Exam Questions JN0-351

Enterprise Routing and Switching - Specialist (JNCIS-ENT)

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

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

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

B)

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

C)

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

D)

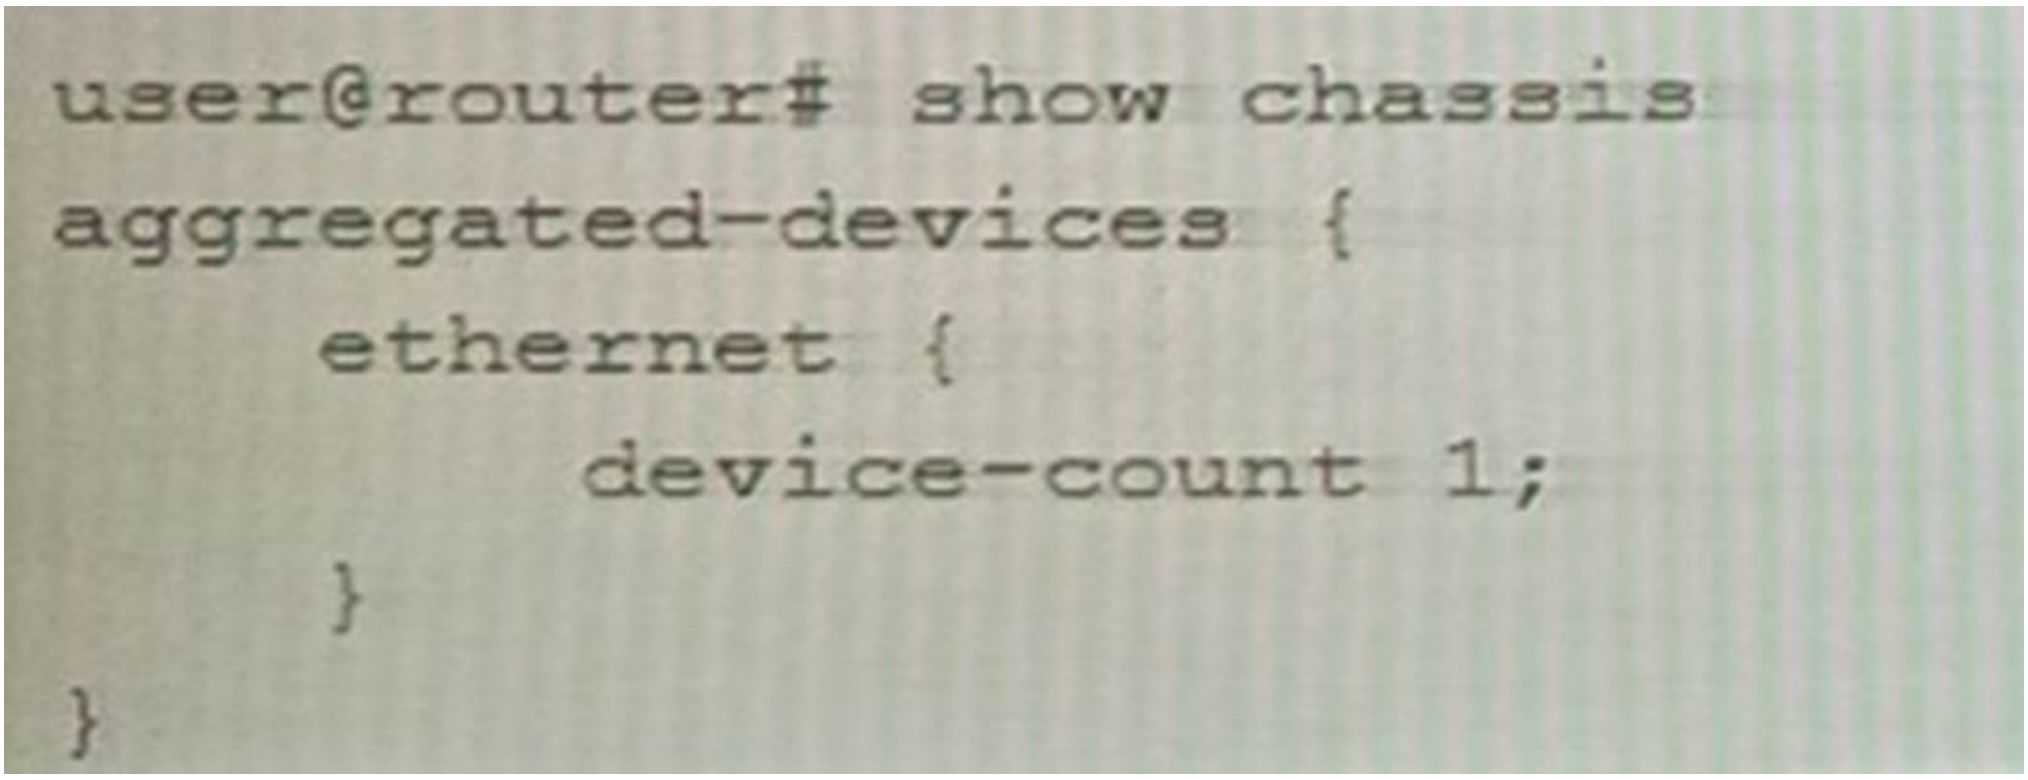

A. Option A
B. Option B
C. Option C
D. Option D

Answer: C

Explanation:
The correct answer to your question is C. Option C. Here is why:
? Option C shows the configuration of the chassis statement, which defines the properties of the router chassis, such as the number of aggregated Ethernet interfaces, the number of FPCs, and the number of PICs1.
? To enable aggregated Ethernet interfaces on a router, you need to specify the aggregated-devices statement under the chassis statement and set the ethernet parameter to the desired number of interfaces2. For example, to enable two aggregated Ethernet interfaces, you can use the following configuration: chassis { aggregated-devices { ethernet { device-count 2; } } }
? Option C shows this configuration with the device-count set to 2, which will enable two aggregated Ethernet interfaces on the router. The other options do not show this configuration and will not enable any aggregated Ethernet interfaces on the router.
? Therefore, option C is the correct answer to your question.

NEW QUESTION 2
Exhibit.

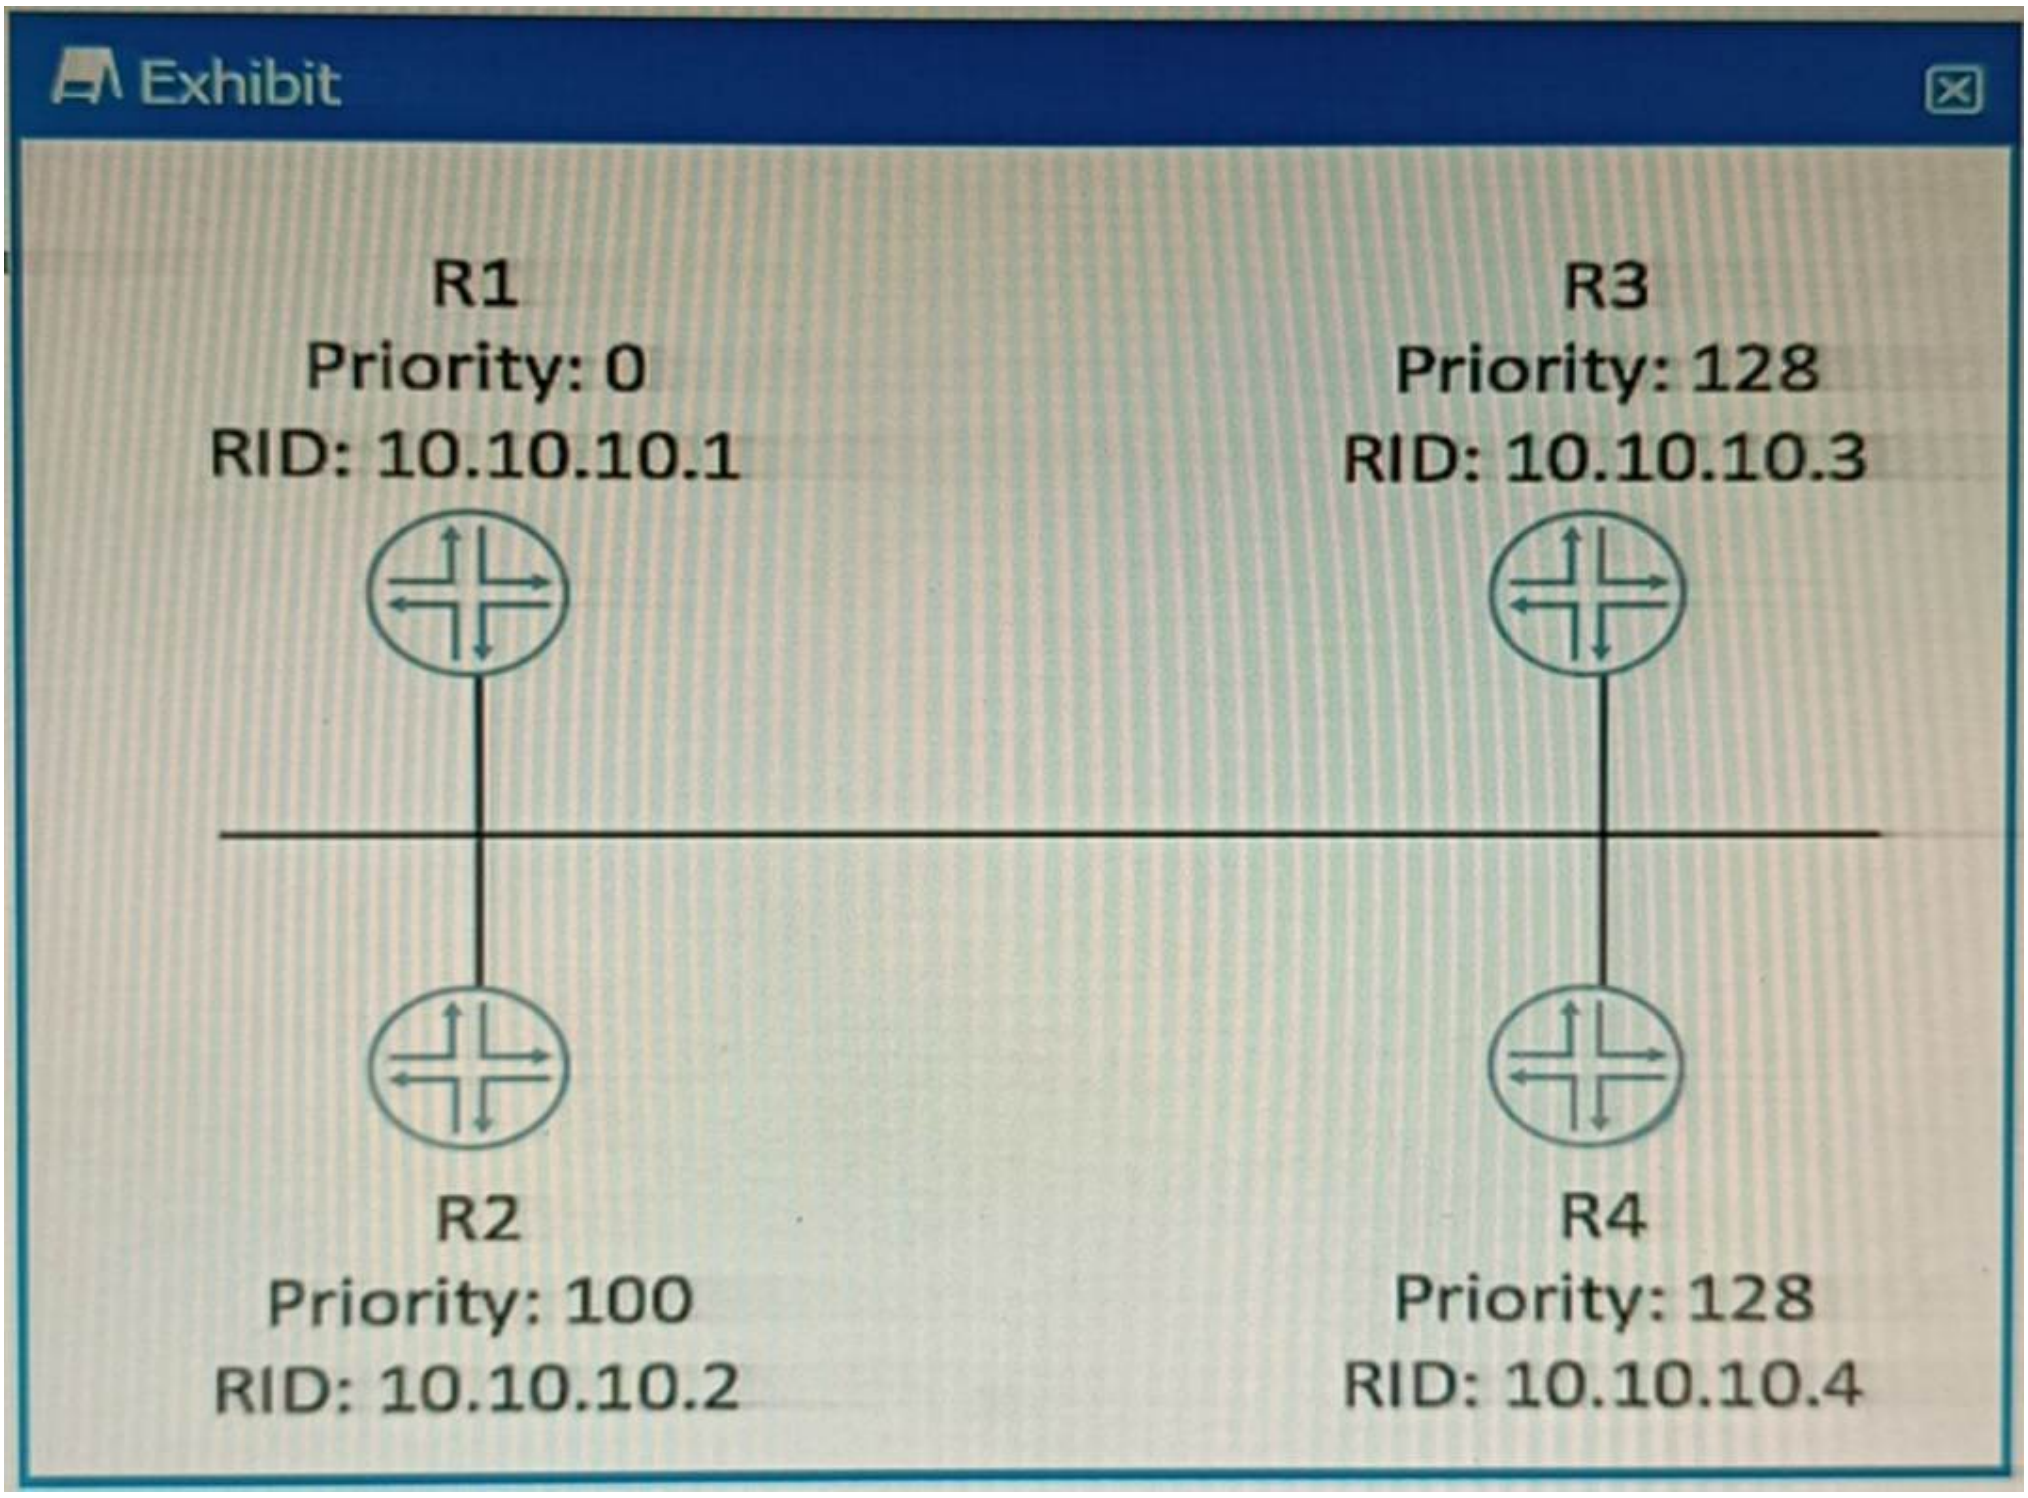

Which router will become the OSPF BDR if all routers are powered on at the same time?

A. R4
B. R1
C. R3
D. R2

Answer: A

Explanation:
OSPF DR/BDR election is a process that occurs on multi-access data links. It is intended to select two OSPF nodes: one to be acting as the Designated Router (DR), and another to be acting as the Backup Designated Router (BDR).The DR and BDR are responsible for generating network LSAs for the multi-access network and synchronizing the LSDB with other routers on the same network1.
The DR/BDR election is based on two criteria: the OSPF priority and the router ID. The OSPF priority is a value between 0 and 255 that can be configured on each interface participating in OSPF. The default priority is 1. A priority of 0 means that the router will not participate in the election and will never become a DR or BDR. The router with the highest priority will become the DR, and the router with the second highest priority will become the BDR. If there is a tie in priority, then the router ID is used as a tie-breaker. The router ID is a 32-bit number that uniquely identifies each router in an OSPF domain.It can be manually configured or automatically derived from the highest IP address on a loopback interface or any active interface2.
In this scenario, all routers have the same priority of 1, so the router ID will determine the outcome of the election. The router IDs are shown in the exhibit as RID values. The highest
RID belongs to R4 (10.10.10.4), so R4 will become the DR. The second highest RID belongs to R3 (10.10.10.3), so R3 will become the BDR.
References:
1:OSPF DR/BDR Election: Process, Configuration, and Tuning2:OSPF Designated Router (DR) and Backup Designated Router (BDR)

NEW QUESTION 3
Exhibit

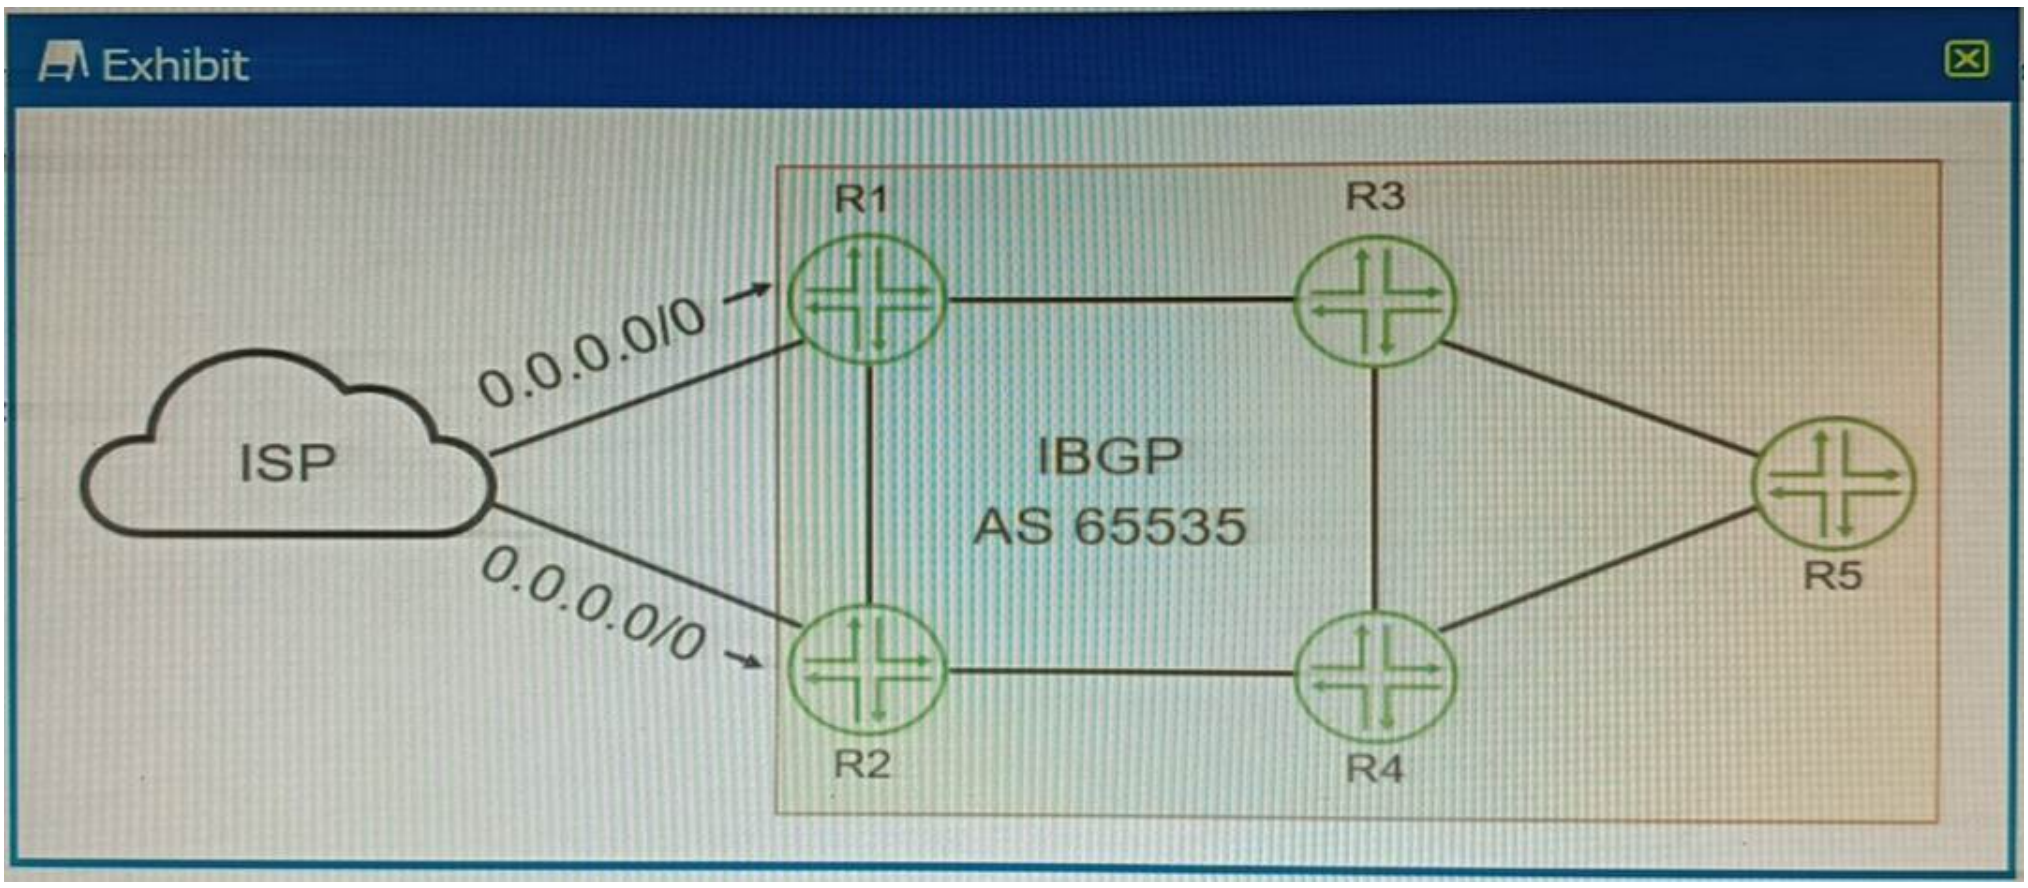

Your ISP is announcing a default route to both R1 and R2. You want your network routers to forward all Internet traffic through the R1 device
Which BGP attribute would you use?

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 4
Exhibit.

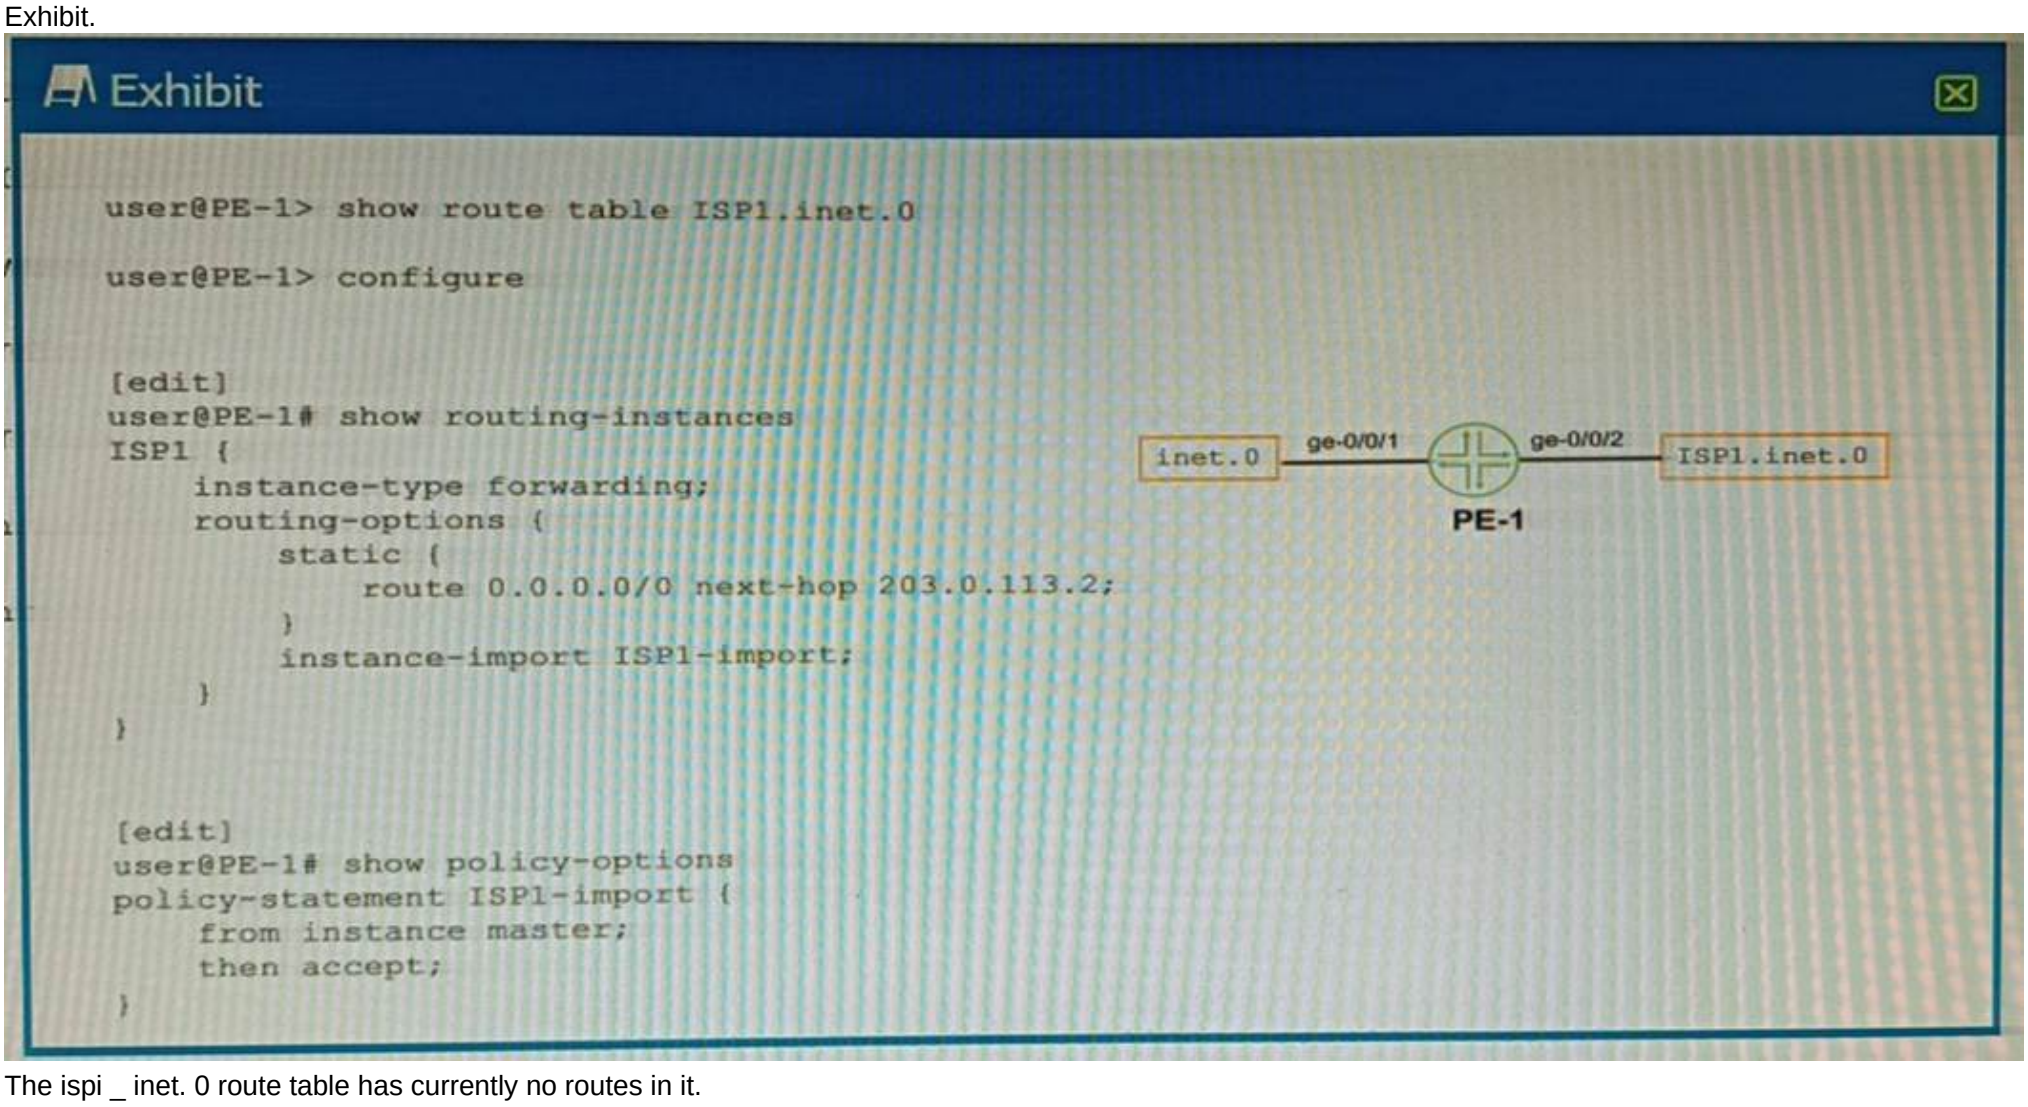

The ispi _ inet. 0 route table has currently no routes in it.
What will happen when you commit the configuration shown on the exhibit?

A. The ine
B. 0 route table will be completely overwritten by the ispi . ine
C. 0 route table.
D. The ine
E. 0 route table will be imported into the ispi . ine
F. 0 route table.

G. The ISPI . ine
H. 0 route table will be completely overwritten by the ine
I. o route table.
J. The ISPI . ine
K. 0 route table will be imported into the ine
L. 0 route table.

Answer: B

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

NEW QUESTION 5
Exhibit

Exhibit

```
user# show protocols bgp


group ext-64501 {
    type external;
    peer-as 64501;
    neighbor 172.30.1.2;
}
group int-64503 {
    type internal;
    local-address 192.168.100.1;
    neighbor 192.168.100.2;
}
bfd-liveness-detection {
    minimum-interval 10;
}
```

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

A. The BFD liveness is set too low.
B. The BFD liveness must be configured on the BGP neighbor.
C. The BFD liveness must be configured on the BGP group.
D. The BFD liveness is set too high.

Answer: B

Explanation:
? The exhibit shows the configuration of BFD liveness detection for BGP at the global level, which applies to all BGP neighbors by default1. However, this configuration does not specify the session mode, which determines whether BFD uses single-hop or multihop mode to communicate with a neighbor2.
? For single-hop BGP neighbors, which are directly connected on the same subnet, the session mode can be either automatic or single-hop. For multihop BGPneighbors, which are not directly connected and require multiple hops to reach, the session mode must be multihop2.
? Since your BGP neighbors are in different countries, they are likely to be multihop neighbors. Therefore, you need to configure the session mode as multihop for each neighbor individually at the [edit protocols bgp group group-name neighbor address bfd-liveness-detection] hierarchy level2. For example:
protocols { bgp { group usa { neighbor 192.0.2.1 { bfd-liveness-detection { session-mode multihop; } } } group france { neighbor 198.51.100.1 { bfd-liveness-detection { session-mode multihop; } } } } }
? If you do not configure the session mode for multihop neighbors, BFD will use the
default mode of automatic, which will try to use single-hop mode and fail to establish a BFD session with the remote neighbor2. This will prevent BGP from using BFD to detect liveliness and failover.

? Therefore, the answer B is correct, as you need to configure the BFD liveness
detection on the BGP neighbor level with the appropriate session mode for multihop neighbors.

NEW QUESTION 6
You are receiving multiple BGP routes from an upstream neighbor and only want to advertise a single summarized prefix to your internal OSPF neighbors. This route should only be advertised when you are receiving these BGP routes from this neighbor.
In this scenario, which type of route should you create?

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

Answer: A

Explanation:
In this scenario, you should create an aggregate route1. Aggregate routes are used for advertising summarized network prefixes1. They help minimize the number of routing tables in an IP network by consolidating selected multiple routes into a single route advertisement1. This approach is in contrast to non-aggregation routing, in which every routing table contains a unique entry for each route1.
Therefore, option A is correct. Options B, C, and D are not correct because:
? Static route using the resolve feature: This type of route uses the resolve feature to install a static route in the routing table only if a specific condition is met1. However, it does not provide the capability to summarize multiple routes into a single prefix.
? Generate route: This type of route generates a route that is always present in the routing table and can be used to summarize routes. However, it does not have the capability to only advertise the route when specific BGP routes are being received from a neighbor1.
? Static route using qualified next hops: This type of route allows for the specification of multiple next-hop addresses for a static route1. However, it does not provide the capability to summarize multiple routes into a single prefix.

NEW QUESTION 7
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.

NEW QUESTION 8
Exhibit.

Exhibit

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

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

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

Answer: B

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

NEW QUESTION 9
Exhibit

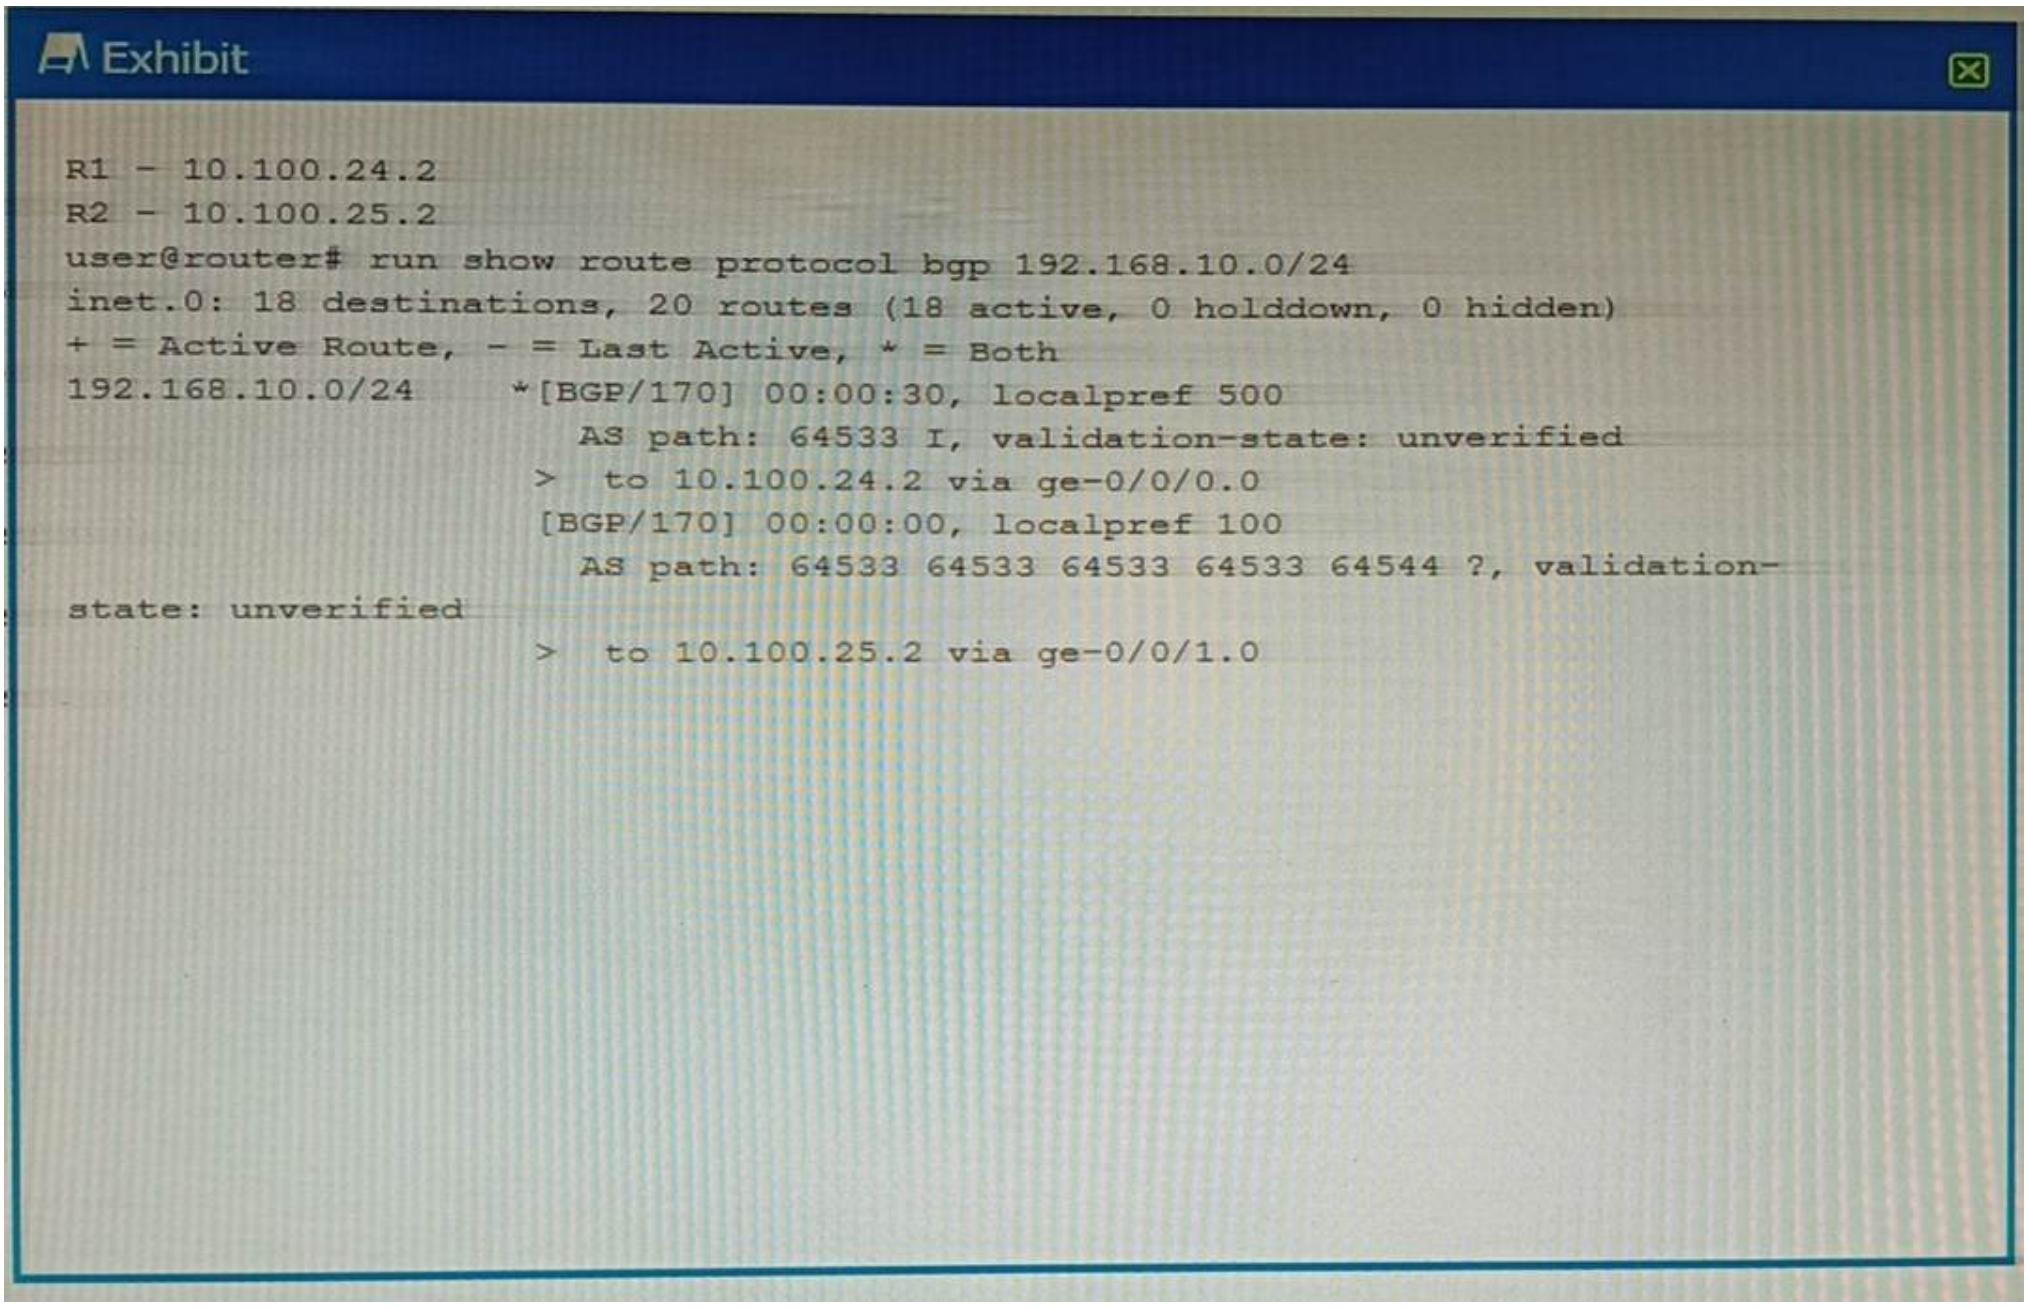

You are troubleshooting an issue where traffic to 192.168.10.0/24 is being sent to R1 instead of your desired path through R2.
Referring to the exhibit, what is the reason for the problem?

A. R2's route is not the best path due to loop prevention.
B. R2's route is not the best path due to a lower origin code.
C. R1's route is the best path due to a higher local preference
D. R1's route is the best path due to the shorter AS path.

Answer: C

Explanation:
? The exhibit shows the output of the command show ip bgp, which displays information about the BGP routes in the routing table1. The output shows two routes for the destination 192.168.10.0/24, one from R1 and one from R2.
? The route from R1 has a local preference of 200, while the route from R2 has a local preference of 100. Local preference is a BGP attribute that indicates the degree of preference for a route within an autonomous system (AS)2. A higher local preference means a more preferred route2.
? BGP uses a best path selection algorithm to choose the best route for each destination among multiple paths. The algorithm compares different attributes of the routes in a specific order of precedence3. The first attribute that is compared is weight, which is a Cisco-specific attribute that is local to the router3. If the weight is equal or not set, the next attribute that is compared is local preference3.
? In this case, both routes have the same weight of 0, which means that they are learned from external BGP (eBGP) peers3. Therefore, the next attribute that is compared is local preference. Since R1??s route has a higher local preference than R2??s route, it is chosen as the best path and installed in the routing table3.
The other attributes, such as origin code and AS path, are not considered in this case.

NEW QUESTION 10
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.

NEW QUESTION 10
Two routers share the same highest priority and start time.

A. In this situation, what is evaluated next when determining the designated router? The router with the lowest router ID become the DR.
B. The router with the highest router ID becomes the DR
C. The routers perform another DR election.
D. The router with the highest MAC address become the DR

Answer: B

Explanation:
? According to the OSPF protocol, the designated router (DR) is the router that acts as the focal point for exchanging routing information on a multi-access network segment, such as a LAN1. The DR election process is based on the following criteria, in order of precedence1:
? In your scenario, two routers share the same highest priority and start time. This means that they have equal chances of becoming the DR based on the first and third criteria. Therefore, the second criterion will be used to break the tie, which is the router ID. The router with the highest router ID will become the DR, and the other router will become the backup designated router (BDR), which is ready to take over the role of DR if it fails1.

NEW QUESTION 14
What is the default MAC age-out timer on an EX Series switch?

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

Answer: D

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

NEW QUESTION 15
Which two types of tunnels are able to be created on all Junos devices? (Choose two.)

A. STP
B. GRE
C. IP-IP
D. IPsec

Answer: BD

Explanation:
Junos devices support various types of tunnels for different purposes12.
? Option B is correct. Generic Routing Encapsulation (GRE) is a tunneling protocol that can encapsulate a wide variety of network layer protocols inside virtual point- to-point links over an Internet Protocol network1. Junos devices support GRE tunnels1.
? Option D is correct. IPsec (Internet Protocol Security) is a protocol suite for securing Internet Protocol (IP) communications by authenticating and encrypting each IP packet of a communication session1. Junos devices support IPsec tunnels1.
? Option A is incorrect. Spanning Tree Protocol (STP) is not a type of tunnel. It??s a network protocol designed to prevent loops in a bridged Ethernet local area network2.
? Option C is incorrect. While Junos devices do support IP-IP (also known as IP tunneling), it??s not supported on all Junos devices1.

NEW QUESTION 19
You implemented the MAC address limit feature with the shutdown action on all interfaces on your switch.
In this scenario, which statement is correct when a violation occurs?

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

Answer: A

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

NEW QUESTION 22
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 26
Which two statements correctly describe RSTP port roles? (Choose two.)

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

Answer: AD

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

NEW QUESTION 29
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 33
Exhibit

Exhibit

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

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

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

Answer: C

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

NEW QUESTION 37
Exhibit.

Exhibit

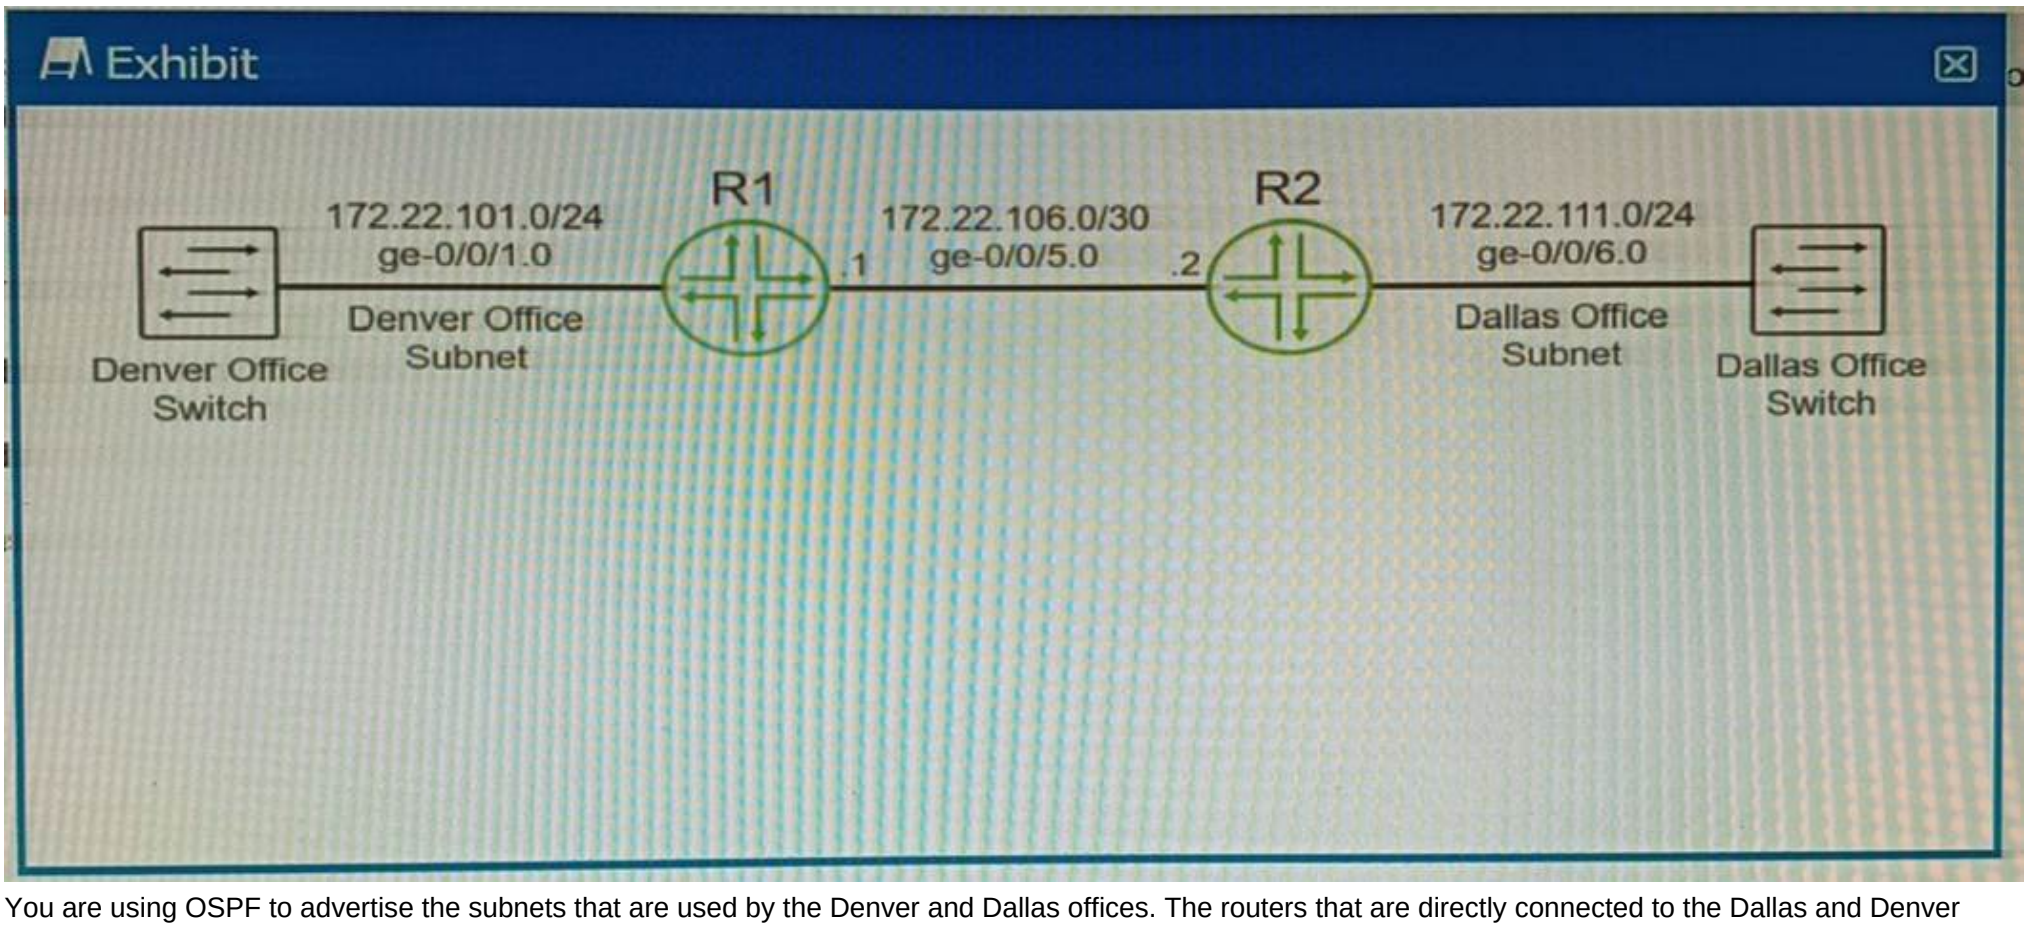

You are using OSPF to advertise the subnets that are used by the Denver and Dallas offices. The routers that are directly connected to the Dallas and Denver subnets are not advertising the connected subnets.
Referring to the exhibit, which two statements are correct? (Choose two.)

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

Answer: CD

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

NEW QUESTION 40
You deployed a new EX Series switch with DHCP snooping enabled and you do not see any entries in the snooping databases for an interface. Which two Juniper configurations for that interface caused this issue? (Choose two.)

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

Answer: AC

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

NEW QUESTION 41
Which statement is correct about graceful Routing Engine switchover (GRES)?

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

Answer: C

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

NEW QUESTION 44
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 45
Which statement about aggregate routes is correct?

A. Aggregate routes can only be used for static routing but not for dynamic routing protocols.
B. Aggregate routes are automatically generated for all of the subnets in a routing table.
C. Aggregate routes are always preferred over more specific routes, even when the specific routes have a better path.
D. Aggregate routes are used for advertising summarized network prefixes.

Answer: D

Explanation:
Aggregate routes are used for advertising summarized network prefixes12. They help minimize the number of routing tables in an IP network by consolidating selected multiple routes into a single route advertisement1. This approach is in contrast to non-aggregation routing, in which every routing table contains a unique entry for each route1.
Therefore, option D is correct. Options A, B, and C are not correct because:
? Aggregate routes can be used with both static routing and dynamic routing protocols1.
? Aggregate routes are not automatically generated for all of the subnets in a routing table. They need to be manually configured1.
? Aggregate routes are not always preferred over more specific routes. The route selection process in Junos OS considers several factors, including route preference and metric, before determining the active route1.

NEW QUESTION 49
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 52
.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

Money Back Guarantee

JN0-351 Practice Exam Features:

* JN0-351 Questions and Answers Updated Frequently

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

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

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