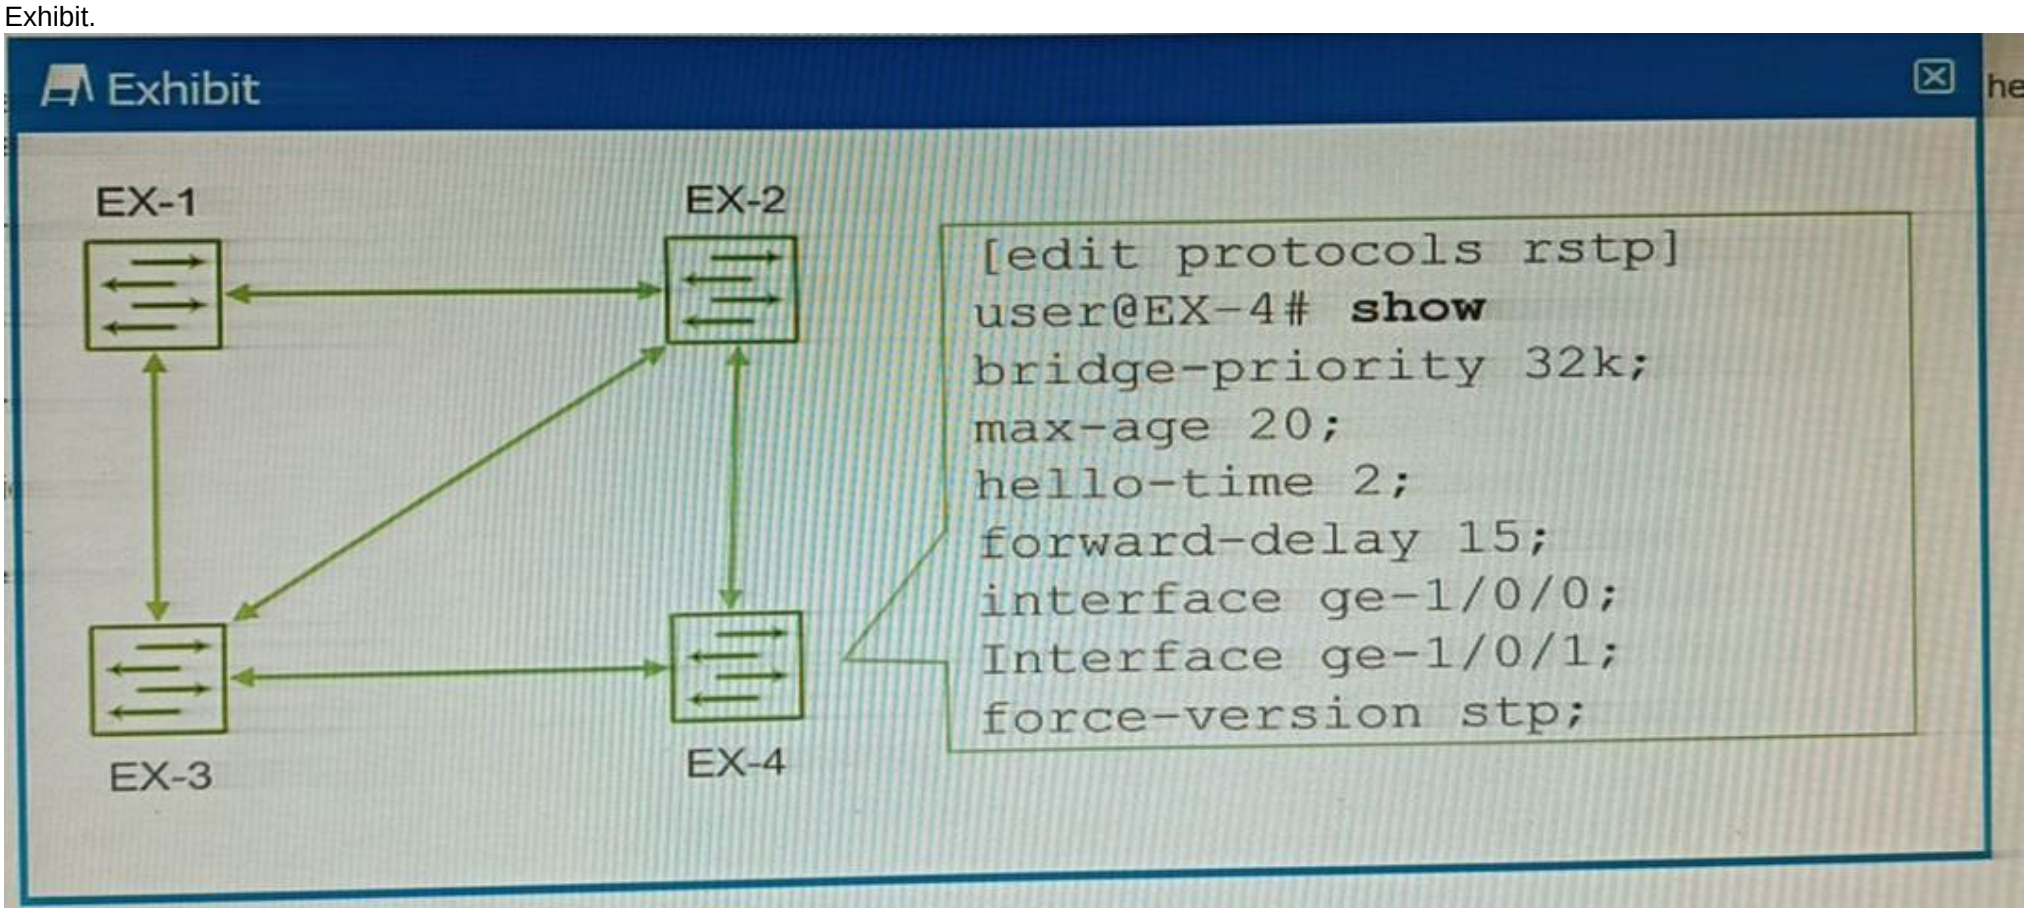

You have configured the four EX Series switches with RSTP, as shown in the exhibit. You discover that whenever a link between switches goes up or down, the switches take longer
than expected for RSTP to converge, using the default settings.
In this scenario, which action would solve the delay in RSTP 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 10
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
Exhibit.

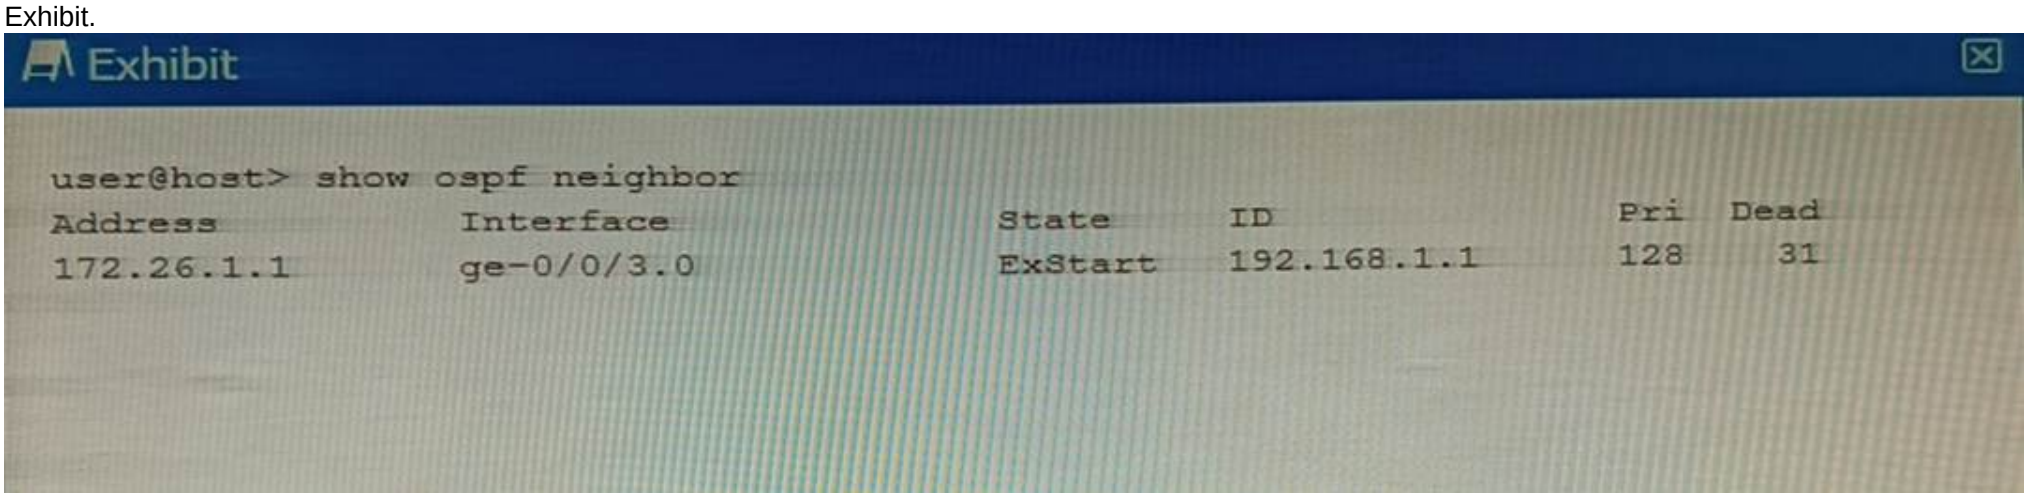

Why is this OSPF adjacency remaining in this state?

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 14
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 19
Exhibit

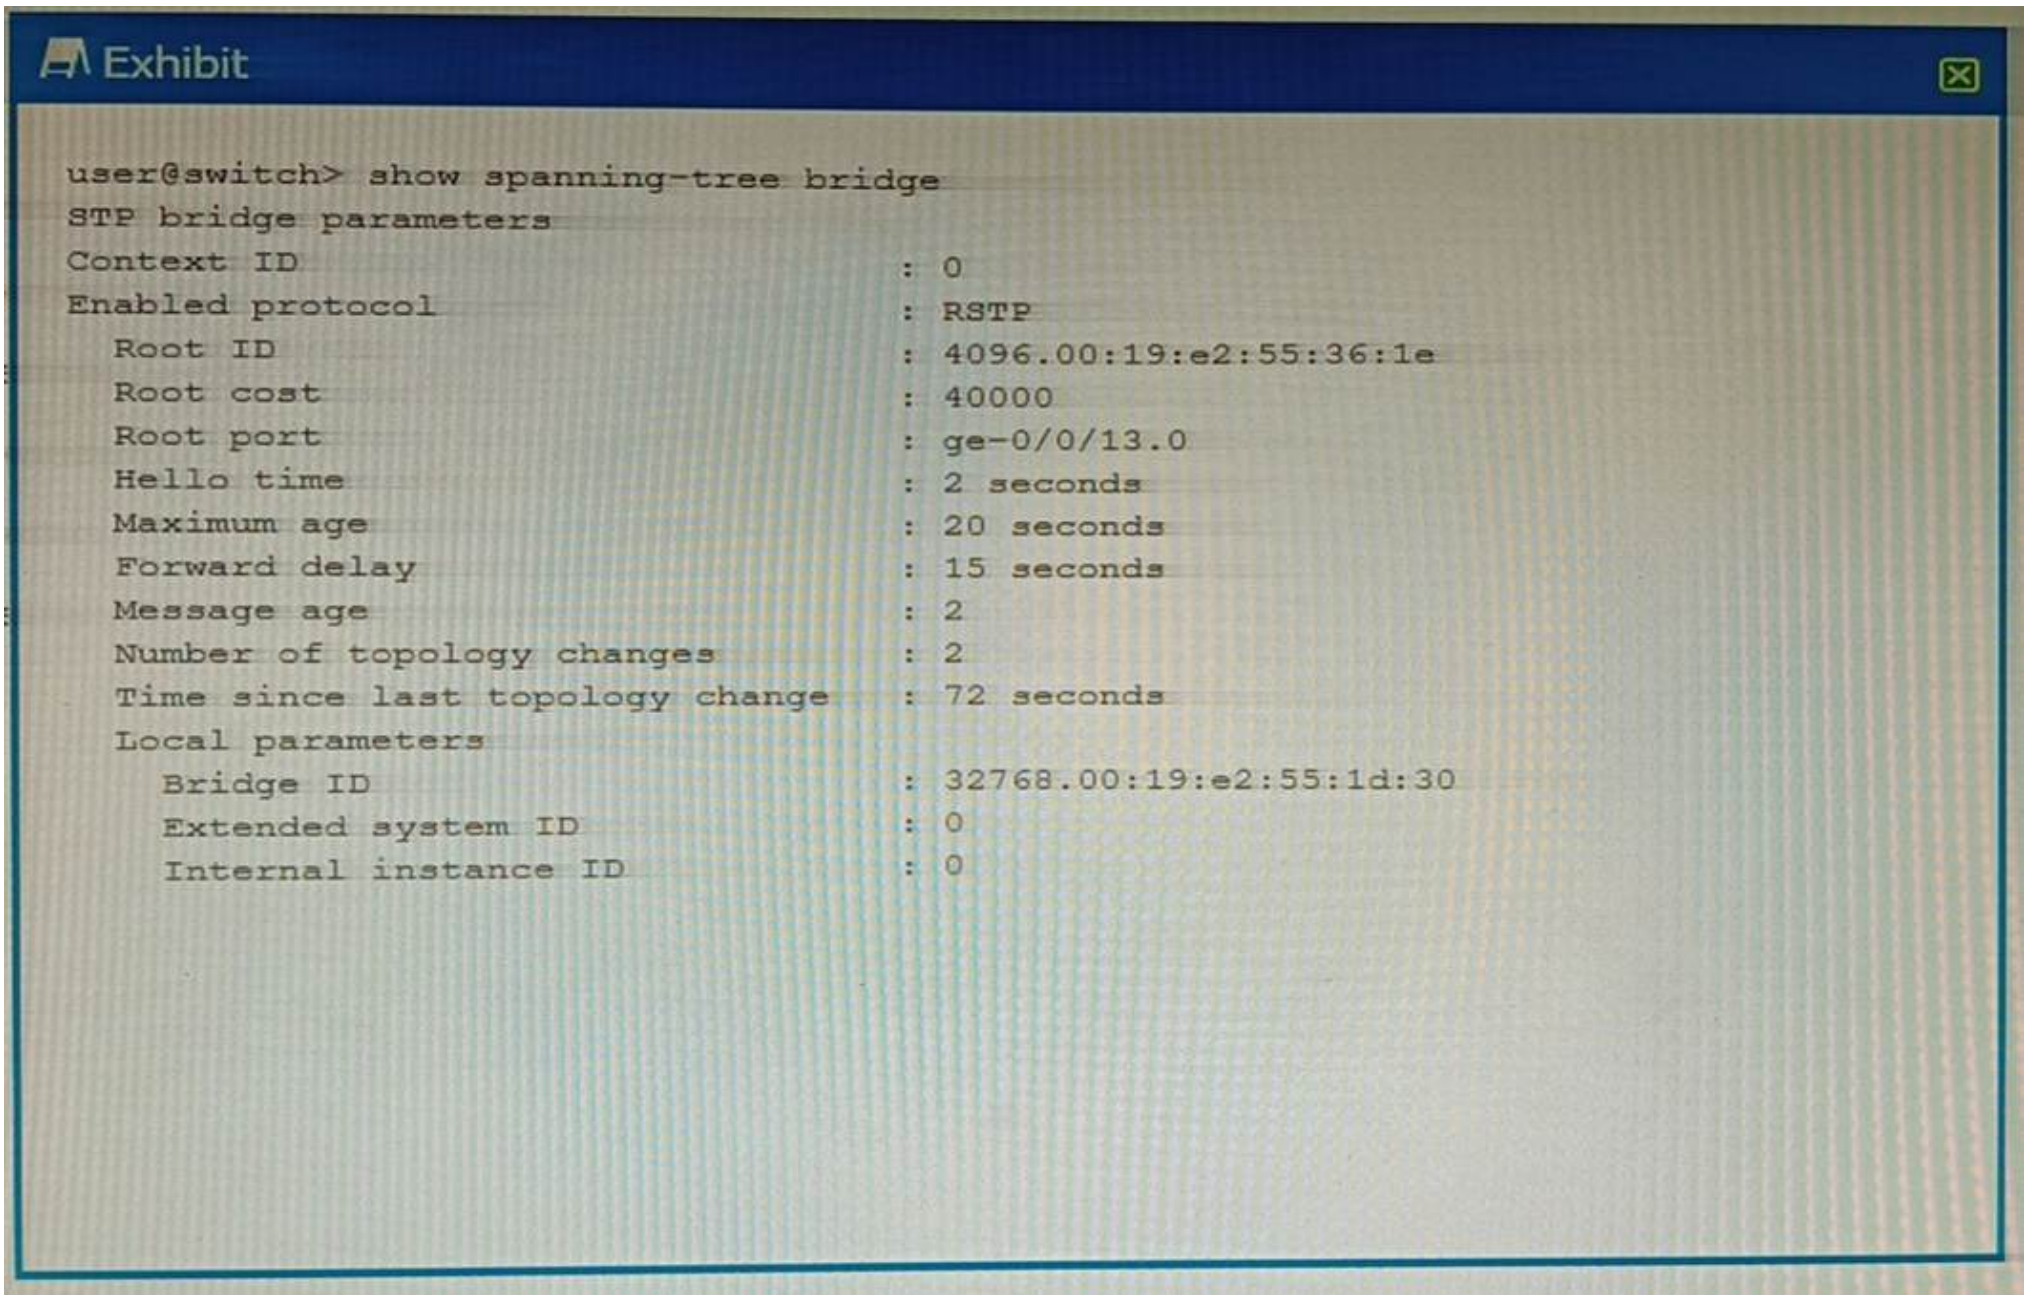

Referring to the exhibit, which statement is correct?

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 24
What is the default MAC age-out timer on an EX Series switch?

A. 30 minutes
B. 30 seconds
C. 300 minutes
D. 300 seconds

Answer: D

Explanation:
The default MAC age-out timer on an EX Series switch is 300 seconds12. The MAC age-out timer is the maximum time that an entry can remain in the MAC table before it ??ages out,?? or is removed31. This configuration can influence efficiency of network resource use by affecting the amount of traffic that is flooded to all interfaces1. When traffic is received for MAC addresses no longer in the Ethernet routing table, the router floods the traffic to all interfaces1.

NEW QUESTION 26
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 31
Which two statements are correct about generated routes? (Choose two.)

A. Generated routes require a contributing route.
B. Generated routes show a next hop in the routing table.
C. Generated routes appear in the routing table as static routes
D. Generated routes cannot be redistributed into dynamic routing protocols.

Answer: AB

Explanation:
? A is correct because generated routes require a contributing route. A contributing route is a route that matches the destination prefix of the generated route and has a valid next hop1. A generated route is only installed in the routing table if there is at least one contributing route available2. This ensures that the generated route is reachable and useful. If there is no contributing route, the generated route is not added to the routing table2.
? B is correct because generated routes show a next hop in the routing table. A generated route inherits the next hop of its primary contributing route, which is the most preferred route among all the contributing routes2. The next hop of the generated route can be either an IP address or an interface name, depending on the type of the contributing route2. The next hop of the generated route can also be modified by a routing policy3.

NEW QUESTION 33
Which two events cause a router to advertise a connected network to OSPF neighbors? (Choose two.)

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

Answer: AD

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

NEW QUESTION 36
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 41
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 44
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 49
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 52
Exhibit.

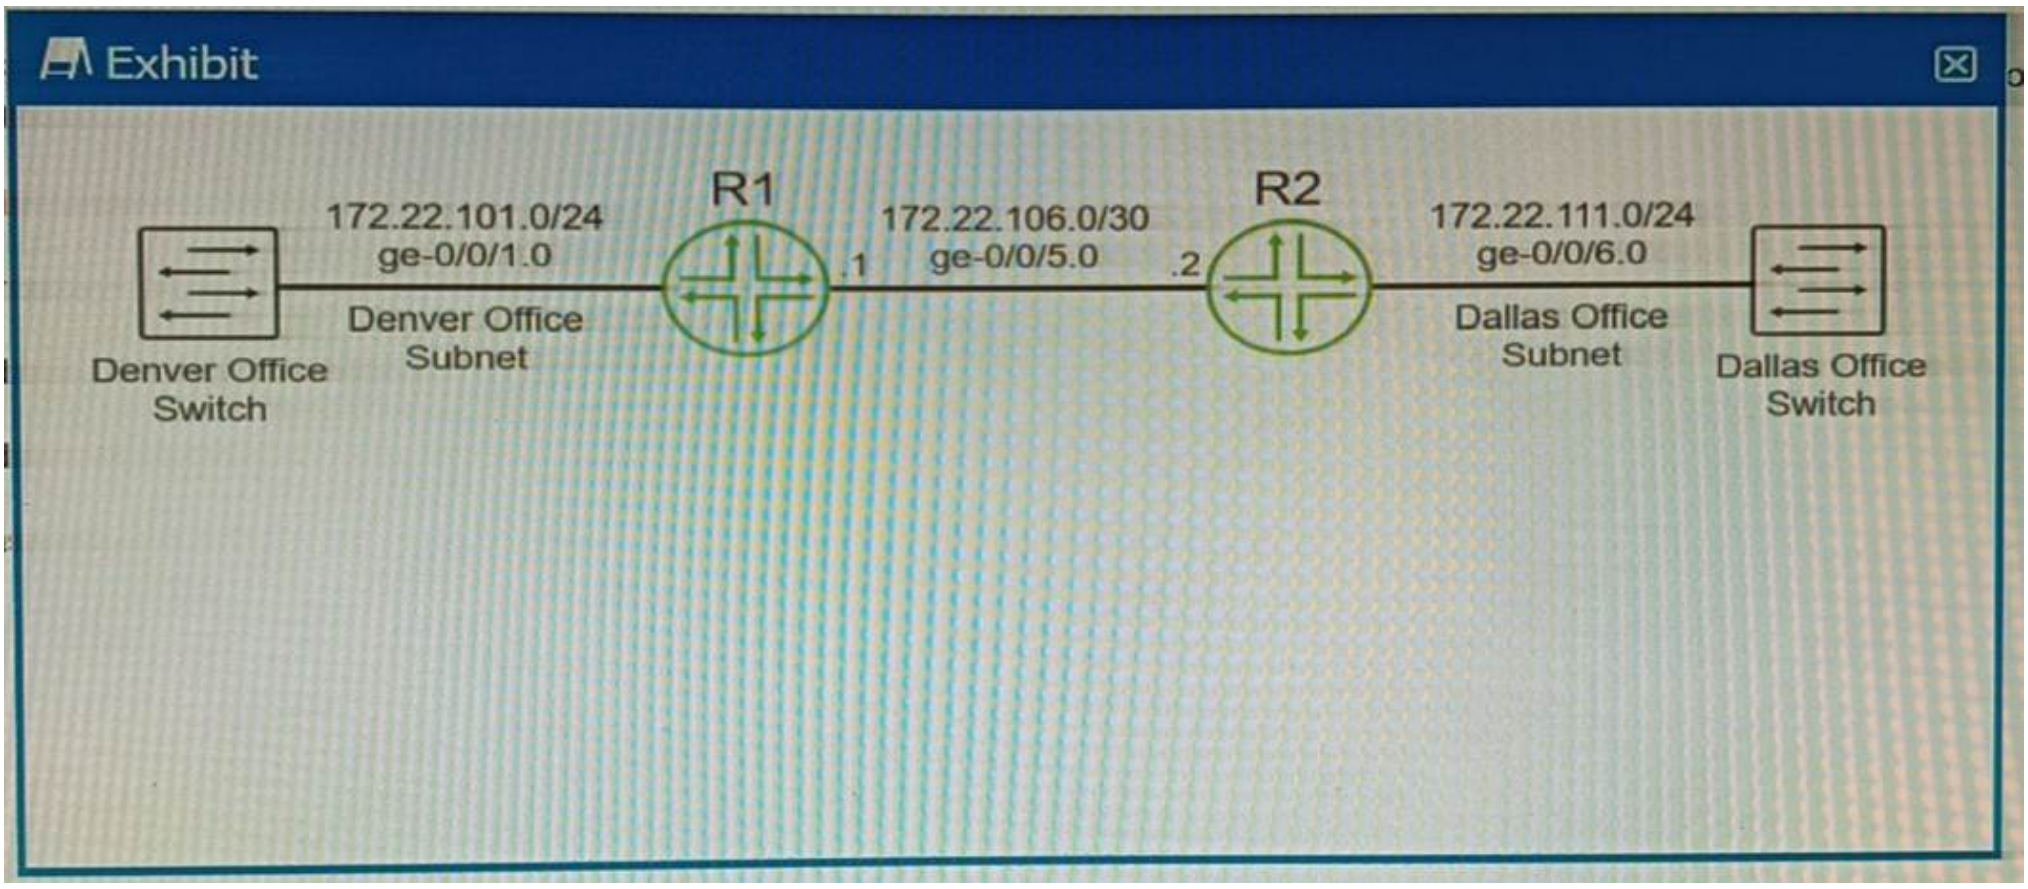

You are using OSPF to advertise the subnets that are used by the Denver and Dallas offices. The routers that are directly connected to the Dallas and Denver subnets are not advertising the connected subnets.
Referring to the exhibit, which two statements are correct? (Choose two.)

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

Answer: CD

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

NEW QUESTION 53
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 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 60
Which statement is correct about graceful Routing Engine switchover (GRES)?

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

Answer: C

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

NEW QUESTION 61
Which two statements are correct about using firewall filters on EX Series switches? (Choose two.)

A. You can deploy only stateless firewall filters on an EX Series switch.
B. You can only apply firewall filters to Layer 2 traffic on an EX Series switch.
C. You can apply firewall filters to both Layer 2 and Layer 3 traffic on an EX Series switch.
D. You can deploy both stateless and stateful firewall filters on an EX Series switch.

Answer: AC

Explanation:
? A is correct because you can deploy only stateless firewall filters on an EX Series switch. A stateless firewall filter is a filter that evaluates each packet individually based on the header information, such as source and destination addresses, protocol, and port numbers1. A stateless firewall filter does not keep track of the state or context of a packet flow, such as the sequence number, flags, or sessioninformation1. EX Series switches support only stateless firewall filters, which are also called access control lists (ACLs) or packet filters2.
? C is correct because you can apply firewall filters to both Layer 2 and Layer 3 traffic on an EX Series switch. Layer 2 traffic is traffic that is switched within a VLAN or a bridge domain, while Layer 3 traffic is traffic that is routed between VLANs or networks3. EX Series switches support three types of firewall filters: port (Layer 2) firewall filters, VLAN firewall filters, and router (Layer 3) firewall filters4. You can apply these filters to different interfaces and directions to control the traffic
entering or exiting the switch.

NEW QUESTION 62
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 65
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 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