ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

Exhibit.

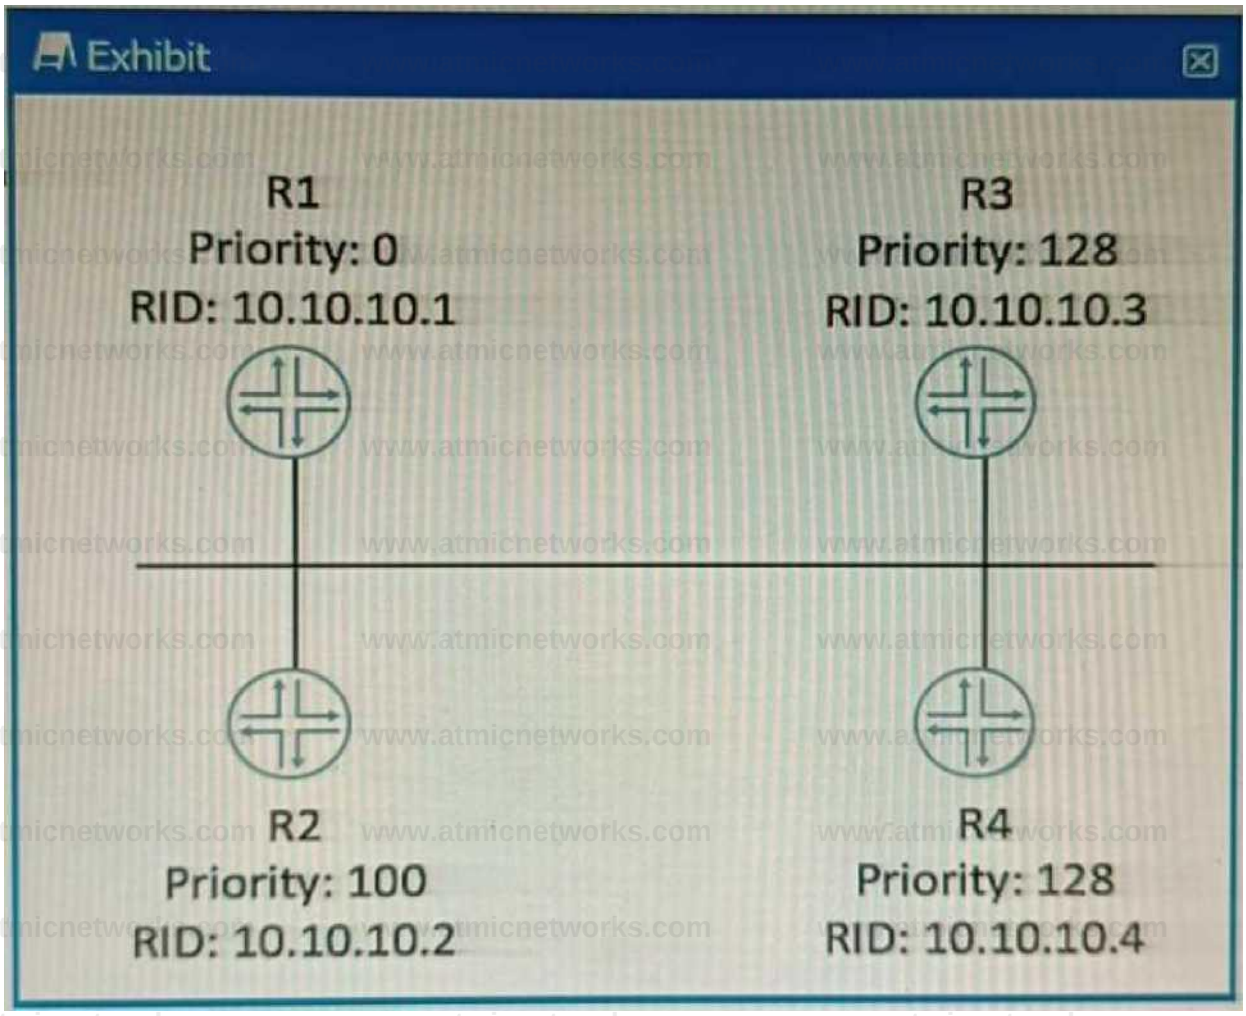

Which router will become the OSPF BDR if all routers are powered on at the same time?

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 multiaccess 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.

Reference:

1: OSPF DR/BDR Election: Process, Configuration, and Tuning 2: OSPF Designated Router (DR) and Backup Designated Router (BDR)

Question: 2

Exhibit.

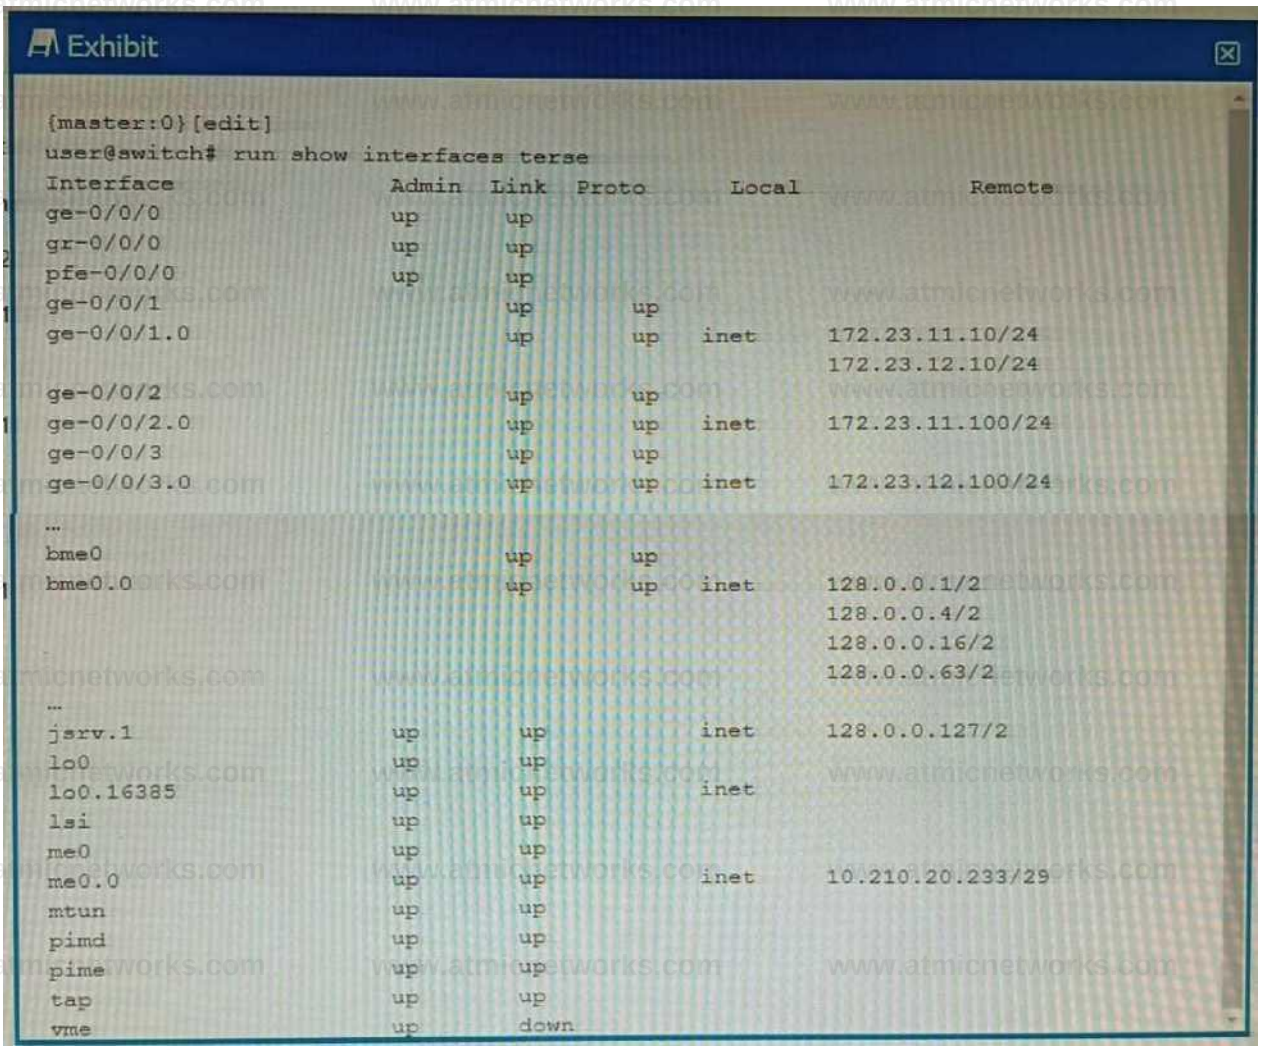

What is the management IP address of the device shown in the exhibit?

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 dedicated management 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 is me0. This interface has an admin status of up, a protocol status of inet, a local address of 172.23.12.100/24, and a remote address of unspecified. 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 is 172.23.12.100.

Reference:

: [Management Interfaces Overview] : [Displaying Interface Status Information]

Question: 3

Which three protocols support BFD? (Choose three.)

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

Answer: BCD

Explanation:

BFD is a protocol that can be used to quickly detect failures in the forwarding path between two adjacent routers or switches. BFD can be integrated with various routing protocols and link aggregation protocols to provide faster convergence and fault recovery.

According to the Juniper Networks documentation, the following protocols support BFD on Junos OS devices1:

BGP: BFD can be used to monitor the connectivity between BGP peers and trigger a session reset if a failure is detected. BFD can be configured for both internal and external BGP sessions, as well as for IPv4 and IPv6 address families2.

OSPF: BFD can be used to monitor the connectivity between OSPF neighbors and trigger a state change if a failure is detected. BFD can be configured for both OSPFv2 and OSPFv3 protocols, as well as for point-to-point and broadcast network types3.

LACP: BFD can be used to monitor the connectivity between LACP members and trigger a link state change if a failure is detected. BFD can be configured for both active and passive LACP modes, as well as for static and dynamic LAGs4.

Other protocols that support BFD on Junos OS devices are:

IS-IS: BFD can be used to monitor the connectivity between IS-IS neighbors and trigger a state change if a failure is detected. BFD can be configured for both level 1 and level 2 IS-IS adjacencies, as well as for point-to-point and broadcast network types.

RIP: BFD can be used to monitor the connectivity between RIP neighbors and trigger a route update if a failure is detected. BFD can be configured for both RIP version 1 and version 2 protocols, as well as for IPv4 and IPv6 address families.

VRRP: BFD can be used to monitor the connectivity between VRRP routers and trigger a priority change if a failure is detected. BFD can be configured for both VRRP version 2 and version 3 protocols, as well as for IPv4 and IPv6 address families.

The protocols that do not support BFD on Junos OS devices are:

RSTP: RSTP is a spanning tree protocol that provides loop prevention and rapid convergence in layer 2 networks. RSTP does not use BFD to detect link failures, but relies on its own hello mechanism that sends BPDU packets every 2 seconds by default.

FTP: FTP is an application layer protocol that is used to transfer files between hosts over a TCP connection. FTP does not use BFD to detect connection failures, but relies on TCP's own retransmission and timeout mechanisms.

Reference:

1: [Configuring Bidirectional Forwarding Detection] 2: [Configuring Bidirectional Forwarding Detection for BGP] 3: [Configuring Bidirectional Forwarding Detection for OSPF] 4: [Configuring Bidirectional Forwarding Detection for Link

Aggregation Control Protocol] : [Configuring Bidirectional Forwarding Detection for IS-IS] : [Configuring Bidirectional Forwarding Detection for RIP] : [Configuring Bidirectional Forwarding Detection for VRRP] : [Understanding Rapid Spanning Tree Protocol] : [Understanding FTP]

Question: 4

Exhibit.

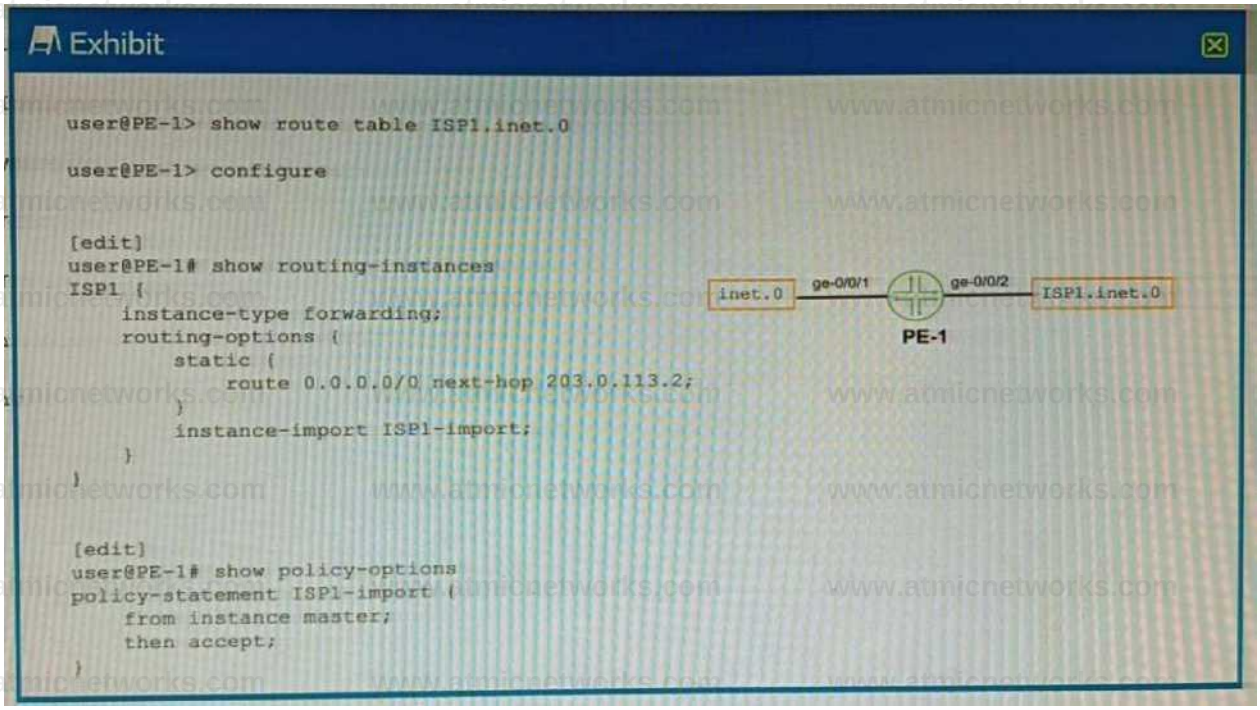

The ispi _ inet. 0 route table has currently no routes in it.

What will happen when you commit the configuration shown on the exhibit?

A. The inet. 0 route table will be completely overwritten by the ispi . inet. 0 route table.
B. The inet. 0 route table will be imported into the ispi . inet. 0 route table.
C. The ISPI . inet. 0 route table will be completely overwritten by the inet. o route table.
D. The ISPI . inet. 0 route table will be imported into the inet. 0 route table.

Answer: B

Explanation:

The configuration shown in the exhibit is an example of a routing instance of type virtual-router. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters that create a separate routing domain on a Juniper device1. A virtual-router routing instance allows administrators to divide a device into multiple independent virtual routers, each with its own routing table2.

The configuration also includes a rib-group statement, which is used to import routes from one routing table to another. A rib-group consists of an import-rib statement, which specifies the source routing table, and an export-rib statement, which specifies the destination routing table.

In this case, the rib-group name is inet-to-ispi, and the import-rib statement specifies inet.0 as the source routing table. The export-rib statement specifies ispi.inet.0 as the destination routing table.

This means that the routes from inet.0 will be imported into ispi.inet.0.

Therefore, the correct answer is B. The inet.0 route table will be imported into the ispi.inet.0 route table.

Reference:

1: Routing Instances Overview 2: Virtual Routing Instances : [rib-group (Routing Options)]

Question: 5

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.

Question: 6

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.

Reference:

1: rib-groups | Junos OS | Juniper Networks

Question: 7

Exhibit.

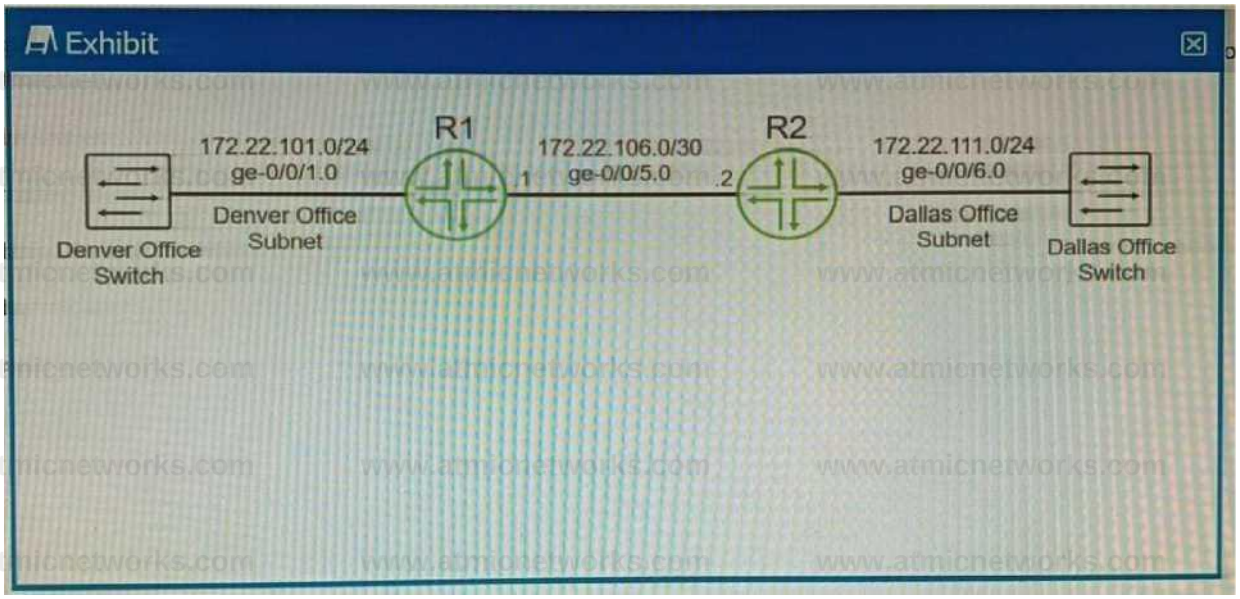

You are using OSPF to advertise the subnets that are used by the Denver and Dallas offices. The routers that are directly connected to the Dallas and Denver subnets are not advertising the connected subnets.

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.

Question: 8

Exhibit.

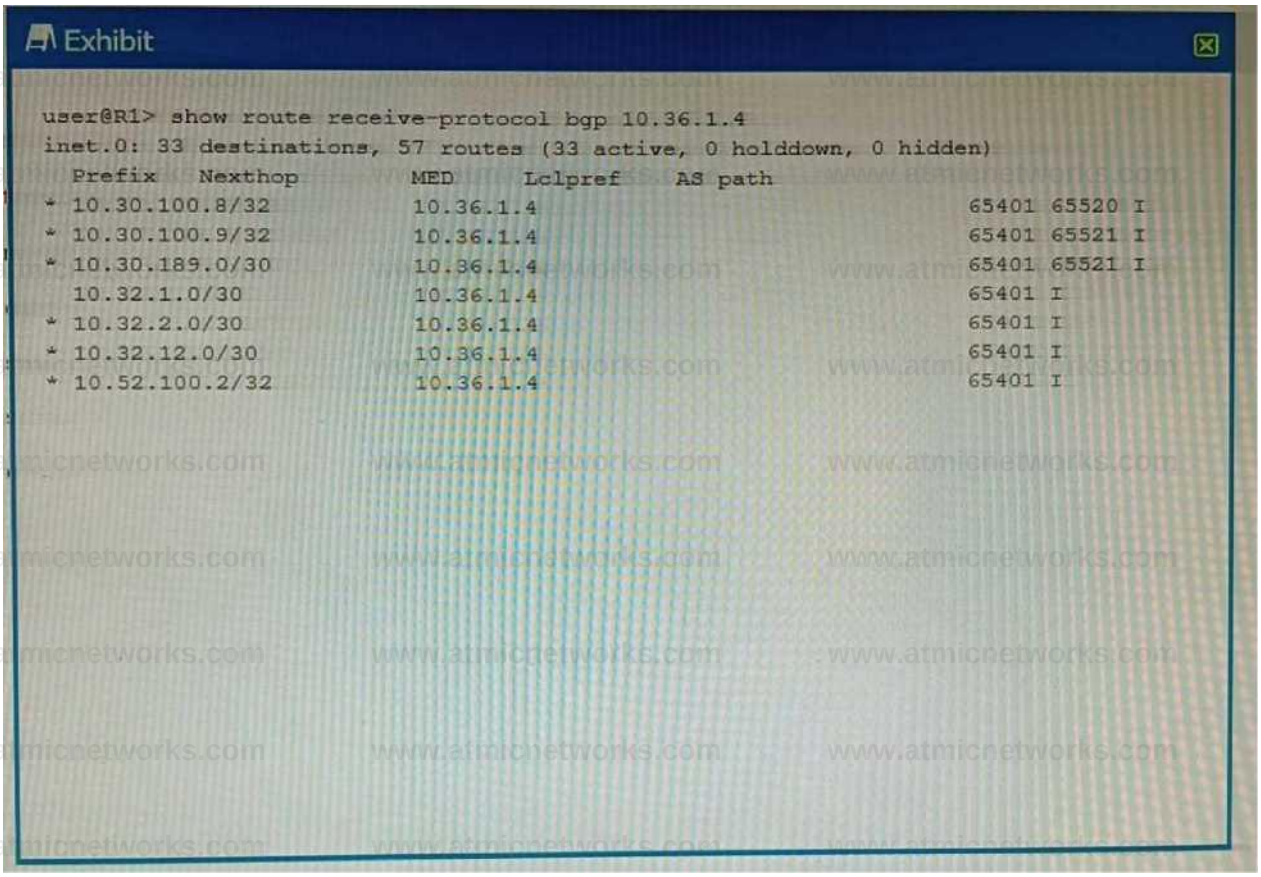

You want to verify prefix information being sent from 10.36.1.4.
Which two statements are correct about the output shown in the exhibit? (Choose two.)

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.

Question: 9

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.

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.

Reference:

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]

Question: 11

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.

Reference:

1: IP-IP Tunneling 2: What is tunneling? | Tunneling in networking 3: IPv4 - Header

Question: 12

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

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

Answer: AD

Explanation:

Option A suggests that the MTU is not at least 1492 bytes. This is correct because IS-IS requires a minimum MTU of

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

Question: 13

You are asked to create a new firewall filter to evaluate Layer 3 traffic that is being sent between VLANs. In this scenario, which two statements are correct? (Choose two.)

A. You should create a family Ethernet-switching firewall filter with the appropriate match criteria and actions.
B. You should apply the firewall filter to the appropriate VLAN.
C. You should create a family inet firewall filter with the appropriate match criteria and actions.
D. You should apply the firewall filter to the appropriate IRB interface.

Answer: CD

Explanation:

A firewall filter is a configuration that defines the rules that determine whether to forward or discard packets at specific processing points in the packet flow. A firewall filter can also modify the attributes of the packets, such as priority, marking, or logging. A firewall filter can be applied to various interfaces, protocols, or routing instances on a Juniper device1.
A firewall filter has a family attribute, which specifies the type of traffic that the filter can evaluate. The family attribute can be one of the following: inet, inet6, mpls, vpls, iso, or ethernet- switching2. The family inet firewall filter is used to evaluate IPv4 traffic, which is the most common type of Layer 3 traffic on a network.
To create a family inet firewall filter, you need to specify the appropriate match criteria and actions for each term in the filter. The match criteria can include various fields in the IPv4 header, such as source address, destination address, protocol, port number, or DSCP value. The actions can include accept, discard, reject, count, log, policer, or next term3.
To apply a firewall filter to Layer 3 traffic that is being sent between VLANs, you need to apply the filter to the appropriate IRB interface. An IRB interface is an integrated routing and bridging interface that provides Layer 3 functionality for a VLAN on a Juniper device. An IRB interface has an IP address that acts as the default gateway for the hosts in the VLAN. An IRB interface can also participate in routing protocols and forward packets to other VLANs or networks4.
Therefore, option C is correct, because you should create a family inet firewall filter with the appropriate match criteria and actions. Option D is correct, because you should apply the firewall filter to the appropriate IRB interface.
Option A is incorrect, because you should not create a family ethernet-switching firewall filter with the appropriate match criteria and actions. A family ethernet-switching firewall filter is used to evaluate Layer 2 traffic on a Juniper device. A family ethernet-switching firewall filter can only match on MAC addresses or VLAN IDs, not on IP addresses or protocols5.
Option B is incorrect, because you should not apply the firewall filter to the appropriate VLAN. A VLAN is a logical grouping of hosts that share the same broadcast domain on a Layer 2 network. A VLAN does not have an IP address or routing capability. A firewall filter cannot be applied directly to a VLAN; it must be applied to an interface that belongs to or connects to the VLAN6.

Reference:

1: Firewall Filters Overview 2: Configuring Firewall Filters 3: Configuring Firewall Filter Match Conditions and Actions 4: Understanding Integrated Routing and Bridging Interfaces 5: Configuring Ethernet-Switching Firewall Filters 6: Understanding VLANs

Question: 14

Exhibit

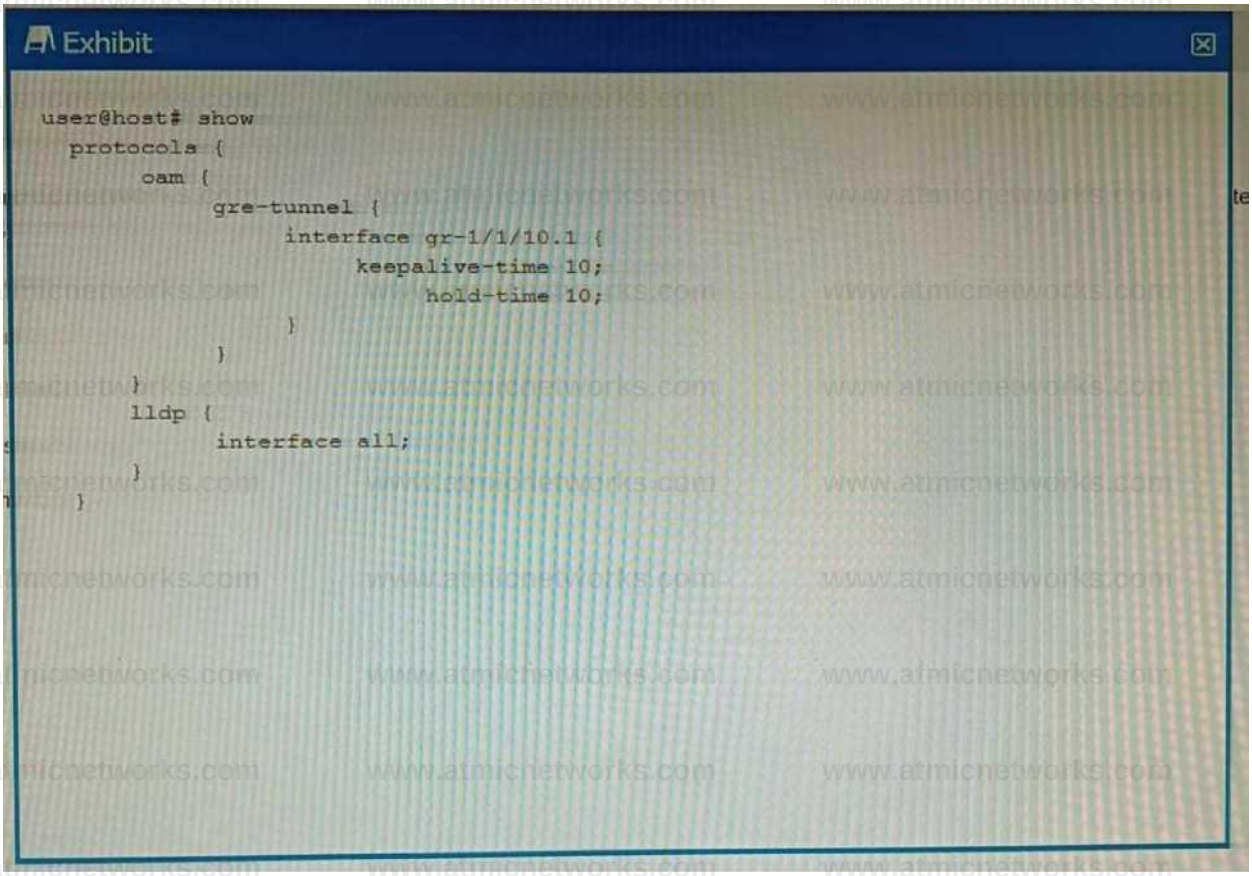

You have configured a GRE tunnel. To reduce the risk of dropping traffic, you have configured a keepalive OAM probe to monitor the state of the tunnel; however, traffic drops are still occurring. Referring to the exhibit, what is the problem?

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

Answer: D

Explanation:

A keepalive OAM probe is a mechanism that can be used to monitor the state of a GRE tunnel and detect any failures in the tunnel path. A keepalive OAM probe consists of sending periodic packets from one end of the tunnel to the other and expecting a reply. If no reply is received within a specified time, the tunnel is considered down and the line protocol of the tunnel interface is changed to down1.

To configure a keepalive OAM probe for a GRE tunnel, you need to specify two parameters: the keepalive-time and the hold-time. The keepalive-time is the interval between each keepalive packet sent by the local router. The hold-time is the maximum time that the local router waits for a reply from the remote router before declaring the tunnel down2.

According to the Juniper Networks documentation, the hold-time value must be two times the keepalive-time value for a GRE tunnel2. This is because the hold-time value must account for both the round-trip time of the keepalive packet and the processing time of the remote router. If the holdtime value is too small, it may cause false positives and unnecessary tunnel flaps.

In the exhibit, the configuration shows that the keepalive-time is set to 10 seconds and the hold-time is set to 15 seconds for the gr-1/1/10.1 interface. This means that the local router will send a keepalive packet every 10 seconds and will wait for 15 seconds for a reply from the remote router. However, this hold-time value is not two

times the keepalive-time value, which violates the recommended configuration. This may cause traffic drops if the remote router takes longer than 15 seconds to reply.

Therefore, option D is correct, because the hold-time value must be two times the keepalive-time value for a GRE tunnel. Option A is incorrect, because BFD is not required for GRE tunnels; BFD is another protocol that can be used to monitor tunnels, but it is not compatible with GRE keepalives3. Option B is incorrect, because the "event link-adjacency-loss" option is not related to GRE tunnels; it is an option that can be used to trigger an action when a link goes down4. Option C is incorrect, because LLDP does not need to be removed from the gr-1/1/10.1 interface; LLDP is a protocol that can be used to discover neighboring devices and their capabilities, but it does not interfere with GRE tunnels5.

Reference:

1: Configuring Keepalive Time and Hold time for a GRE Tunnel Interface 2: keepalive | Junos OS | Juniper Networks 3: Configuring Bidirectional Forwarding Detection 4: event link-adjacency-loss | Junos OS | Juniper Networks 5: Understanding Link Layer Discovery Protocol

Question: 15

Exhibit

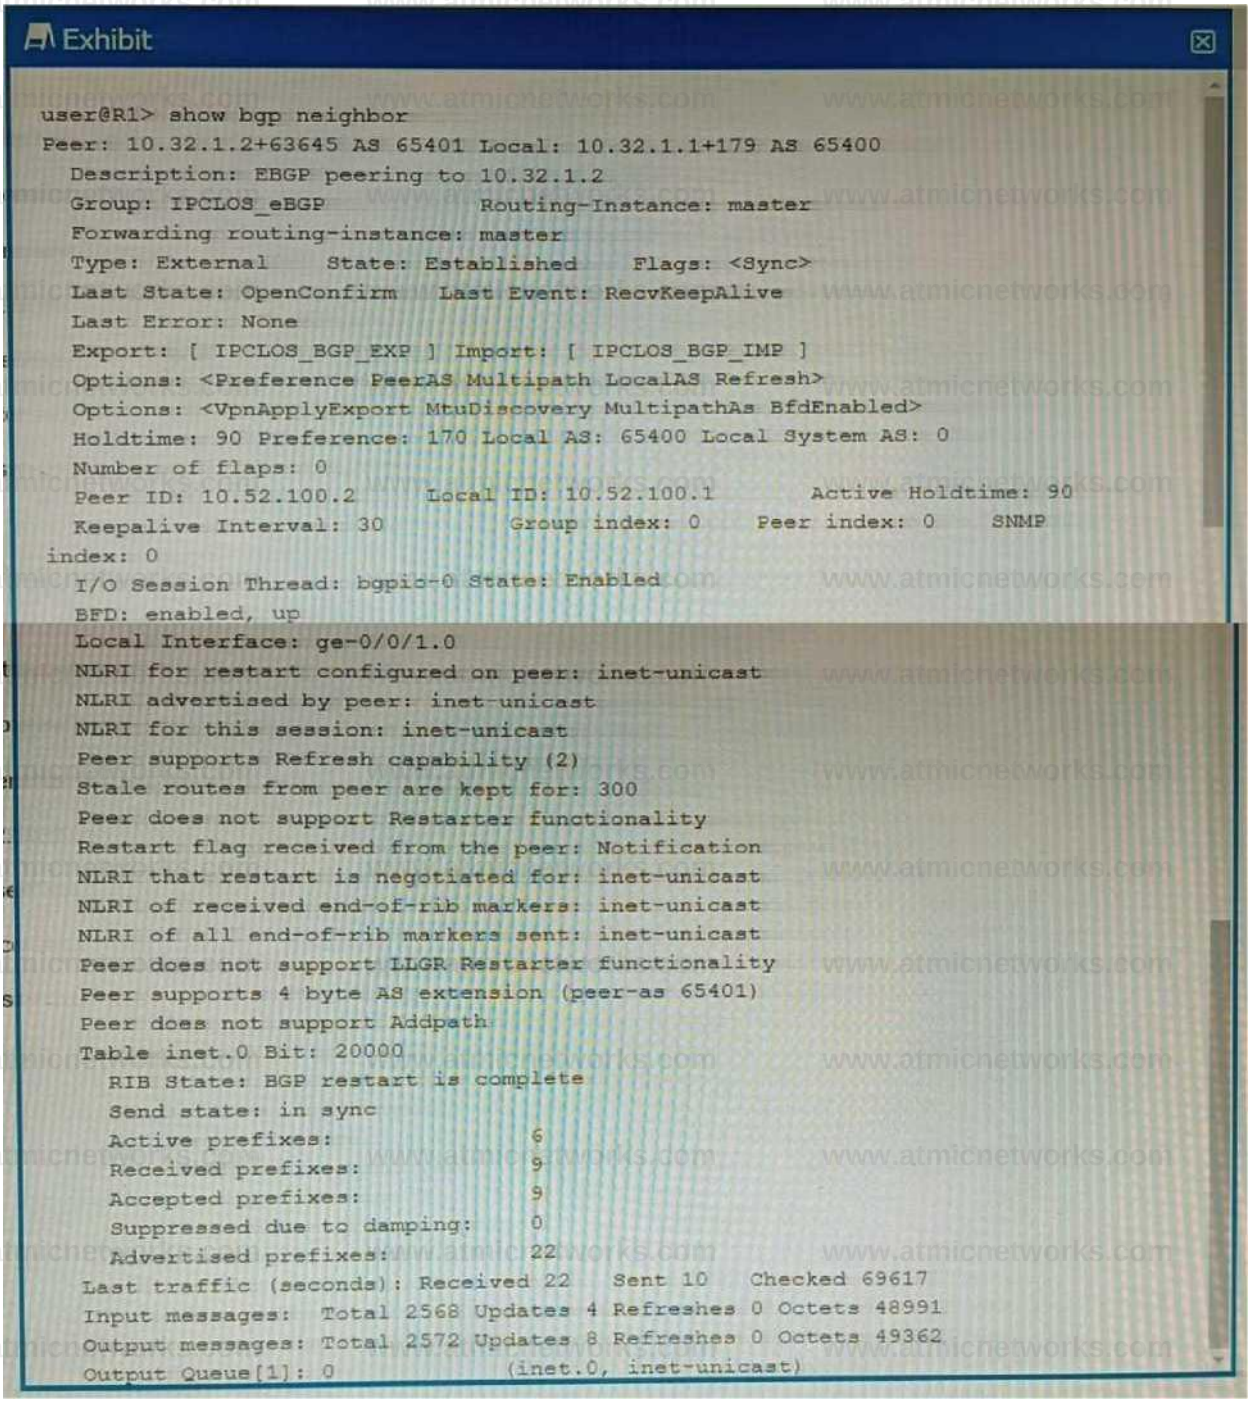

You are a network operator troubleshooting BGP connectivity.
Which two statements are correct about the output shown in the exhibit? (Choose two.)

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.

Question: 16

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

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

Answer: D

Explanation:

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

Question: 17

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.

Question: 18

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

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

Answer: CD

Explanation:

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

Question: 19

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.

Question: 20

Exhibit

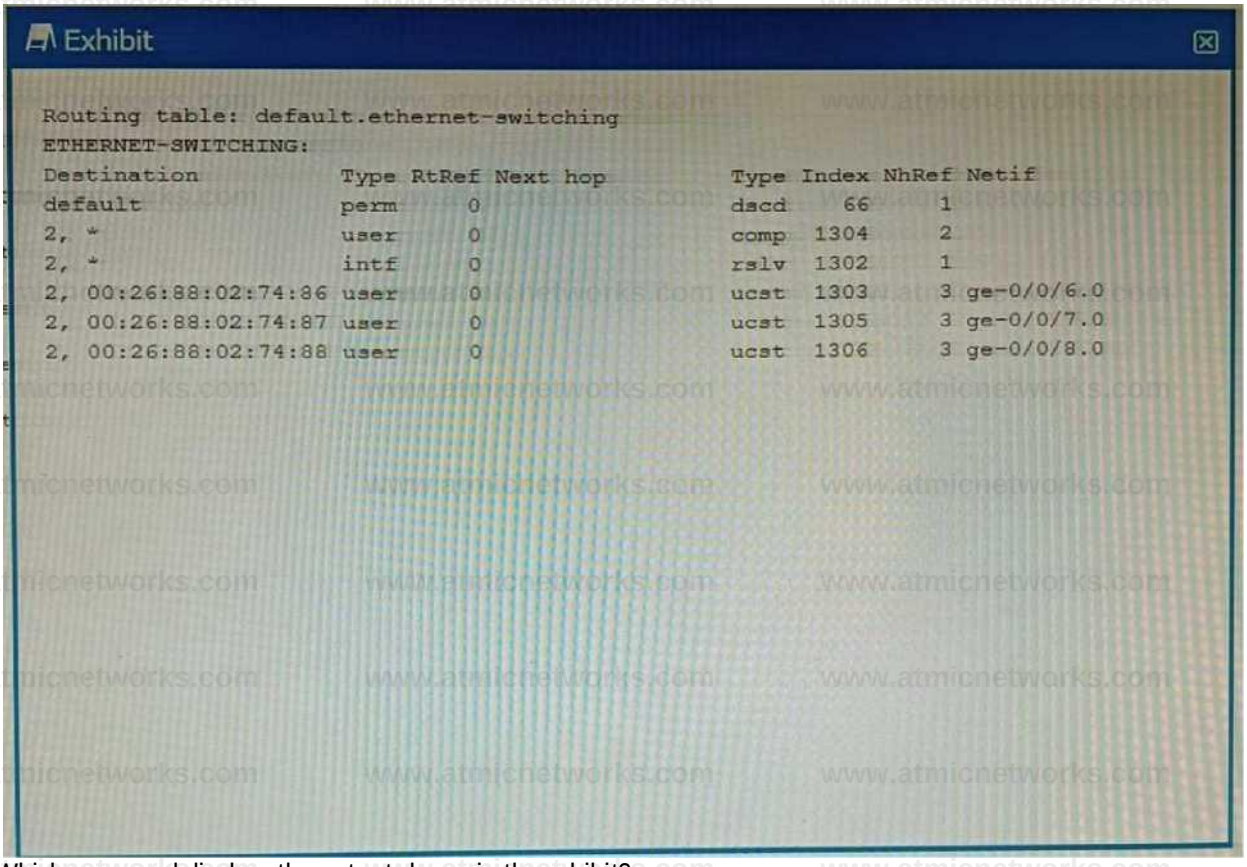

Which command displays the output shown in the exhibit?

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

Answer: B

Explanation:

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

ethernet-switching protocol family, which shows the destination MAC address, the next-hop MAC address, and the interface name3. It does not show the VLAN name or the forwarding state.

Question: 21

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

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

Answer: AC

Explanation:

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

Question: 22

Exhibit

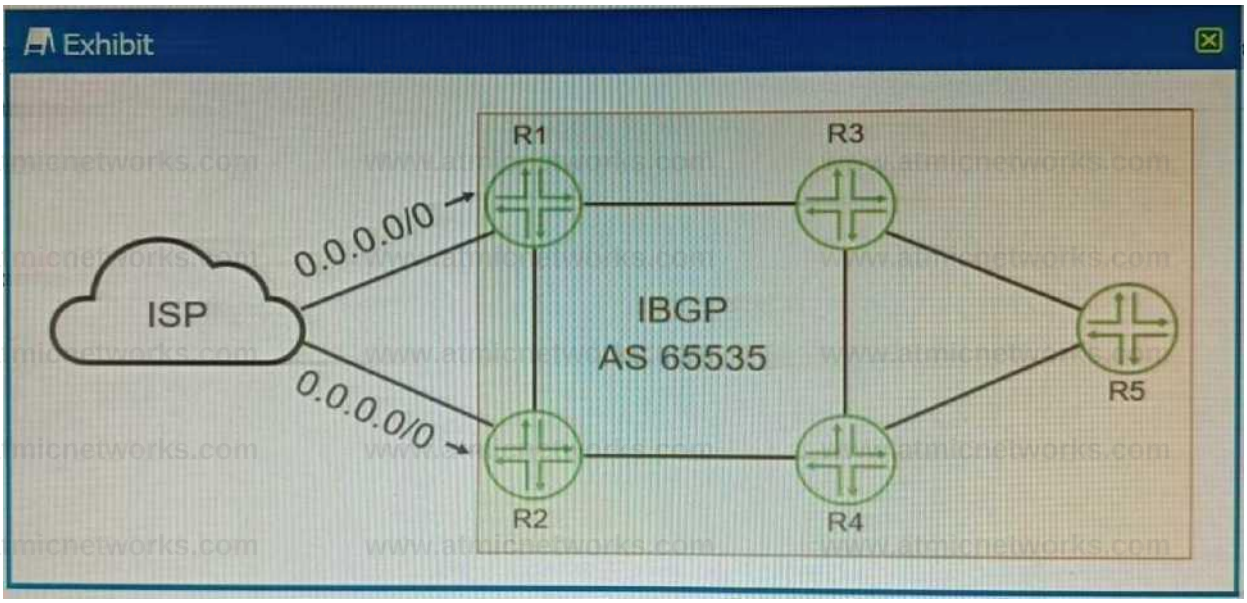

Your ISP is announcing a default route to both R1 and R2. You want your network routers to forward all Internet traffic through the R1 device Which BGP attribute would you use?

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.

Question: 23

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.

Question: 24

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

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

Answer: AC

Explanation:

BGP attributes are properties that BGP uses for route advertisement, path selection, and loop prevention1. There are four categories of BGP attributes123:

Well-known mandatory: Must be recognized by all BGP routers, present in all BGP updates, and passed on to other BGP routers123.

Well-known discretionary: Supported by all BGP implementations, and are optionally included in BGP updates1.

Optional transitive: May not be supported by all implementations of BGP1.

Optional non-transitive: May not be supported by all implementations of BGP1.

The well-known mandatory attributes must be supported by all BGP implementations and must be included in every update123. These include the AS path and next hop attributes23. Therefore, options A and C are correct.

Question: 25

Exhibit

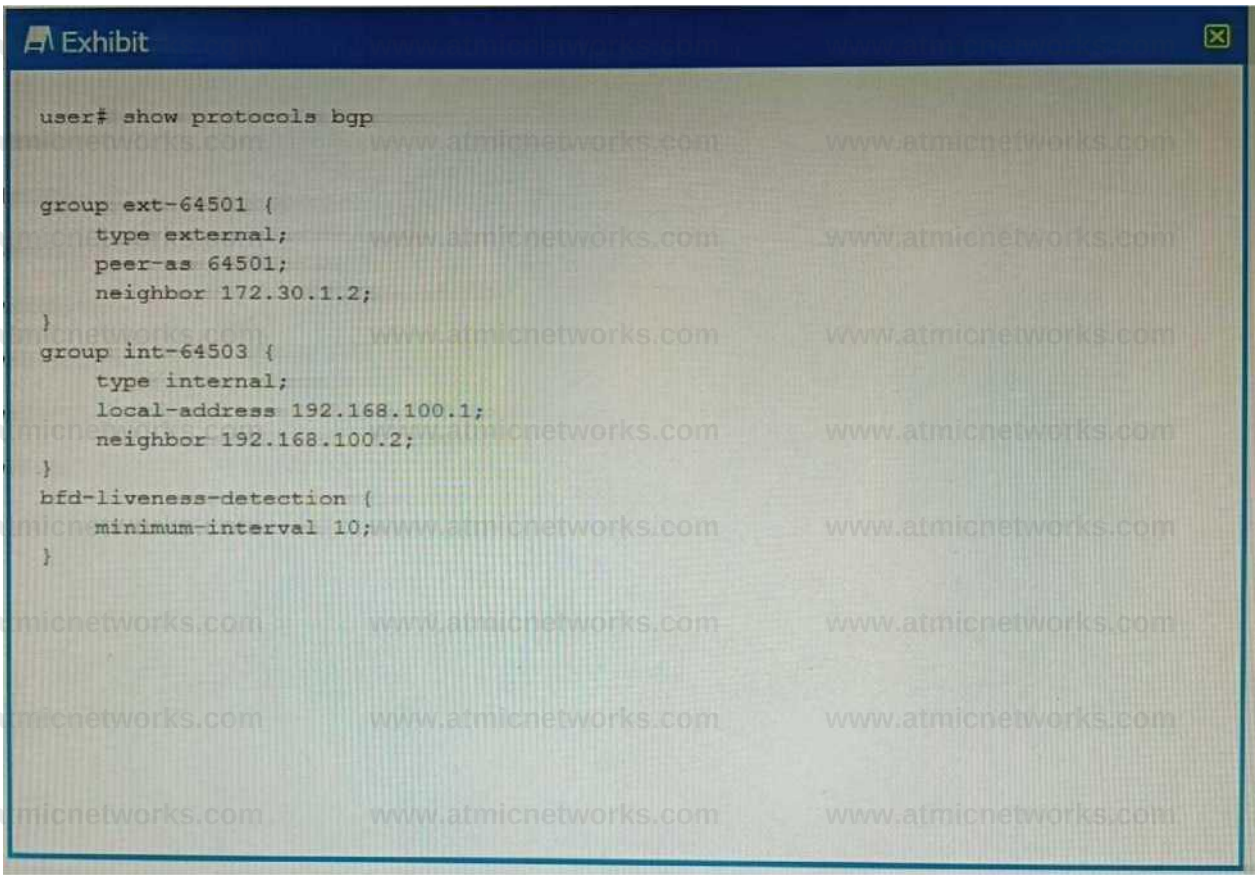

Your BGP neighbors, one in the USA and one in France, are not establishing a connection with each other. Referring to the exhibit, which statement is correct?

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 BGP neighbors, 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.

Question: 26

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.

Question: 27

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.

Question: 28

Exhibit

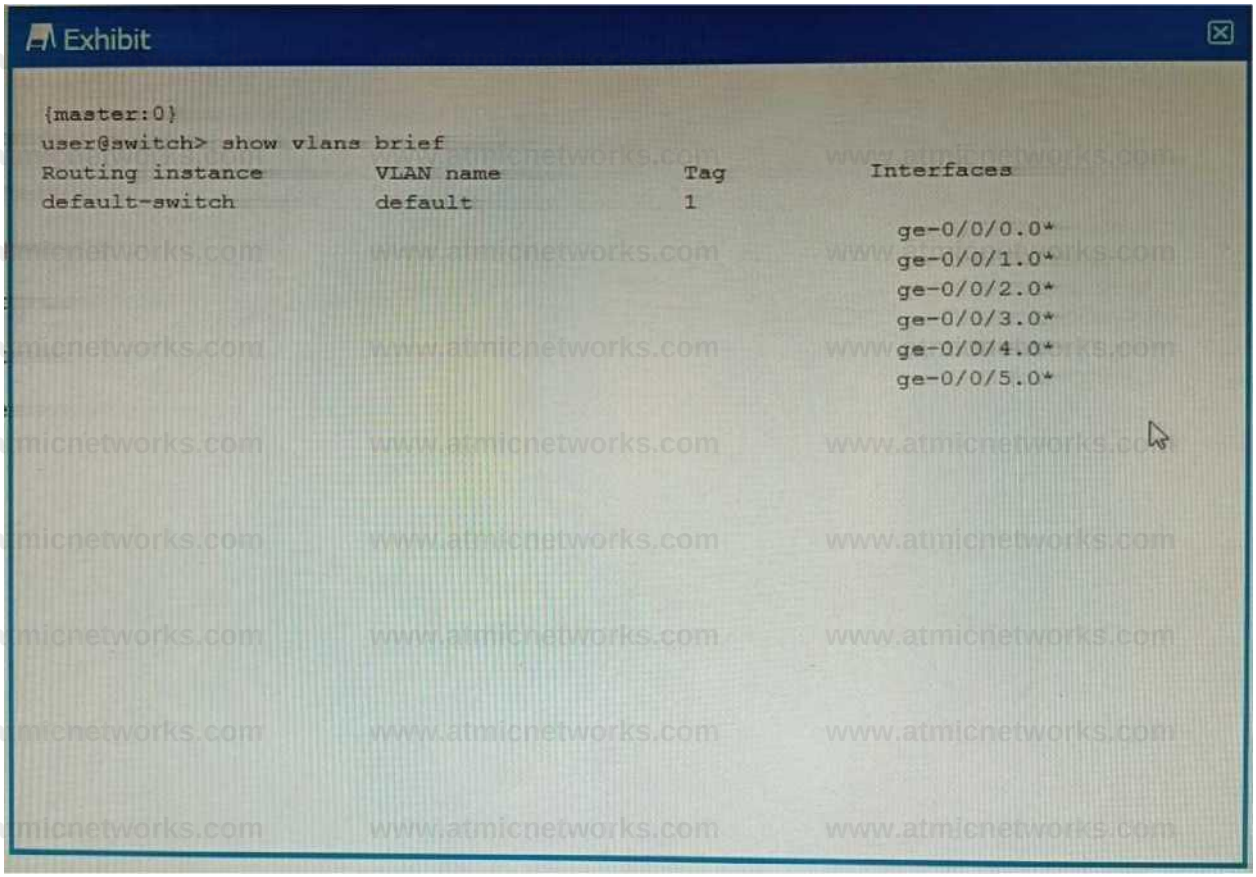

What does the * indicate in the output shown in the exhibit?

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

Answer: C

Explanation:

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

Question: 29

You have two OSPF routers forming an adjacency. R1 has a priority of 32 and a router ID of 192.168.1.2. R2 has a priority of 64 and a router ID of 192.168.1.1. The routers were started at the same time and all other OSPF settings are the default settings.
Which statement is correct in this scenario?

A. At least three routers are required for a DR/BDR election B. Router IDs must match for an adjacency to form.

C. R2 will be the BDR.

D. R1 will be the BDR.

Answer: D

Explanation:

In OSPF, the Designated Router (DR) and Backup Designated Router (BDR) are elected based on the priority of the routers1. The router with the highest priority becomes the DR, and the router with the second highest priority becomes the BDR1. If there is a tie in priority, then the router with the highest Router ID is chosen1.
In this scenario, R2 has a higher priority (64) than R1 (32), so R2 will become the DR1. Since R1 has the second highest priority, it will become the BDR1. Therefore, option D is correct.

Question: 30

Exhibit.

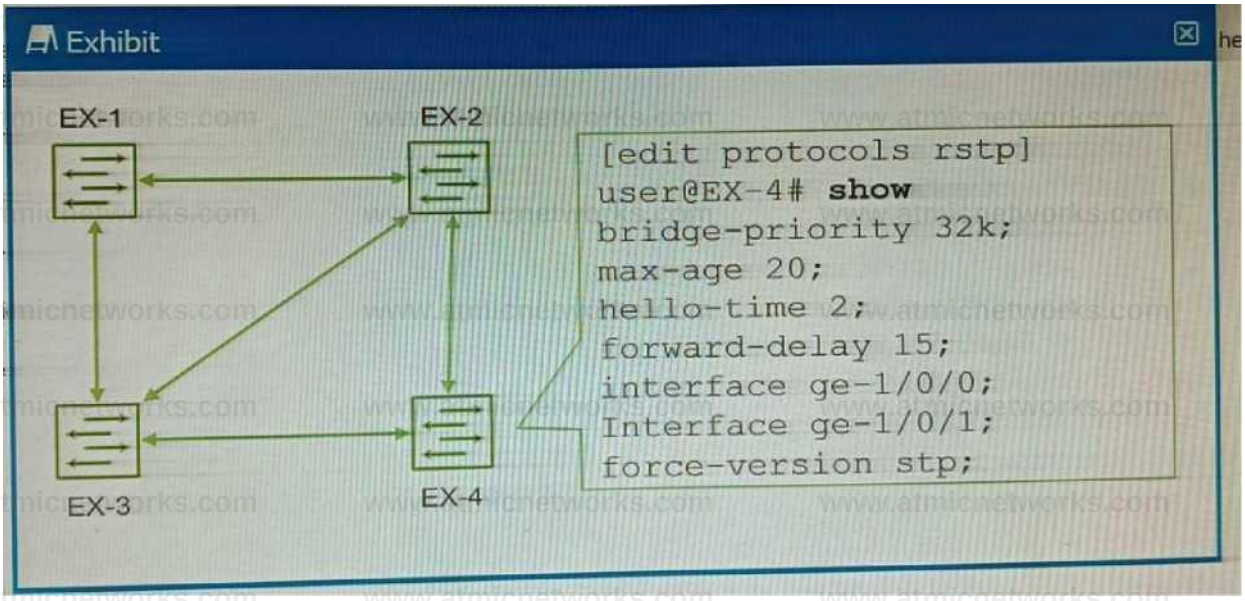

You have configured the four EX Series switches with RSTP, as shown in the exhibit. You discover that whenever a link between switches goes up or down, the switches take longer than expected for RSTP to converge, using the default settings.
In this scenario, which action would solve the delay in RSTP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

Question: 31

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.

Question: 32

Exhibit

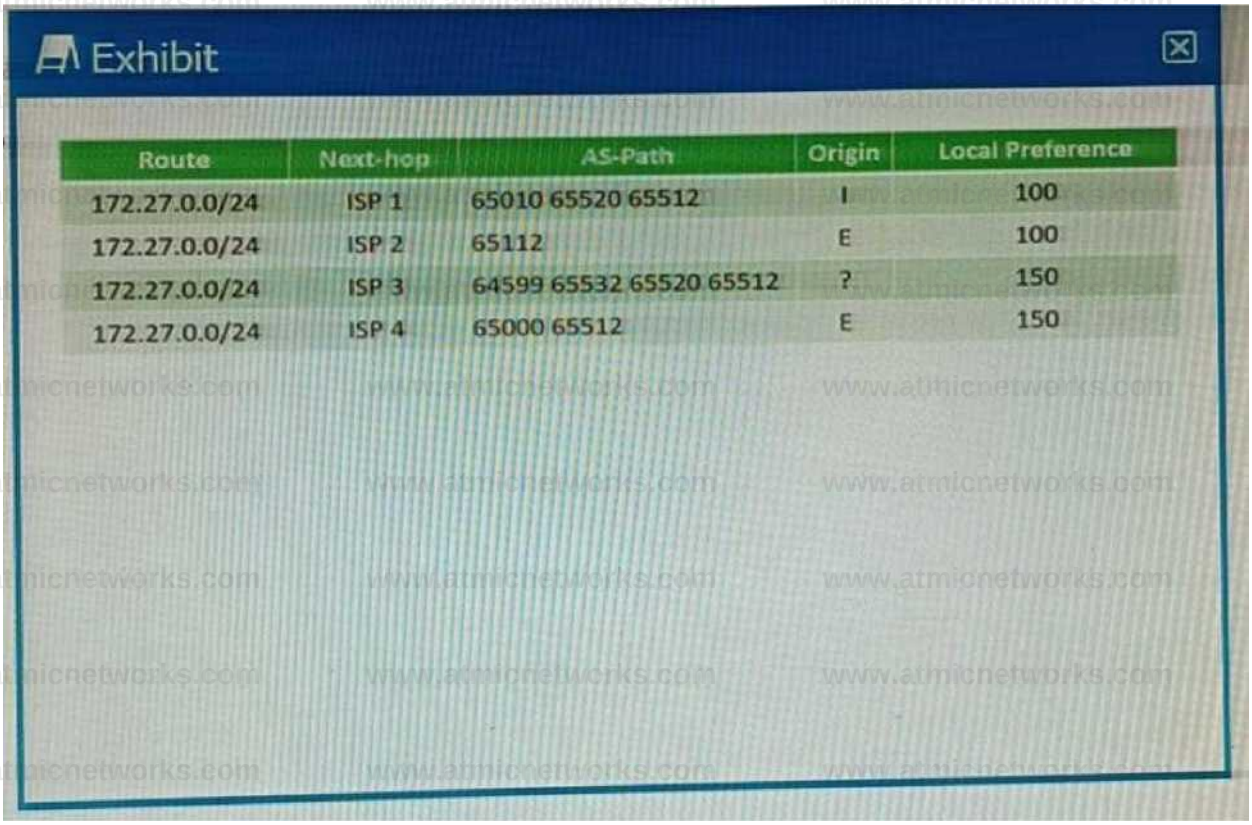

You are receiving the BGP route shown in the exhibit from four different upstream ISPs. Referring to the exhibit, which ISP will be selected as the active path?

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.

Question: 33

Exhibit.

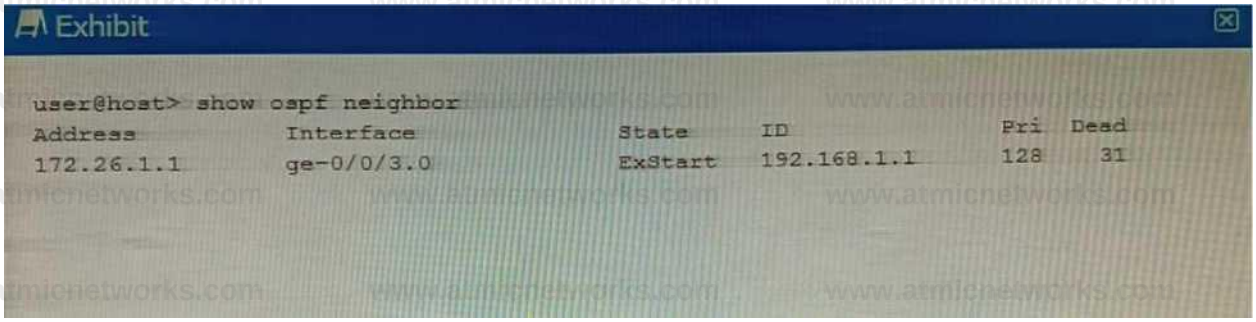

Why is this OSPF adjacency remaining in this state?

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 2Way 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, they may 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.

Question: 34

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

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

Answer: CDE

Explanation:

IS-IS (Intermediate System to Intermediate System) is a routing protocol that is designed to move information efficiently within a computer network12. It supports multiple protocol topologies, including IPv4, IPv6, and

multicast12. Therefore, options C, E, and D are correct.

Question: 35

Refer to the exhibit.

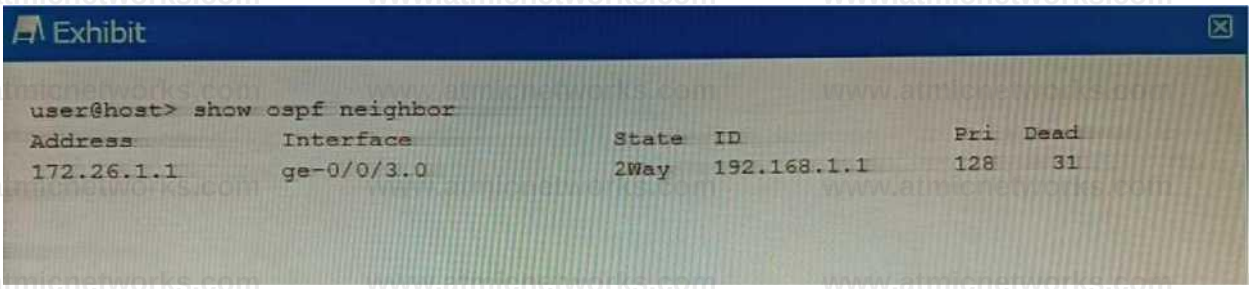

Referring to the output shown in the exhibit, which statement is correct?

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.

Question: 36

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.

Question: 37

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

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

Answer: BCD

Explanation:

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

Question: 38

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:

The router with the highest OSPF priority becomes the DR. The default priority is 1, and a priority of 0 means the router will not participate in the election.
If there is a tie in priority, the router with the highest router ID becomes the DR. The router ID is a 32- bit number that uniquely identifies a router in an OSPF domain. It can be manually configured or automatically derived from the highest IP address of a loopback interface or a physical interface.

If there is a tie in router ID, the router that was first to become an OSPF neighbor becomes the DR.
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.

Question: 39

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

A. Redundant trunk groups load-balance traffic across two designated uplink interfaces.
B. If the active link fails, then the secondary link automatically takes over.
C. Layer 2 control traffic is permitted on the secondary link
D. Redundant trunk groups must be connected to the same aggregation switch.

Answer: BD

Explanation:

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

Question: 40

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.

Question: 41

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.

Question: 42

Which statement is correct about the IS-IS ISO NET address?

A. An ISO NET address defined with a system ID of 0000.0000.0000 must be selected as the DIS.
B. An ISO NET address must be unique for each device in the network.

C. You can only define a single ISO NET address per device.
D. The Area ID must match on all devices within a L2 area.

Answer: B

Explanation:

An ISO NET address is a type of network address used by the IS-IS routing protocol. It identifies a point of connection to the network, such as a router interface, and is also called a Network Service Access Point (NSAP)1.
An ISO NET address consists of three parts: an area ID, a system ID, and a selector2. The area ID identifies the IS-IS area to which the device belongs. The system ID uniquely identifies the device within the area. The selector identifies a specific service or function on the device, such as routing or management2.
An ISO NET address must be unique for each device in the network, because it is used by IS-IS to establish adjacencies, exchange routing information, and compute shortest paths2. If two devices have the same ISO NET address, they will not be able to communicate with each other or with other devices in the network. Therefore, it is important to assign different ISO NET addresses to each device in the network.

Question: 43

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.

Question: 44

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.

Question: 45

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.

Question: 46

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.

Question: 47

Exhibit.

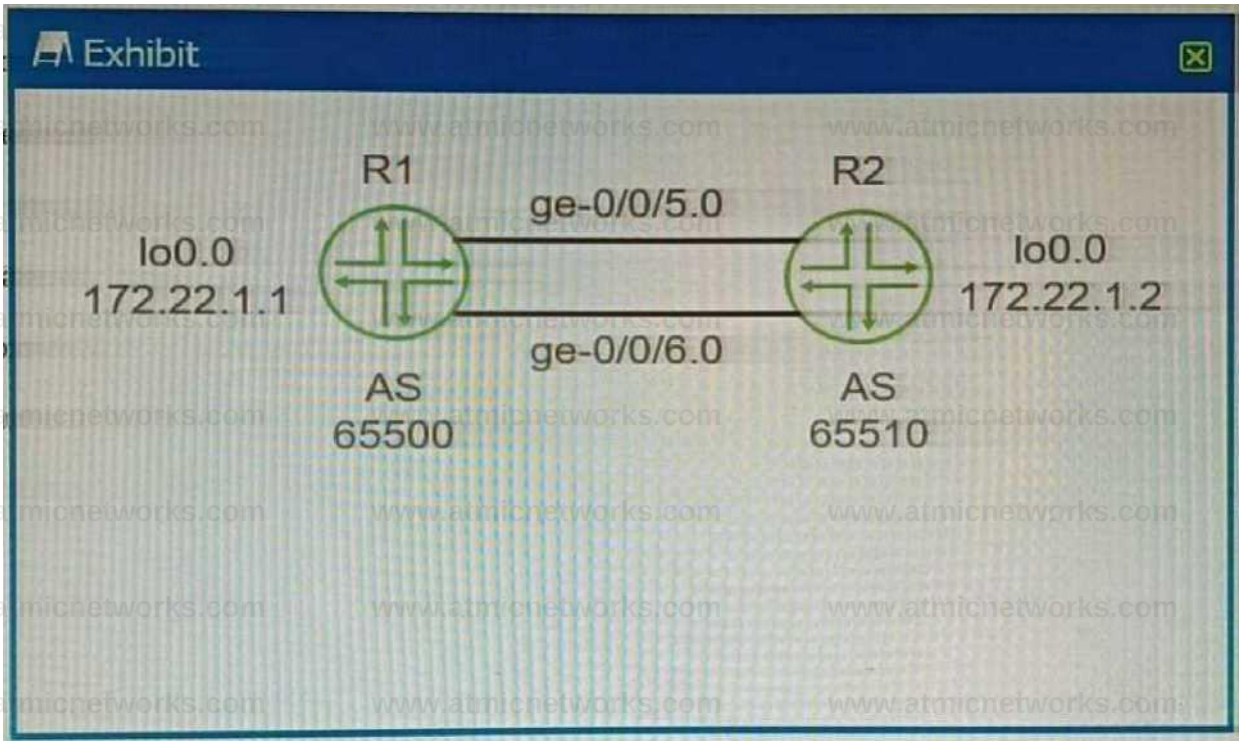

You want to enable redundancy for the EBGP peering between the two routers shown in the exhibit. Which three actions will you perform in this scenario? (Choose three.)

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 and reachable 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.

Question: 48

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.

Question: 49

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.

Question: 50

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 session information1. 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.

Question: 51

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 inet. o routing table.
D. RIB groups are used to hide routes in the inet. 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.

Question: 52

You want to ensure traffic is routed through a GRE tunnel.
In this scenario, which two statements will satisfy this requirement? (Choose two.)

A. Tunnel endpoints must have a route that directs traffic into the tunnel.
B. All intermediary devices must have a route to the tunnel endpoints.
C. Keepalives must be used on stateless tunneling protocols.
D. BFD must be used on the stateless tunneling protocols.

Answer: AB

Explanation:

Option A is correct. For traffic to be sent through a GRE tunnel, there must be a route that directs the traffic into the tunnel. This is typically accomplished through the use of a static route or a dynamic routing protocol.

Option B is correct. All intermediary devices must have a route to the tunnel endpoints34. In real- world scenarios, the tunnel endpoints for a tunnel going over the Internet must have globally reachable internet addresses. Otherwise, intermediate routers in the Internet cannot forward the tunneled packets.

Question: 53

Exhibit

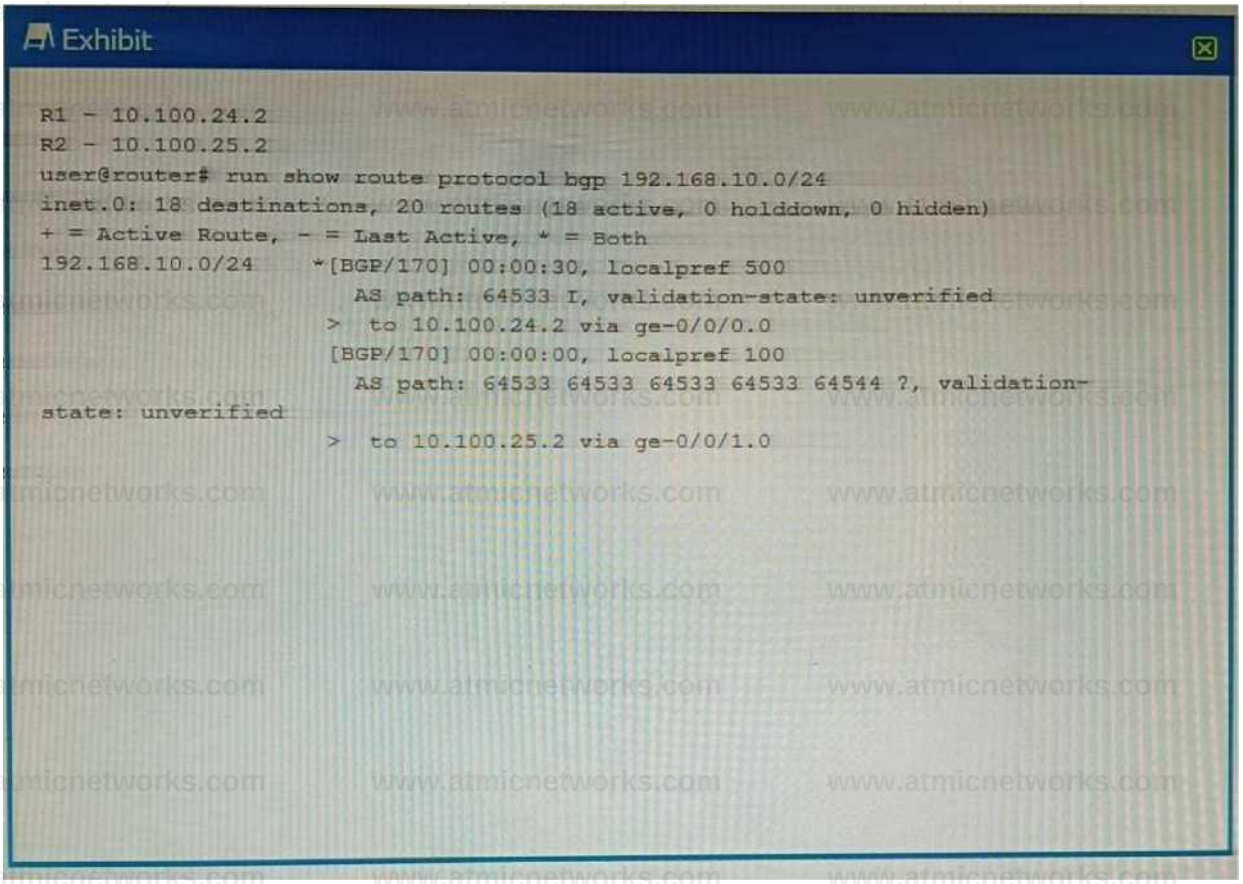

You are troubleshooting an issue where traffic to 192.168.10.0/24 is being sent to R1 instead of your desired path through R2.
Referring to the exhibit, what is the reason for the problem?

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.

Question: 54

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.

Question: 55

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

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

Answer: BC

Explanation:

B is correct because the loops do not exist is one of the conditions that are verified by the receiving router to ensure that the received BGP route is valid. A loop in BGP means that a route has been advertised by the same AS more than once, which can cause routing instability and inefficiency1. To prevent loops, BGP uses the AS-path attribute, which lists the AS numbers that a route has traversed from the origin to the destination2. The receiving router checks the AS-path attribute of the received route and discards it if it finds its own AS number in the list2. This way, BGP

avoids accepting routes that contain loops.

C is correct because the next hop is reachable is one of the conditions that are verified by the receiving router to ensure that the received BGP route is valid. The next hop is the IP address of the next router that is used to forward packets to the destination network3. The receiving router checks the next hop attribute of the received route and verifies that it has a valid route to reach it3. If the next hop is not reachable, the received route is not usable and is rejected by the receiving router3. This way, BGP ensures that only feasible routes are accepted.

Question: 56

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.

Question: 57

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.

Question: 58

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:
8) 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.

Question: 59

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

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

Answer: BC

Explanation:

The DHCP snooping feature in Juniper Networks' EX Series switches works by building a binding database that maps the IP address, MAC address, lease time, binding type, VLAN number, and interface information1. This database is used to filter and validate DHCP messages from untrusted sources1.
However, there are certain conditions that could prevent entries from being automatically created in the snooping database for an interface:

MAC limiting: If MAC limiting is enabled on the interface, it could potentially interfere with the operation of DHCP snooping. MAC limiting restricts the number of MAC addresses that can be learned on a physical interface to prevent MAC flooding attacks1. This could inadvertently limit the number of DHCP clients that can be learned on an interface, thus preventing new entries from being added to the DHCP snooping database.

Static IP address: If the device connected to the interface is configured with a static IP address, it will not go through the DHCP process and therefore will not have an entry in the DHCP snooping database1. The DHCP snooping feature relies on monitoring DHCP messages to build its database1, so devices with static IP addresses that do not send DHCP messages will not have their information added.

Therefore, options B and C are correct. Options A and D are not correct because performing a DHCPRELEASE would simply remove an existing entry from the database1, and Dynamic ARP inspection (DAI) uses the information stored in the DHCP snooping binding database but does not prevent entries from being created1.

Question: 60

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 autorecovery from port error disabled conditions, you must manually clear the 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.

Question: 61

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 nondefault 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.

Question: 62

Exhibit

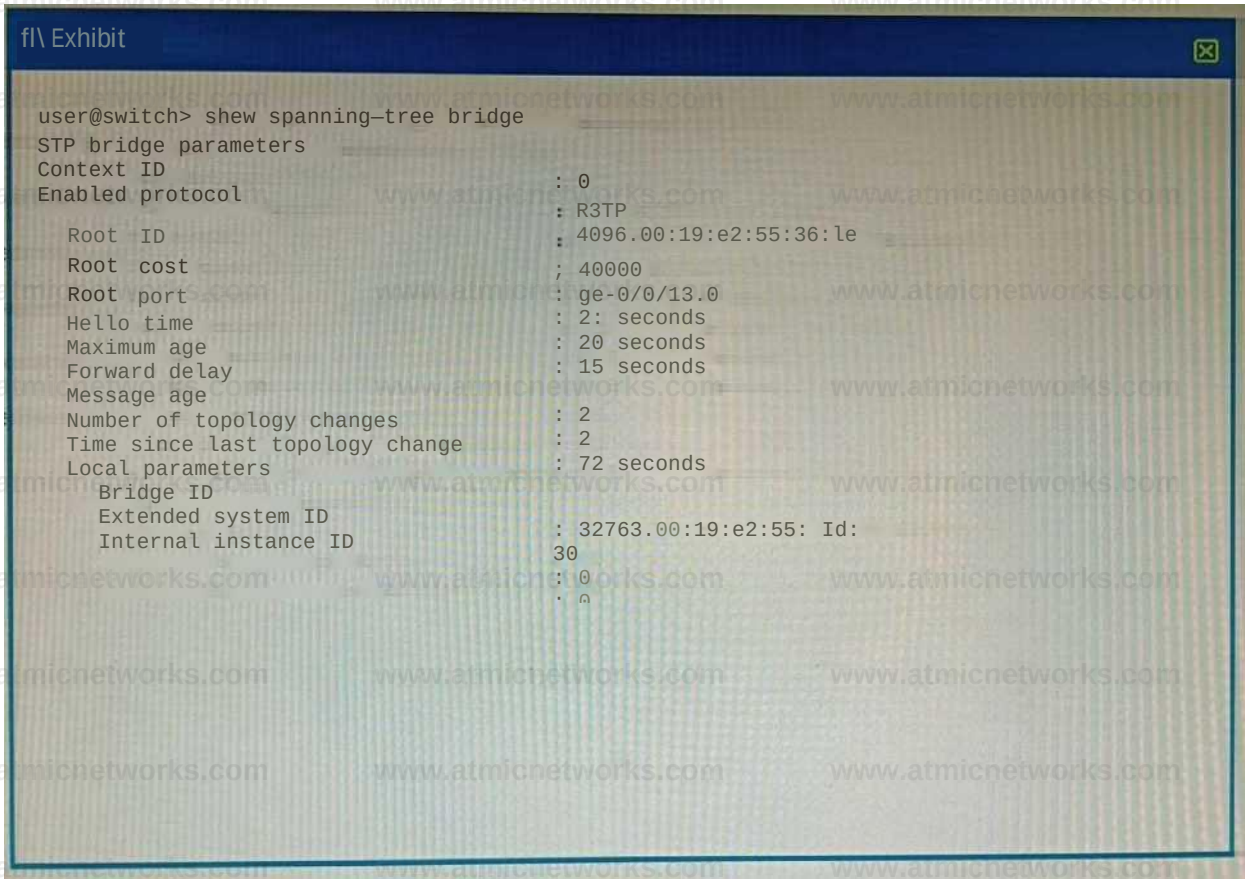

Referring to the exhibit, which statement is correct?

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.

Question: 63

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.

Question: 64

Exhibit

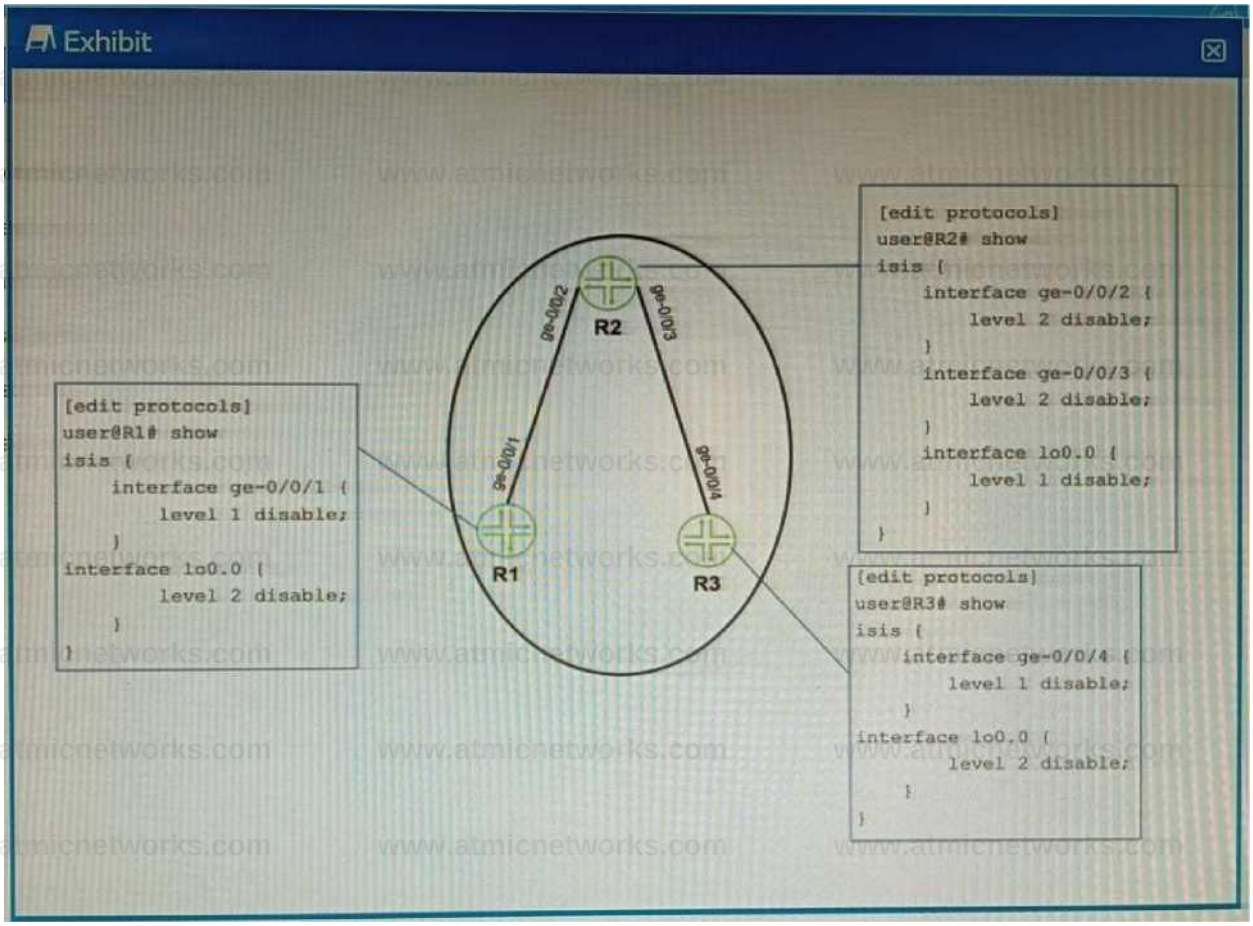

Referring to the exhibit, which two configuration changes must you apply for packets to reach from R1 to R3 using IS-IS? (Choose two.)

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.

Question: 65

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 nexthop addresses for a static route1. However, it does not provide the capability to summarize multiple routes into a single prefix.