Juniper

Exam Questions JN0-351

Enterprise Routing and Switching - Specialist (JNCIS-ENT)

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

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

B)

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

C)

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

D)

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

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

Answer: C

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

NEW QUESTION 2
Exhibit.

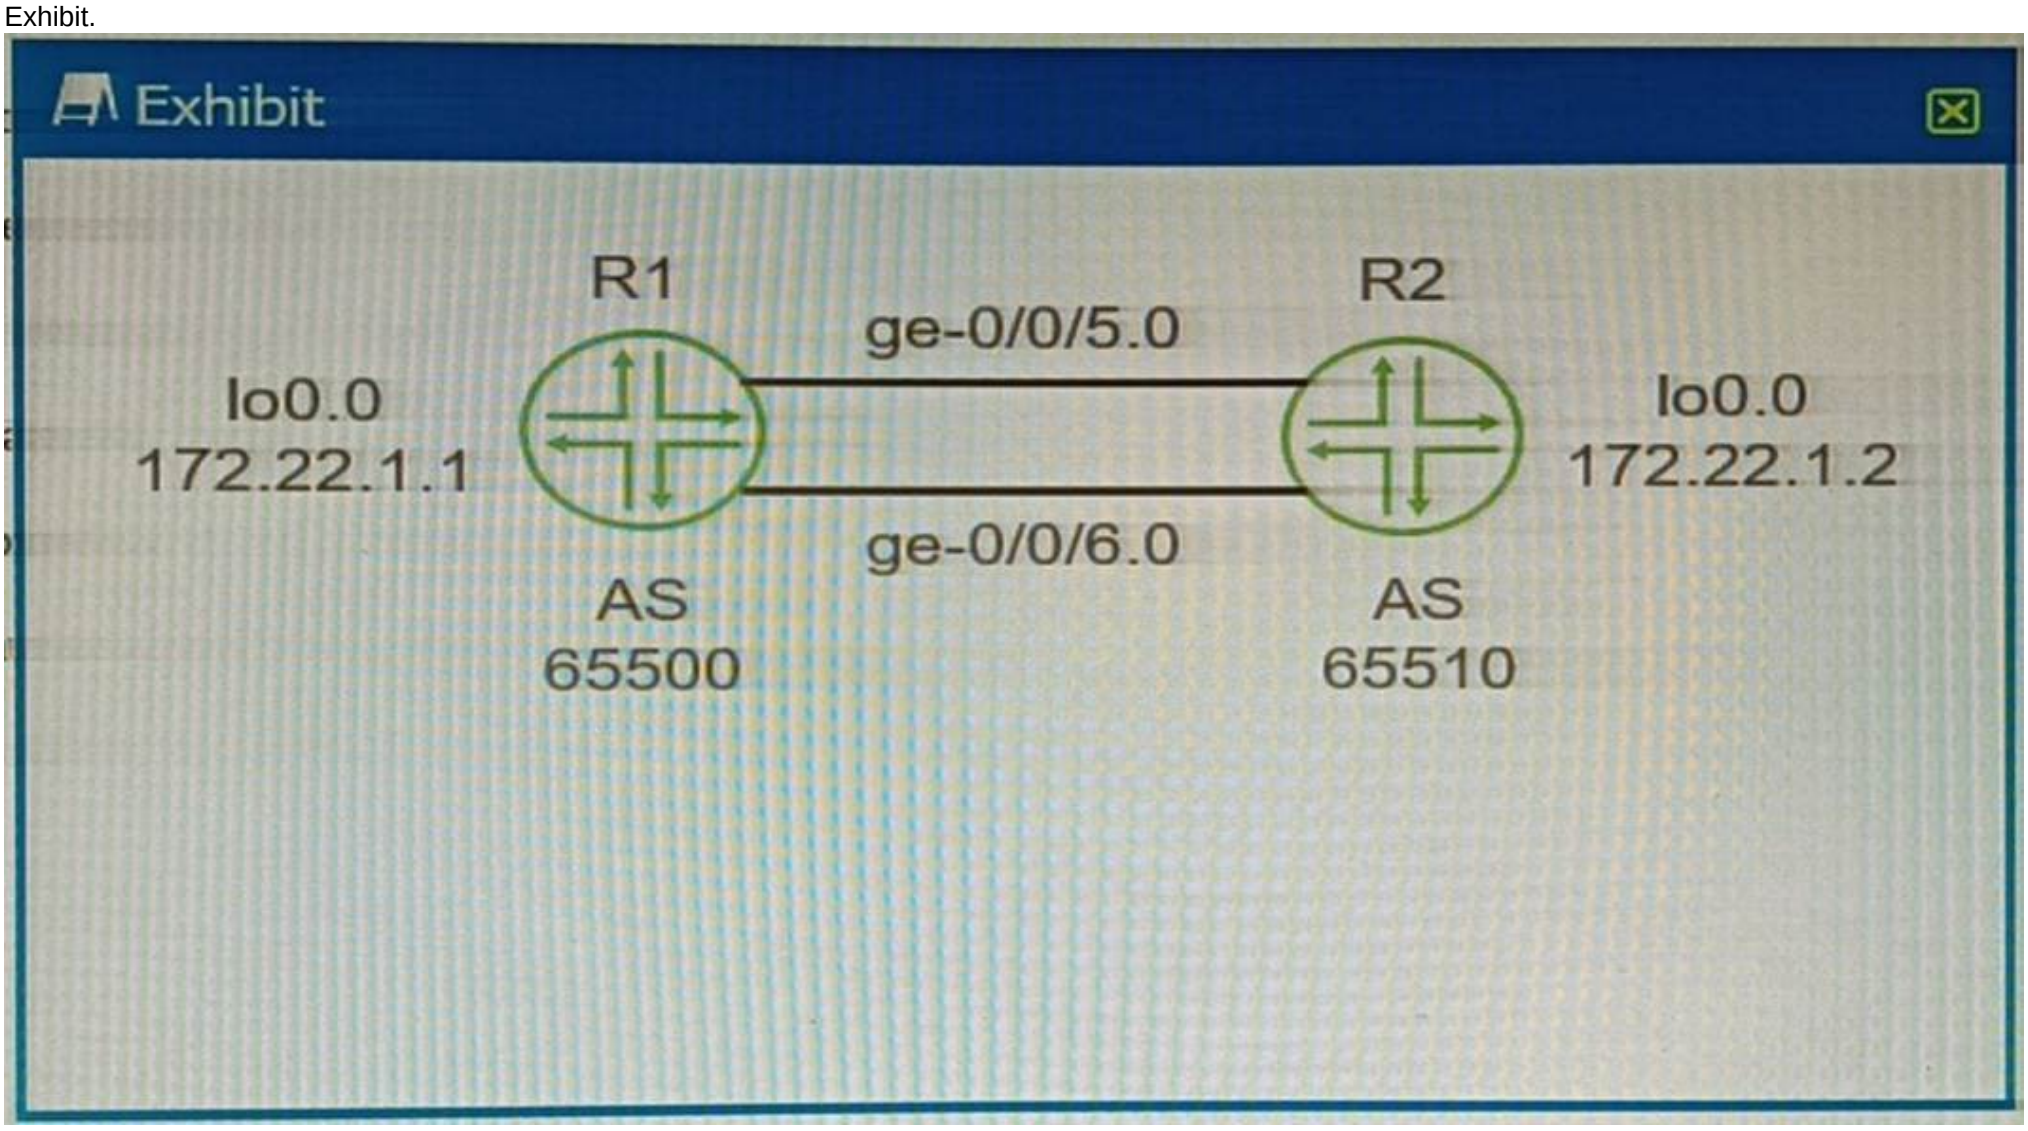

You want to enable redundancy for the EBGP peering between the two routers shown in the exhibit. Which three actions will you perform in this scenario? (Choose three.)

A. Configure BGP multihop.
B. Configure loopback interface peering.
C. Configure routes for the peer loopback interface IP addresses.

D. Configure an MD5 peer authentication.
E. Configure a cluster ID.

Answer: ABC

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

NEW QUESTION 3
Exhibit.

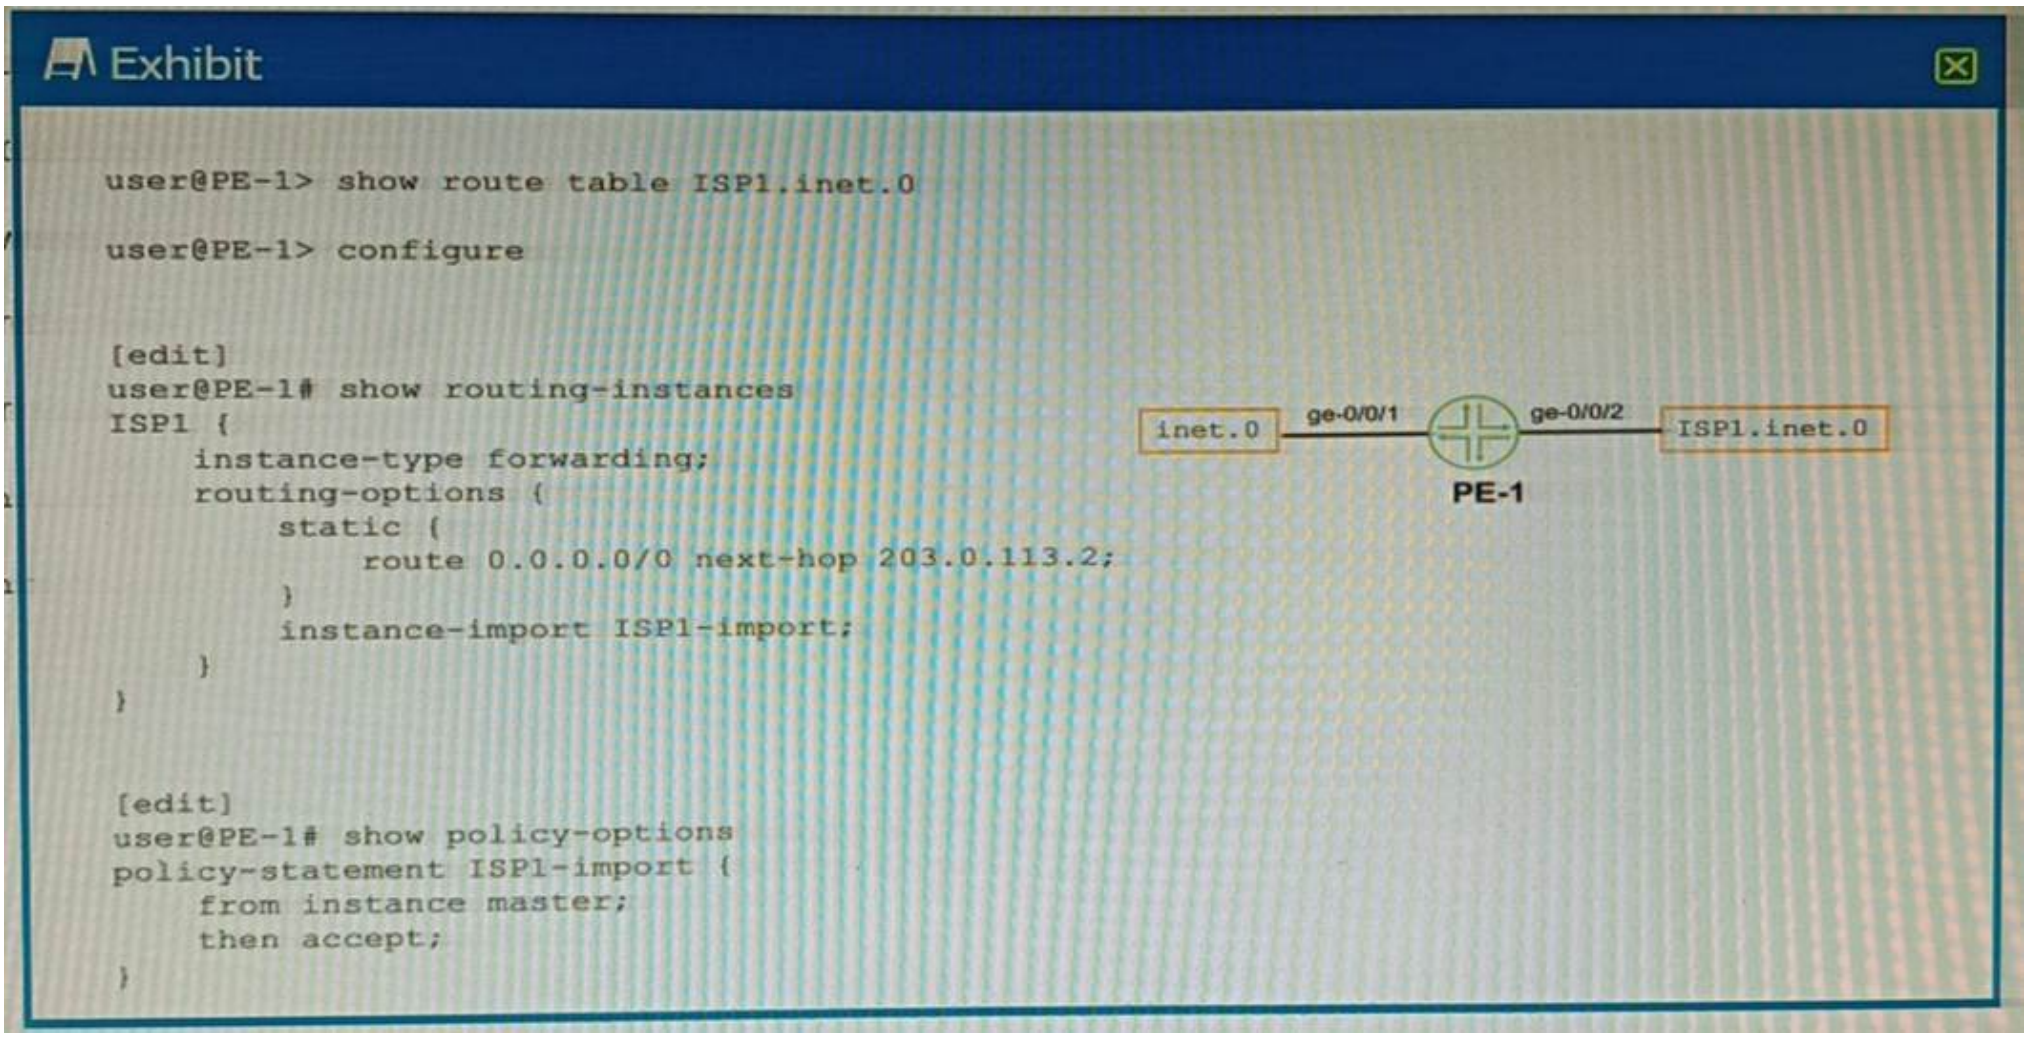

The ispi _ inet. 0 route table has currently no routes in it.
What will happen when you commit the configuration shown on the exhibit?

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 4
Exhibit

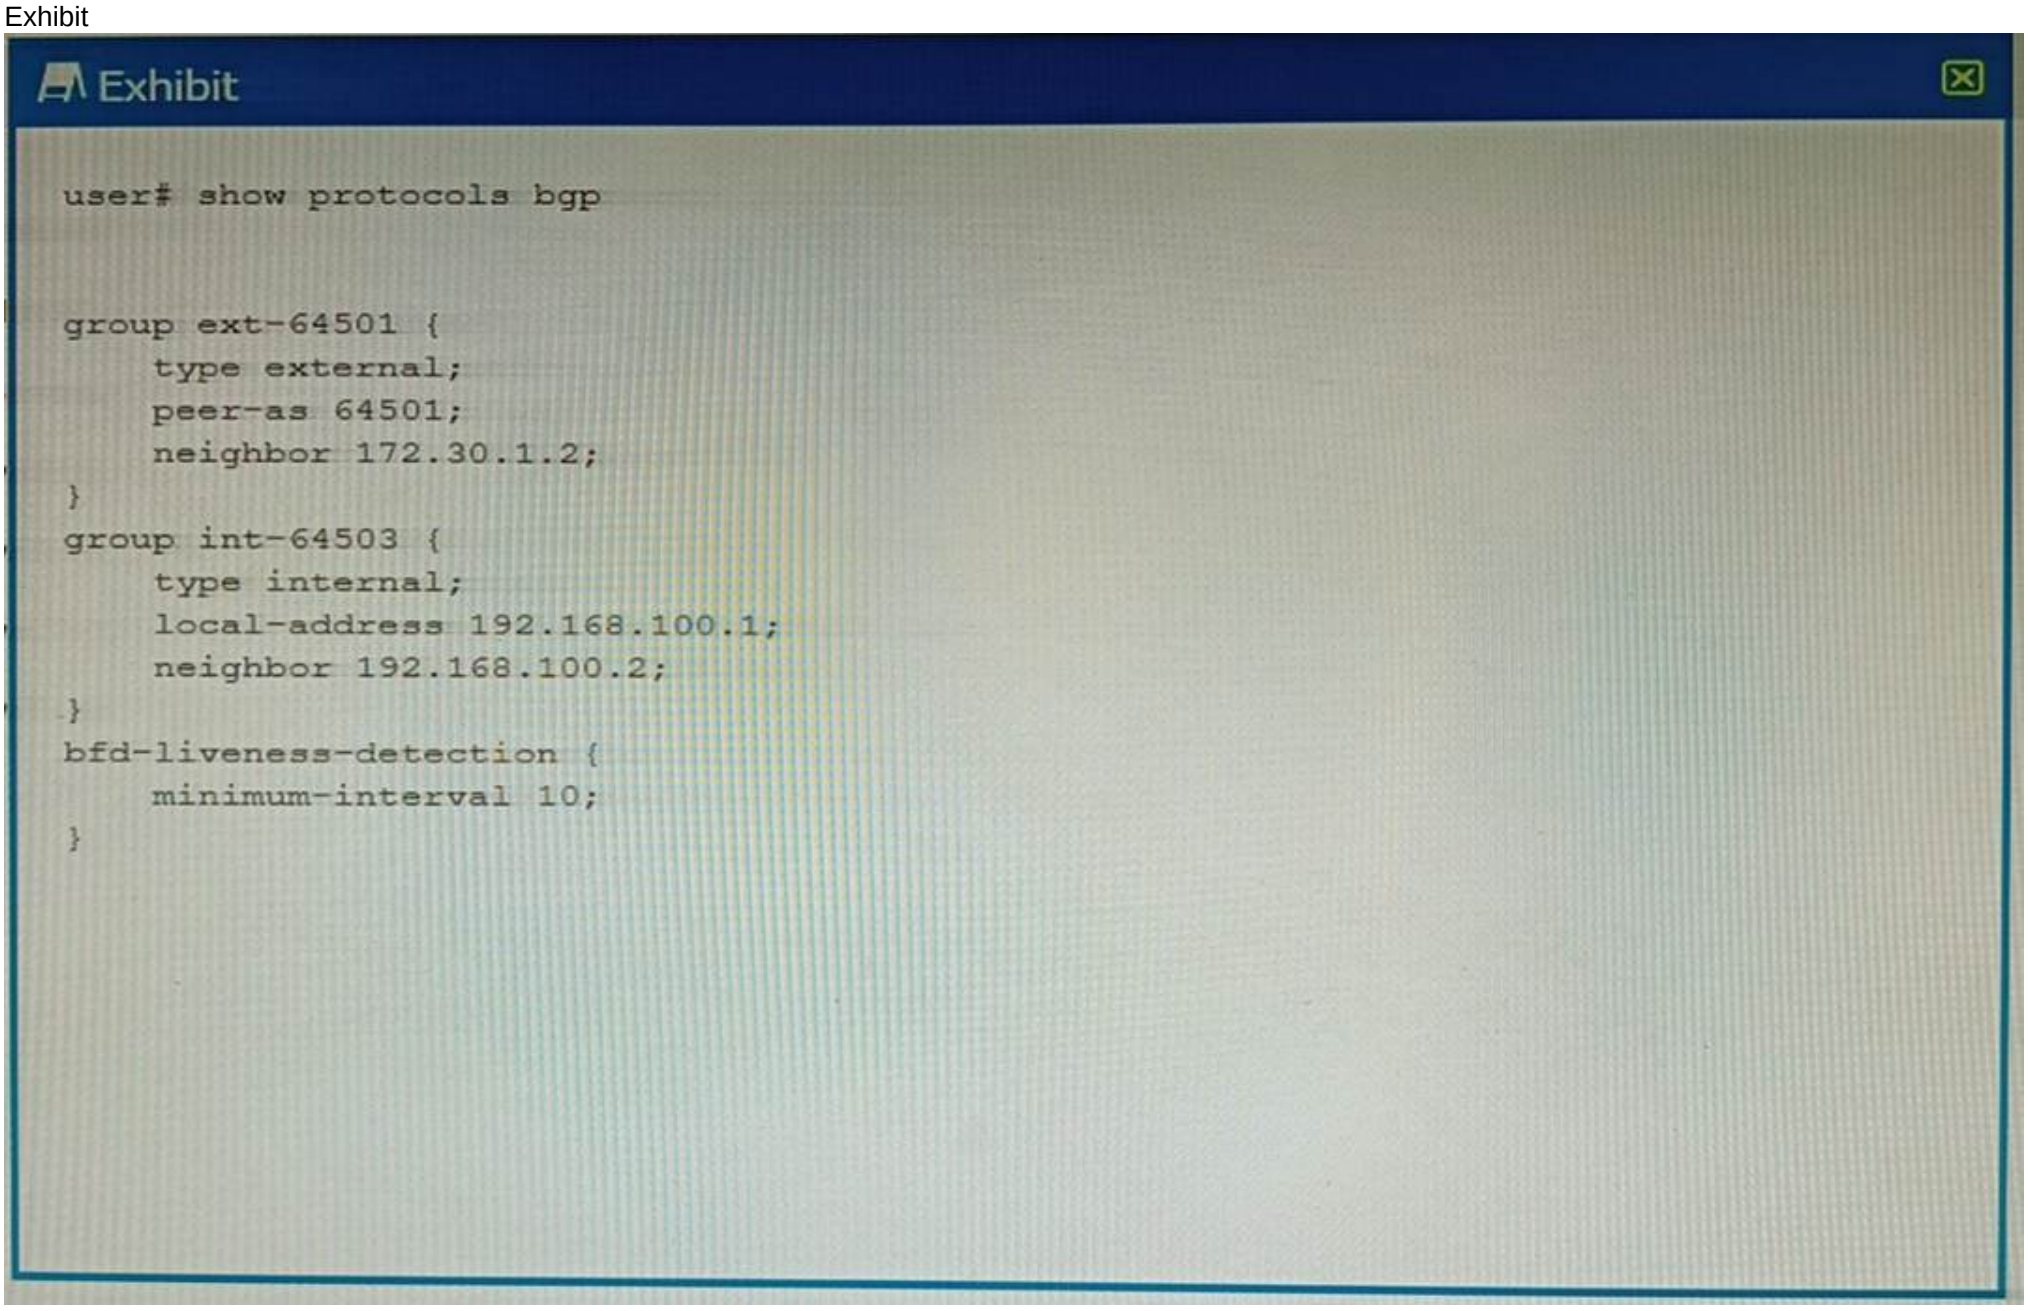

Your BGP neighbors, one in the USA and one in France, are not establishing a connection with each other.
Referring to the exhibit, which statement is correct?

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 BGPneighbors, which are not directly connected and require multiple hops to reach, the session mode must be multihop2.
? Since your BGP neighbors are in different countries, they are likely to be multihop neighbors. Therefore, you need to configure the session mode as multihop for each neighbor individually at the [edit protocols bgp group group-name neighbor address bfd-liveness-detection] hierarchy level2. For example:
protocols { bgp { group usa { neighbor 192.0.2.1 { bfd-liveness-detection { session-mode multihop; } } } group france { neighbor 198.51.100.1 { bfd-liveness-detection { session-mode multihop; } } } } }
? If you do not configure the session mode for multihop neighbors, BFD will use the
default mode of automatic, which will try to use single-hop mode and fail to establish a BFD session with the remote neighbor2. This will prevent BGP from using BFD to detect liveliness and failover.
? Therefore, the answer B is correct, as you need to configure the BFD liveness
detection on the BGP neighbor level with the appropriate session mode for multihop neighbors.

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

A. Redundant trunk groups use spanning tree to provide loop-free redundant uplinks.
B. Redundant trunk groups load balance traffic across two designated uplink interfaces.
C. Layer 2 control traffic is permitted on the secondary link.
D. If the active link fails, then the secondary link automatically takes over.

Answer: CD

Explanation:
? C is correct because Layer 2 control traffic is permitted on the secondary link of a redundant trunk group (RTG) on EX Series switches. Layer 2 control traffic includes protocols such as LLDP, LACP, and STP, which are used to exchange information and coordinate actions between switches1. According to the Juniper Networks documentation2, Layer 2 control traffic is allowed to pass through both the active and the secondary links of an RTG, but data traffic is only forwarded through the active link. This allows the switches to maintain their Layer 2 adjacencies and monitor the link status on both links.

? D is correct because if the active link fails, then the secondary link automatically
takes over in an RTG on EX Series switches. An RTG consists of two trunk links: an active or primary link, and a secondary or backup link2. The active link is used to forward data traffic, while the secondary link is in standby mode. If the active link fails or becomes unavailable, the secondary link immediately transitions to a forwarding state and takes over the data traffic without waiting for normal STP convergence2. This provides fast recovery and redundancy for the network.

NEW QUESTION 6
Exhibit

Exhibit

```
user@R1> show bgp neighbor
Peer: 10.32.1.2+63645 AS 65401 Local: 10.32.1.1+179 AS 65400
  Description: EBGP peering to 10.32.1.2
  Group: IPCLOS_eBGP                Routing-Instance: master
  Forwarding routing-instance: master
  Type: External    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ IPCLOS_BGP_EXP ] Import: [ IPCLOS_BGP_IMP ]
  Options: <Preference PeerAS Multipath LocalAS Refresh>
  Options: <VpnApplyExport MtuDiscovery MultipathAs BfdEnabled>
  Holdtime: 90 Preference: 170 Local AS: 65400 Local System AS: 0
  Number of flaps: 0
  Peer ID: 10.52.100.2      Local ID: 10.52.100.1        Active Holdtime: 90
  Keepalive Interval: 30          Group index: 0    Peer index: 0     SNMP
index: 0
  I/O Session Thread: bgpio-0 State: Enabled
  BFD: enabled, up
  Local Interface: ge-0/0/1.0
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  Restart flag received from the peer: Notification
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer does not support LLGR Restarter functionality
  Peer supports 4 byte AS extension (peer-as 65401)
  Peer does not support Addpath
  Table inet.0 Bit: 20000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              6
    Received prefixes:            9
    Accepted prefixes:            9
    Suppressed due to damping:    0
    Advertised prefixes:          22
Last traffic (seconds): Received 22   Sent 10   Checked 69617
Input messages:  Total 2568 Updates 4 Refreshes 0 Octets 48991
Output messages: Total 2572 Updates 8 Refreshes 0 Octets 49362
Output Queue[1]: 0              (inet.0, inet-unicast)
```

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

A. Peer 10.32.1.2 is configured for AS 63645.
B. The BGP session is not established.
C. The R1 is configured for AS 65400.
D. The routers are exchanging IPv4 routes.

Answer: BC

Explanation:
Option B suggests that the BGP session is not established. This is correct because in the output, the state of the BGP session is shown as ??Idle??. In BGP, an ??Idle?? state means that the BGP session is not currently established1.
Option C suggests that R1 is configured for AS 65400. This is also correct because in the output, it??s shown that the local AS number is 654001. The local AS number represents the Autonomous System (AS) number of the router on which you??re checking the BGP session1.

NEW QUESTION 7
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.

NEW QUESTION 8
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 9
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.

NEW QUESTION 10
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 10

Refer to the exhibit.

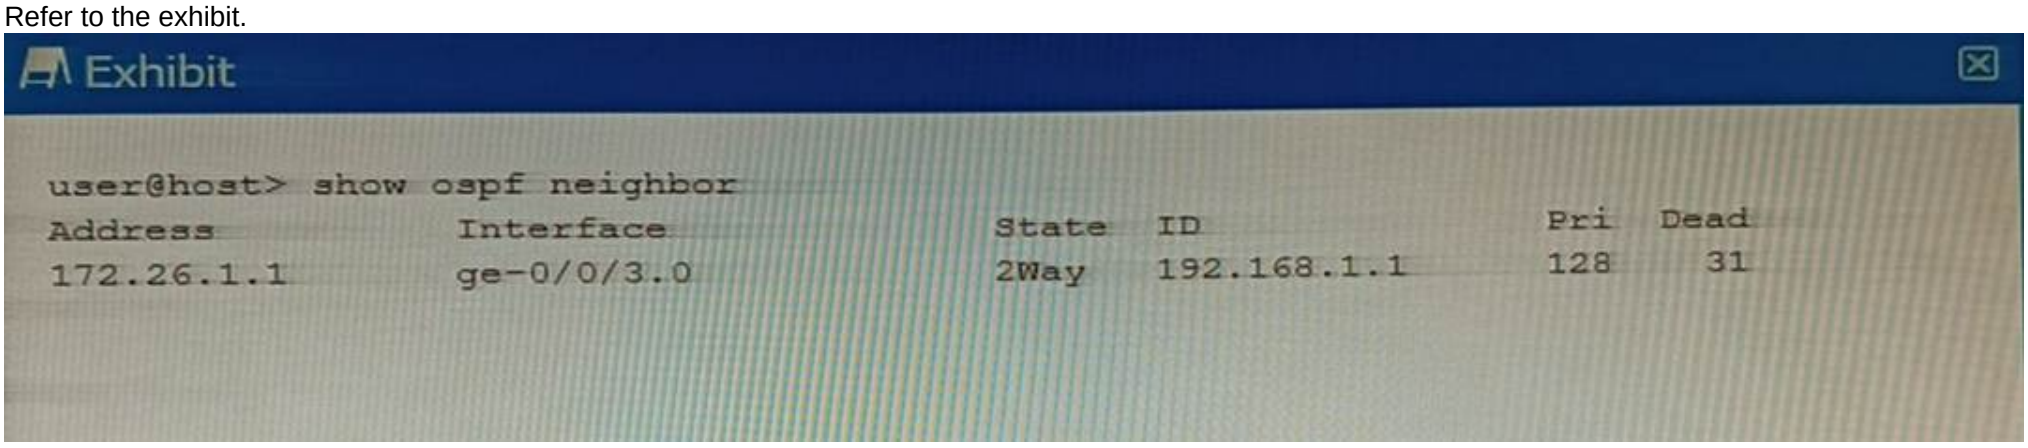

Referring to the output shown in the exhibit, which statement is correct?

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

Answer: B

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

NEW QUESTION 11
Which two types of tunnels are able to be created on all Junos devices? (Choose two.)

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

Answer: BD

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

NEW QUESTION 16
Which two events cause a router to advertise a connected network to OSPF neighbors? (Choose two.)

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

Answer: AD

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

NEW QUESTION 18
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 21
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 24
Which two statements are true about the default VLAN on Juniper switches? (Choose two.)

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

Answer: AD

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

NEW QUESTION 26
Which statement is correct about controlling the routes installed by a RIB group?

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

Answer: A

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

NEW QUESTION 31
Exhibit

Exhibit

```
user@host# show
  protocols {
        oam {
              gre-tunnel {
                    interface gr-1/1/10.1 {
                          keepalive-time 10;
                                hold-time 10;
                    }
              }
        }
        lldp {
              interface all;
        }
  }
```

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 hold-time 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.
References:
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

NEW QUESTION 33
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 36
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 38
.....

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