Juniper

Exam Questions JN0-351

Enterprise Routing and Switching - Specialist (JNCIS-ENT)

NEW QUESTION 1
Exhibit

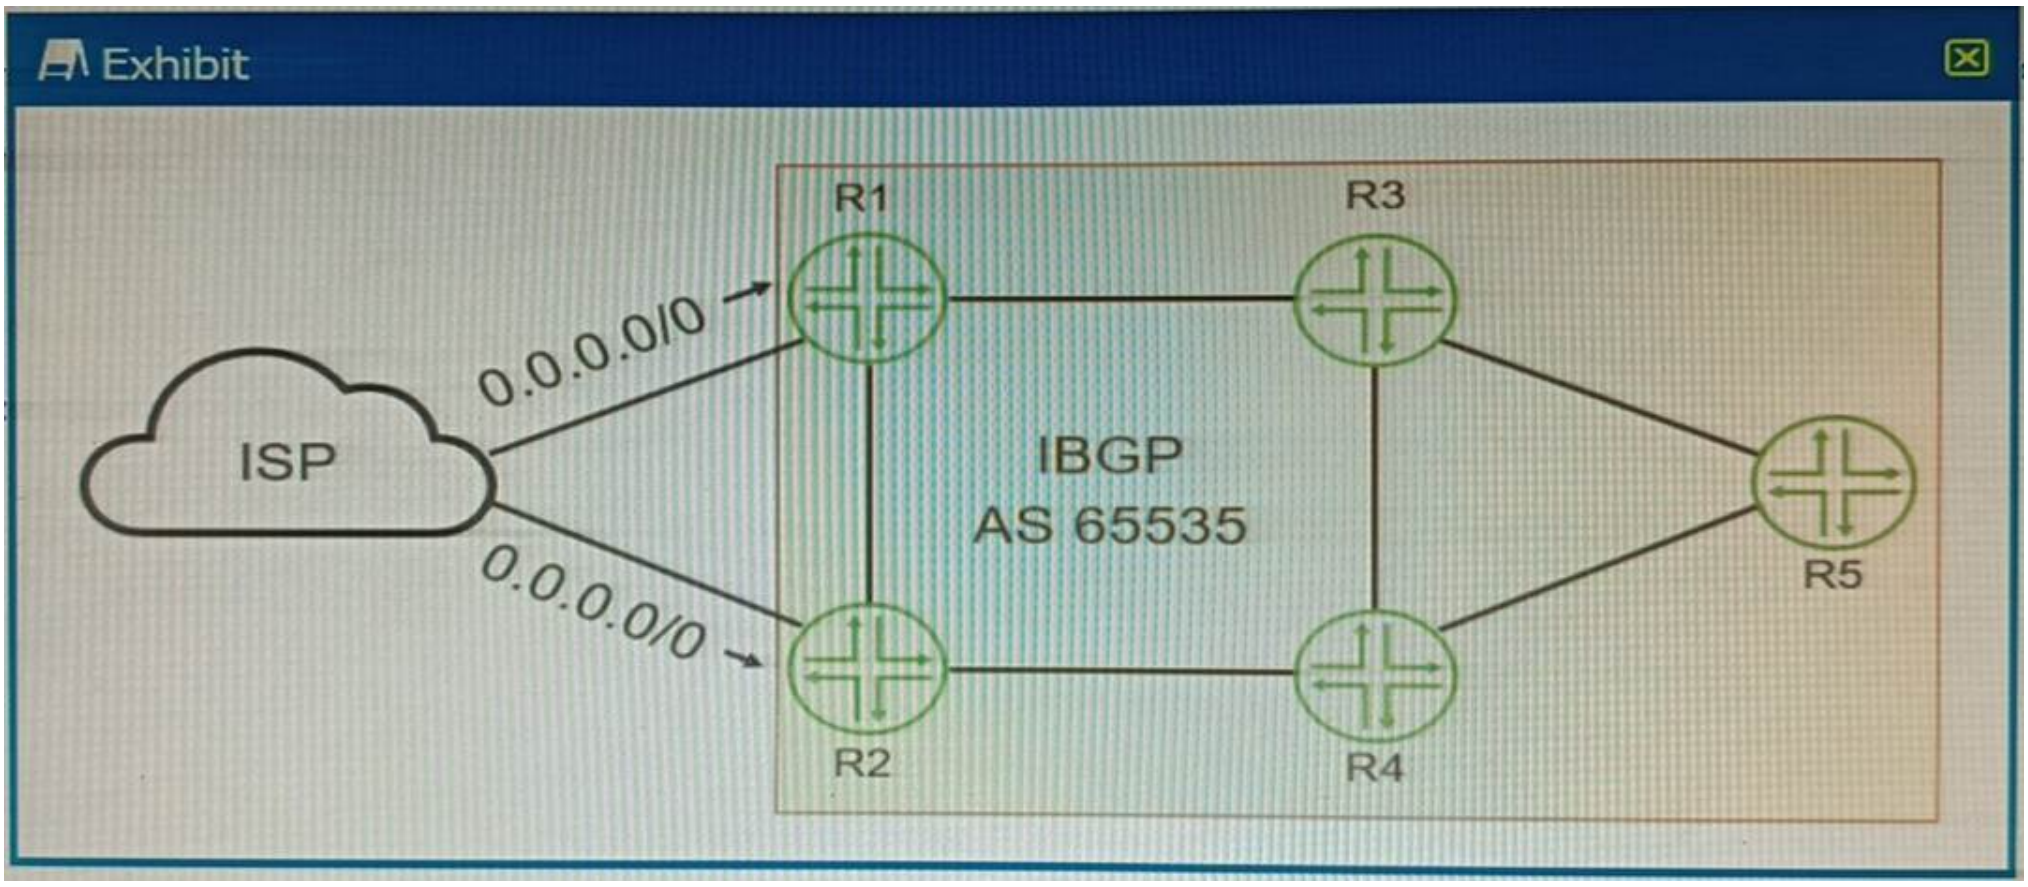

Your ISP is announcing a default route to both R1 and R2. You want your network routers to forward all Internet traffic through the R1 device
Which BGP attribute would you use?

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

Answer: C

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

NEW QUESTION 2
What is the default keepalive time for BGP?

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

Answer: B

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

NEW QUESTION 3
You are receiving multiple BGP routes from an upstream neighbor and only want to advertise a single summarized prefix to your internal OSPF neighbors. This route should only be advertised when you are receiving these BGP routes from this neighbor.
In this scenario, which type of route should you create?

A. aggregate route
B. static route using the resolve feature
C. generate route
D. static route using qualified next hops

Answer: A

Explanation:
In this scenario, you should create an aggregate route1. Aggregate routes are used for advertising summarized network prefixes1. They help minimize the number of routing tables in an IP network by consolidating selected multiple routes into a single route advertisement1. This approach is in contrast to non-aggregation routing, in which every routing table contains a unique entry for each route1.
Therefore, option A is correct. Options B, C, and D are not correct because:
? Static route using the resolve feature: This type of route uses the resolve feature to install a static route in the routing table only if a specific condition is met1. However, it does not provide the capability to summarize multiple routes into a single prefix.
? Generate route: This type of route generates a route that is always present in the routing table and can be used to summarize routes. However, it does not have

the capability to only advertise the route when specific BGP routes are being received from a neighbor1.
? Static route using qualified next hops: This type of route allows for the specification of multiple next-hop addresses for a static route1. However, it does not provide the capability to summarize multiple routes into a single prefix.

NEW QUESTION 4
Exhibit.

Exhibit

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

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

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

Answer: B

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

NEW QUESTION 5
Which two statements about redundant trunk groups on EX Series switches are correct? (Choose two.)

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

Answer: BD

Explanation:
Redundant Trunk Groups (RTGs) on EX Series switches provide a simple solution for network recovery when a trunk port on a switch goes down1. They are configured on the access switch and contain two links: a primary or active link, and a secondary link1. Therefore, option B is correct because if the active link fails, the secondary link automatically starts forwarding data traffic without waiting for normal spanning-tree protocol convergence1.

Option D is also correct. In a typical enterprise network composed of distribution and access layers, RTGs are used where one Access switch is connected to two different uplink switches2. This implies that RTGs must be connected to the same aggregation switch2.

NEW QUESTION 6
After receiving a BGP route, which two conditions are verified by the receiving router to ensure that the received route is valid? (Choose two)

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

Answer: BC

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

NEW QUESTION 7
You have DHCP snooping enabled but no entries are automatically created in the snooping database for an interface on your EX Series switch. What are two reasons for the problem? (Choose two.)

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

Answer: BC

Explanation:
The DHCP snooping feature in Juniper Networks?? EX Series switches works by building a binding database that maps the IP address, MAC address, lease time, binding type, VLAN number, and interface information1. This database is used to filter and validate DHCP messages from untrusted sources1.
However, there are certain conditions that could prevent entries from being automatically created in the snooping database for an interface:
? MAC limiting: If MAC limiting is enabled on the interface, it could potentially
interfere with the operation of DHCP snooping. MAC limiting restricts the number of MAC addresses that can be learned on a physical interface to prevent MAC flooding attacks1. This could inadvertently limit the number of DHCP clients that can be learned on an interface, thus preventing new entries from being added to the DHCP snooping database.
? Static IP address: If the device connected to the interface is configured with a
static IP address, it will not go through the DHCP process and therefore will not have an entry in the DHCP snooping database1. The DHCP snooping feature relies on monitoring DHCP messages to build its database1, so devices with static IP addresses that do not send DHCP messages will not have their information added.
Therefore, options B and C are correct. Options A and D are not correct because performing a DHCPRELEASE would simply remove an existing entry from the database1, and Dynamic ARP inspection (DAI) uses the information stored in the DHCP snooping binding database but does not prevent entries from being created1.

NEW QUESTION 8
Exhibit

Exhibit

```
user@switch> show spanning-tree bridge
STP bridge parameters
Context ID                                : 0
Enabled protocol                          : RSTP
  Root ID                                 : 4096.00:19:e2:55:36:1e
  Root cost                               : 40000
  Root port                               : ge-0/0/13.0
  Hello time                              : 2 seconds
  Maximum age                             : 20 seconds
  Forward delay                           : 15 seconds
  Message age                             : 2
  Number of topology changes              : 2
  Time since last topology change         : 72 seconds
  Local parameters
    Bridge ID                             : 32768.00:19:e2:55:1d:30
    Extended system ID                    : 0
    Internal instance ID                  : 0
```

Referring to the exhibit, which statement is correct?

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

Answer: D

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

NEW QUESTION 9
Refer to the exhibit.

Exhibit

```
user@host> show ospf neighbor
Address          Interface              State   ID               Pri  Dead
172.26.1.1       ge-0/0/3.0             2Way    192.168.1.1      128    31
```

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

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

Answer: B

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

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

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

Answer: BD

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

NEW QUESTION 10
Which two mechanisms are part of building and maintaining a Layer 2 bridge table? (Choose two.)

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

Answer: BC

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

NEW QUESTION 15
An update to your organization's network security requirements document requires management traffic to be isolated in a non-default routing-instance. You want to implement
this requirement on your Junos-based devices.
Which two commands enable this behavior? (Choose two.)

A. set routing—instances mgmtjunoa interface ge-0/0/0.0
B. set routing—instances mgmt_junos interface em1
C. set system management—instance
D. set routing—instances mgmt_junos

Answer: CD

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

NEW QUESTION 17
A new network requires multiple topology support. You decide to use IS-IS in this situation. Which three protocol topologies are supported in this scenario? (Choose three.)

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

Answer: CDE

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

NEW QUESTION 19
What is a purpose of using a spanning tree protocol?

A. to look up MAC addresses
B. to eliminate broadcast storms
C. to route IP packets
D. to tunnel Ethernet frames

Answer: B

Explanation:
? A broadcast storm is a network condition where a large number of broadcast packets are sent and received by multiple devices, causing congestion and performance degradation1. A broadcast storm can occur when there are loops in the network topology, meaning that there are multiple paths between two devices2.
? A spanning tree protocol is a network protocol that prevents loops from being formed when switches or bridges are interconnected via multiple paths. It does this by creating a logical tree structure that spans all the devices in the network, and disabling or blocking the links that are not part of the tree, leaving a single active path between any two devices3.
? By eliminating loops, a spanning tree protocol also eliminates broadcast storms, as broadcast packets will not be forwarded endlessly along the looped paths. Instead, broadcast packets will be sent only along the tree structure, reaching each device once and avoiding congestion3.

NEW QUESTION 23
Which statement is correct about the storm control feature?

A. The storm control feature is enabled in the factory-default configuration on EX Series switches.
B. The storm control feature requires a special license on EX Series switches.
C. The storm control feature is not supported on aggregate Ethernet interfaces.
D. The storm control configuration only applies to traffic being sent between the forwarding and control plane.

Answer: A

Explanation:
? Option A is correct. The storm control feature is enabled in the factory-default configuration on EX Series switches12. On EX2200, EX3200, EX3300, EX4200, and EX6200 switches, the factory default configuration enables storm control for broadcast and unknown unicast traffic on all switch interfaces2. On EX4300 switches, the factory default configuration enables storm control on all Layer 2 switch interfaces1.
? Option B is incorrect. The storm control feature does not require a special license on EX Series switches34.
? Option C is incorrect. There??s no information available that suggests the storm control feature is not supported on aggregate Ethernet interfaces.
? Option D is incorrect. The storm control configuration applies to traffic at the ingress of an interface5, not just between the forwarding and control plane.

NEW QUESTION 27
Which statement is correct about IP-IP tunnels?

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

Answer: A

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

NEW QUESTION 29
Exhibit.

Exhibit

```
user@R1> show route receive-protocol bgp 10.36.1.4
inet.0: 33 destinations, 57 routes (33 active, 0 holddown, 0 hidden)
  Prefix    Nexthop          MED     Lclpref    AS path
* 10.30.100.8/32           10.36.1.4                          65401 65520 I
* 10.30.100.9/32           10.36.1.4                          65401 65521 I
* 10.30.189.0/30           10.36.1.4                          65401 65521 I
  10.32.1.0/30             10.36.1.4                          65401 I
* 10.32.2.0/30             10.36.1.4                          65401 I
* 10.32.12.0/30            10.36.1.4                          65401 I
* 10.52.100.2/32           10.36.1.4                          65401 I
```

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

A. The routes displayed have traversed one or more autonomous systems.
B. The output shows routes that were received prior to the application of any BGP import policies.
C. The output shows routes that are active and rejected by an import policy.
D. The routes displayed are being learned from an I BGP peer.

Answer: AB

Explanation:
The output shown in the exhibit is the result of the command ??show ip bgp neighbor 10.36.1.4 received-routes??, which displays all received routes (both accepted and rejected) from the specified neighbor.
Option A is correct, because the routes displayed have traversed one or more autonomous systems. This can be seen from the AS_PATH attribute, which shows the sequence of AS numbers that the route has passed through. For example, the route 10.0.0.0/8 has an AS_PATH of 65001 65002, which means that it has traversed AS 65001 and AS 65002 before reaching the local router.
Option B is correct, because the output shows routes that were received prior to the application of any BGP import policies. This can be seen from the fact that some routes have a status code of ??r??, which means that they are rejected by an import policy. The??received-routes?? keyword shows the routes coming from a given neighbor before the inbound policy has been applied. To see the routes after the inbound policy has been applied, the ??routes?? keyword should be used instead.
Option C is incorrect, because the output does not show routes that are active and rejected by an import policy. The status code of ??r?? means that the route is rejected by an import policy, but it does not mean that it is active. The status code of ??>?? means that the route is active and selected as the best path. None of the routes in the output have both ??>?? and ??r?? status codes.
Option D is incorrect, because the routes displayed are not being learned from an IBGP peer. An IBGP peer is a BGP neighbor that belongs to the same AS as the local router. The output shows that the neighbor 10.36.1.4 has a remote AS of 65001, which is different from the local AS of 65002. Therefore, the neighbor is an EBGP peer, not an IBGP peer.

NEW QUESTION 32
You are configuring an IS-IS IGP network and do not see the IS-IS adjacencies established. In this scenario, what are two reasons for this problem? (Choose two.)

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

Answer: AD

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

NEW QUESTION 36
What are two characteristics of RSTP alternate ports? (Choose two.)

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

Answer: AC

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

NEW QUESTION 39
Which two statements are correct about using firewall filters on EX Series switches? (Choose two.)

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

Answer: AC

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

NEW QUESTION 41
You need to configure a LAG between your switches. In this scenario, which two statements are correct? (Choose two.)

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

Answer: BC

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

NEW QUESTION 43
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.
References:
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

NEW QUESTION 47
Exhibit

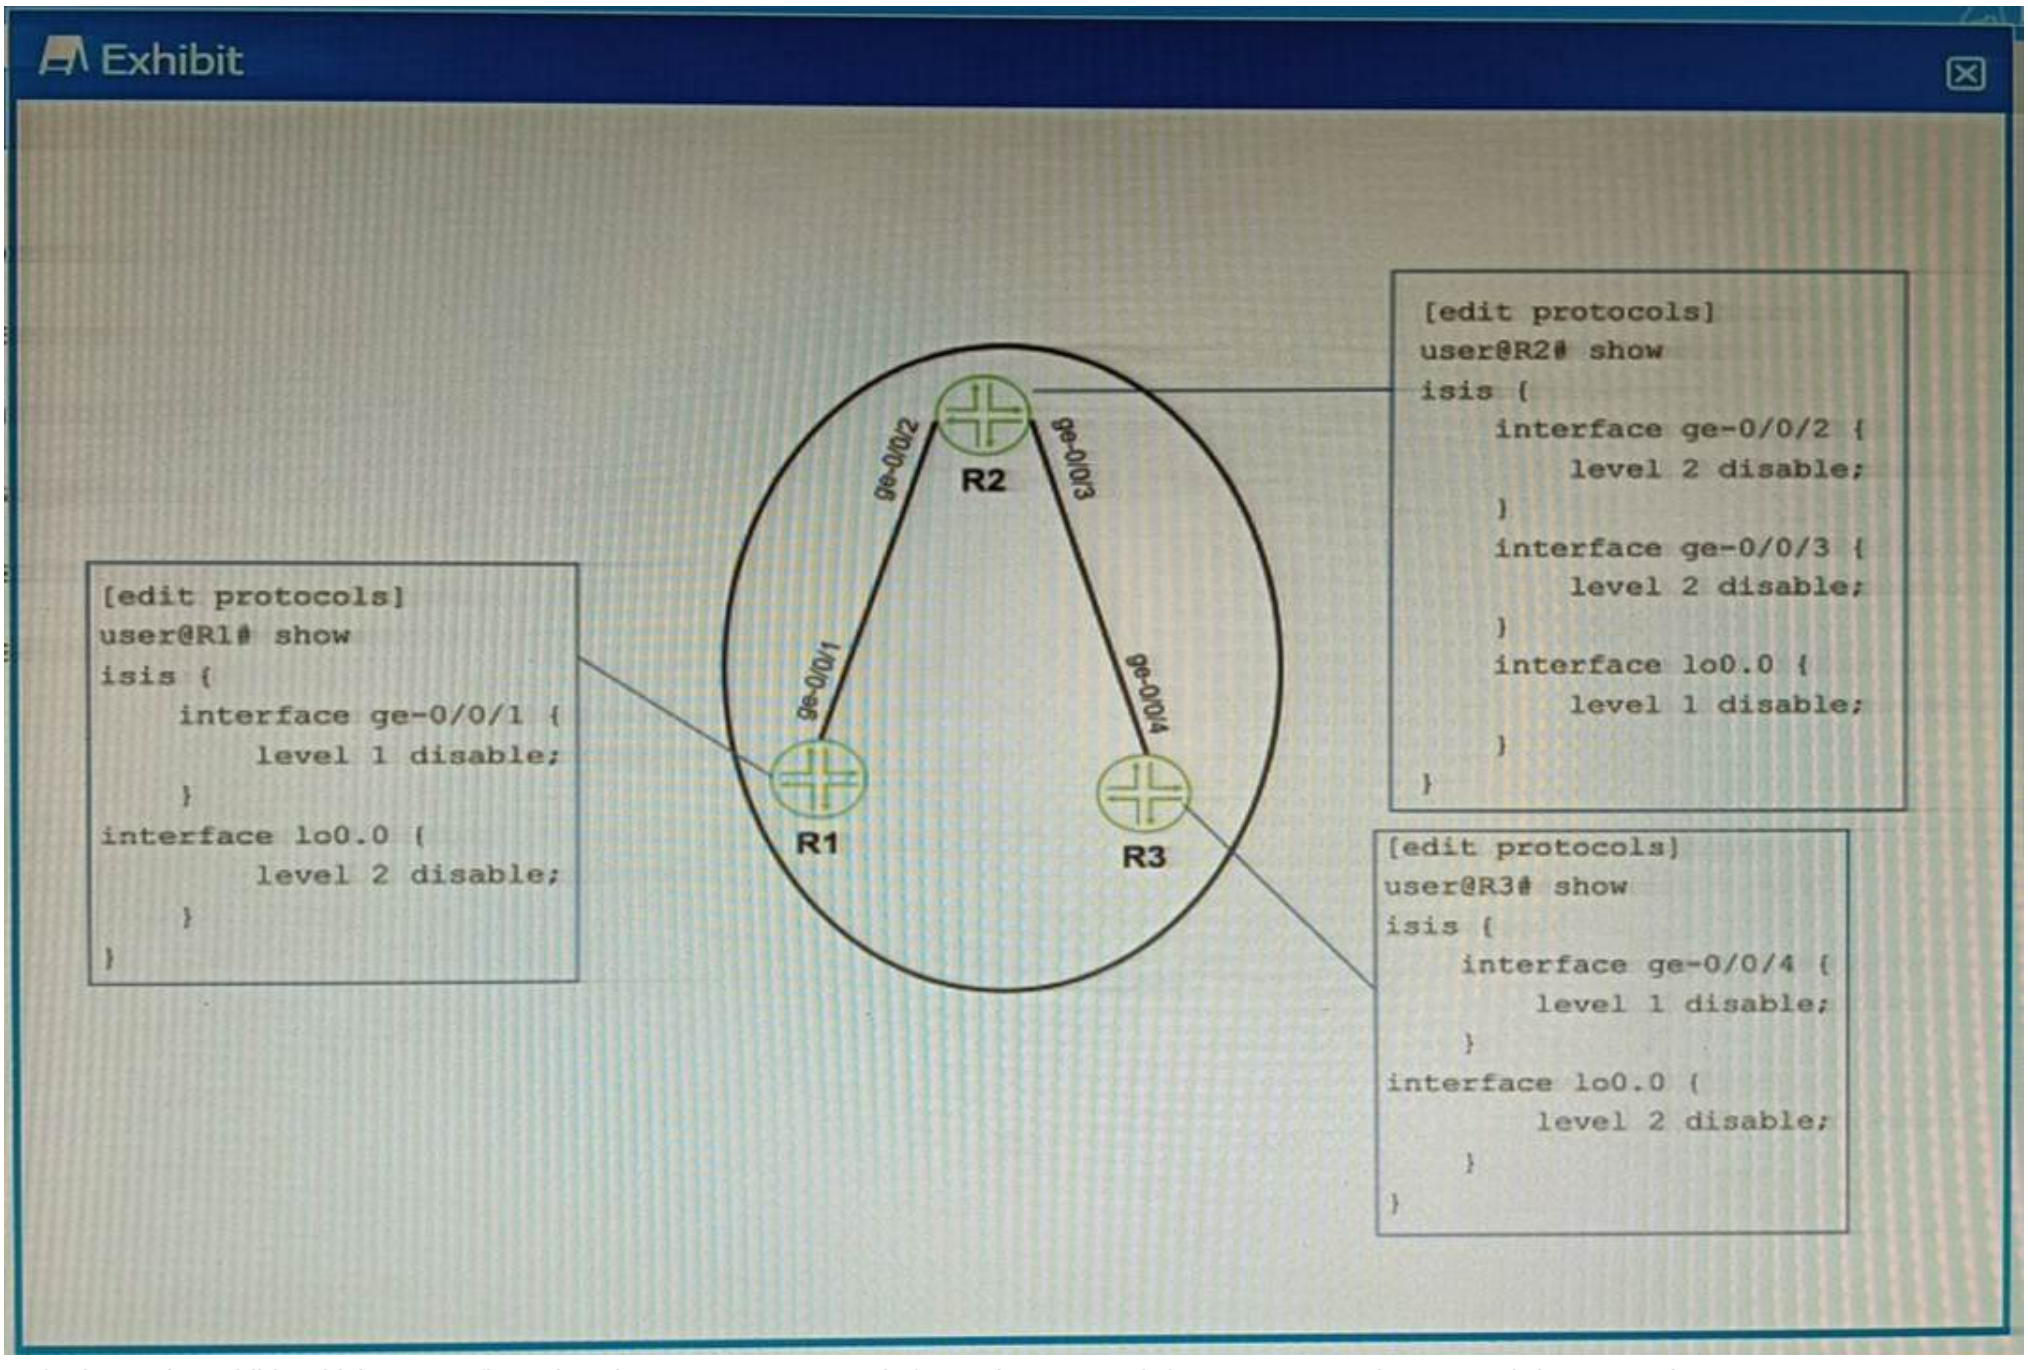

Referring to the exhibit, which two configuration changes must you apply for packets to reach from R1 to R3 using IS-IS? (Choose two.)

A. On R1, enable Level 1 on the ge-0/0/1 interface.
B. On R3 disable Level 2 on the ge-0/0/4 interface.
C. On R1, disable Level 2 on the ge-0/0/1 interface.
D. On R3 enable Level 1 on the ge-0/0/4 interface

Answer: AD

Explanation:
A. On R1, enable Level 1 on the ge-0/0/1 interface. In IS-IS, both levels (Level 1 and Level 2) are enabled by default when you enable IS-IS on an interface1. Level 1 systems route within an area2. If the destination is outside an area, Level 1 systems route toward a Level 2 system2. Therefore, enabling Level 1 on the ge-0/0/1 interface on
R1 would allow packets to reach from R1 to R3.
* D. On R3 enable Level 1 on the ge-0/0/4 interface Similarly, enabling Level 1 on the ge- 0/0/4 interface on R3 would allow packets to reach from R1 to R3.
These explanations are based on the IS-IS configuration documents and learning resources available at Juniper Networks1 and Cisco34.

NEW QUESTION 51
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

JN0-351 Practice Exam Features:

* JN0-351 Questions and Answers Updated Frequently

* JN0-351 Practice Questions Verified by Expert Senior Certified Staff

* JN0-351 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* JN0-351 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year