

Exam Code: JN0-351

Exam Name: Juniper Enterprise Routing and Switching, Specialist



Exam A

QUESTION 1

What is the default keepalive time for BGP?

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

Correct Answer: B

Section:

Explanation:

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

QUESTION 2

Which two statements are correct about tunnels? (Choose two.)

- A. BFD cannot be used to monitor tunnels.
- B. Tunnel endpoints must have a valid route to the remote tunnel endpoint.
- C. IP-IP tunnels are stateful.
- D. Tunnels add additional overhead to packet size.



Correct Answer: B, D

Section:

Explanation:

A tunnel is a connection between two computer networks, in which data is sent from one network to another through an encrypted link. Tunnels are commonly used to secure data communications between two networks or to connect two networks that use different protocols.

Option B is correct, because tunnel endpoints must have a valid route to the remote tunnel endpoint. A tunnel endpoint is the device that initiates or terminates a tunnel connection. For a tunnel to be established, both endpoints must be able to reach each other over the underlying network. This means that they must have a valid route to the IP address of the remote endpoint.

Option D is correct, because tunnels add additional overhead to packet size. Tunnels work by encapsulating packets: wrapping packets inside of other packets. This means that the original packet becomes the payload of the surrounding packet, and the surrounding packet has its own header and trailer. The header and trailer of the surrounding packet add extra bytes to the packet size, which is called overhead. Overhead can reduce the efficiency and performance of a network, as it consumes more bandwidth and processing power.

Option A is incorrect, because BFD can be used to monitor tunnels. BFD is a protocol that can be used to quickly detect failures in the forwarding path between two adjacent routers or switches. BFD can be integrated with various routing protocols and link aggregation protocols to provide faster convergence and fault recovery. BFD can also be used to monitor the connectivity of tunnels, such as GRE, IPsec, or MPLS.

Option C is incorrect, because IP-IP tunnels are stateless. IP-IP tunnels are a type of tunnels that use IP as both the encapsulating and encapsulated protocol. IP-IP tunnels are simple and easy to configure, but they do not provide any security or authentication features. IP-IP tunnels are stateless, which means that they do not keep track of the state or status of the tunnel connection. Stateless tunnels do not require any signaling or negotiation between the endpoints, but they also do not provide any error detection or recovery mechanisms.

1: What is Tunneling? | Tunneling in Networking 2: What Is Tunnel In Networking, Its Types, And Its Benefits?: [Configuring Bidirectional Forwarding Detection] : [IP-IP Tunneling]

QUESTION 3

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.

Correct Answer: A

Section:

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.
1:IP-IP Tunneling2:What is tunneling? | Tunneling in networking3: IPv4 - Header

QUESTION 4

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.

Correct Answer: A, D

Section:

Explanation:

Option A suggests that the MTU is not at least 1492 bytes. This is correct because IS-IS requires a minimum MTU of 1492 bytes to establish adjacencies¹. If the MTU is less than this, IS-IS adjacencies will not be established¹.

Option D suggests that the lo0 interface is not included as an IS-IS interface. This is also correct because the loopback interface (lo0) is typically used as the router ID in IS-IS¹. If the loopback interface is not included in IS-IS, it could prevent IS-IS adjacencies from being established¹.

Therefore, options A and D are correct.

QUESTION 5

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.

Correct Answer: C, D

Section:

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.

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

QUESTION 6

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

Correct Answer: D

Section:**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.

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

QUESTION 7

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
```


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.

Correct Answer: B, C

Section:

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 established.

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 65400. The local AS number represents the Autonomous System (AS) number of the router on which you're checking the BGP session.

QUESTION 8

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

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

Correct Answer: D

Section:

Explanation:

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

QUESTION 9

You are a network operator who wants to add a second ISP connection and remove the default route to the existing ISP. You decide to deploy the BGP protocol in the network.

What two statements are correct in this scenario? (Choose two.)

- A. IBGP updates the next-hop attribute to ensure reachability within an AS.
- B. IBGP peers advertise routes received from EBGP peers to other IBGP peers.
- C. IBGP peers advertise routes received from IBGP peers to other IBGP peers.
- D. EBGP peers advertise routes received from IBGP peers to other EBGP peers.

Correct Answer: A, B

Section:

Explanation:

A is correct because IBGP updates the next-hop attribute to ensure reachability within an AS. This is because the next-hop attribute is the IP address of the router that advertises the route to a BGP peer. If the next-hop attribute is not changed by IBGP, it would be the IP address of an external router, which may not be reachable by all routers within the AS. Therefore, IBGP updates the next-hop attribute to the IP address of the



router that received the route from an EBGp peer¹.

It's correct because IBGP peers advertise routes received from EBGp peers to other IBGP peers. This is because BGP follows the rule of advertising only the best route to a destination, and EBGp routes have a higher preference than IBGP routes. Therefore, IBGP peers advertise routes learned from an EBGp peer to all BGP peers, including both EBGp and IBGP peers¹.

QUESTION 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.

Correct Answer: C, D

Section:

Explanation:

In BGP, routing policies are used to control the flow of routing information between BGP peers¹.

Option C suggests that import policies are applied after the RIB-In table. This is correct because import policies in BGP are applied to routes that are received from a BGP peer, before they are installed in the local BGP Routing Information Base (RIB-In)¹. The RIB-In is a database that stores all the routes that are received from all peers¹.

Option D suggests that export policies are applied after the RIB-Local table. This is correct because export policies in BGP are applied to routes that are being advertised to a BGP peer, after they have been selected from the local BGP Routing Information Base (RIB-Local)¹. The RIB-Local is a database that stores all the routes that the local router is using¹.

Therefore, options C and D are correct.

QUESTION 11

You are asked to connect an IP phone and a user computer using the same interface on an EX Series switch. The traffic from the computer does not use a VLAN tag, whereas the traffic from the IP phone uses a VLAN tag.

Which feature enables the interface to receive both types of traffic?

- A. native VLAN
- B. DHCP snooping
- C. MAC limiting
- D. voice VLAN

Correct Answer: D

Section:

Explanation:

The feature that enables an interface on an EX Series switch to receive both untagged traffic (from the computer) and tagged traffic (from the IP phone) is the voice VLAN¹².

The voice VLAN feature in EX-series switches enables access ports to accept both data (untagged) and voice (tagged) traffic and separate that traffic into different VLANs¹². This allows the switch to differentiate between voice and data traffic, ensuring that voice traffic can be treated with a higher priority¹². Therefore, option D is correct.

QUESTION 12

Exhibit

```
Routing table: default.ethernet-switching
```

```
ETHERNET-SWITCHING:
```

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		dscd	66	1	
2, *	user	0		comp	1304	2	
2, *	intf	0		rslv	1302	1	
2, 00:26:88:02:74:86	user	0		ucst	1303	3	ge-0/0/6.0
2, 00:26:88:02:74:87	user	0		ucst	1305	3	ge-0/0/7.0
2, 00:26:88:02:74:88	user	0		ucst	1306	3	ge-0/0/8.0



Which command displays the output shown in the exhibit?

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

Correct Answer: B

Section:

Explanation:

The output shown in the exhibit is a brief display of the Ethernet switching table, which shows the learned Layer 2 MAC addresses for each VLAN and interface1.

The command `show ethernet-switching table` displays the Ethernet switching table with brief information, such as the destination MAC address, the VLAN name, the forwarding state, and the interface name1.

The command `show route forwarding-table` displays the routing table information for each protocol family, such as inet, inet6, mpls, iso, and so on2. It does not show the Ethernet switching table or the MAC addresses.

The command `show ethernet-switching table extensive` displays the Ethernet switching table with extensive information, such as the destination MAC address, the VLAN name, the forwarding state, the interface name, the VLAN index, and the tag type1. It shows more details than the brief output shown in the exhibit.

The command `show route forwarding-table family ethernet-switching` displays the routing table information for the ethernet-switching protocol family, which shows the destination MAC address, the next-hop MAC address, and the interface name3. It does not show the VLAN name or the forwarding state.

QUESTION 13

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.

Correct Answer: A, C

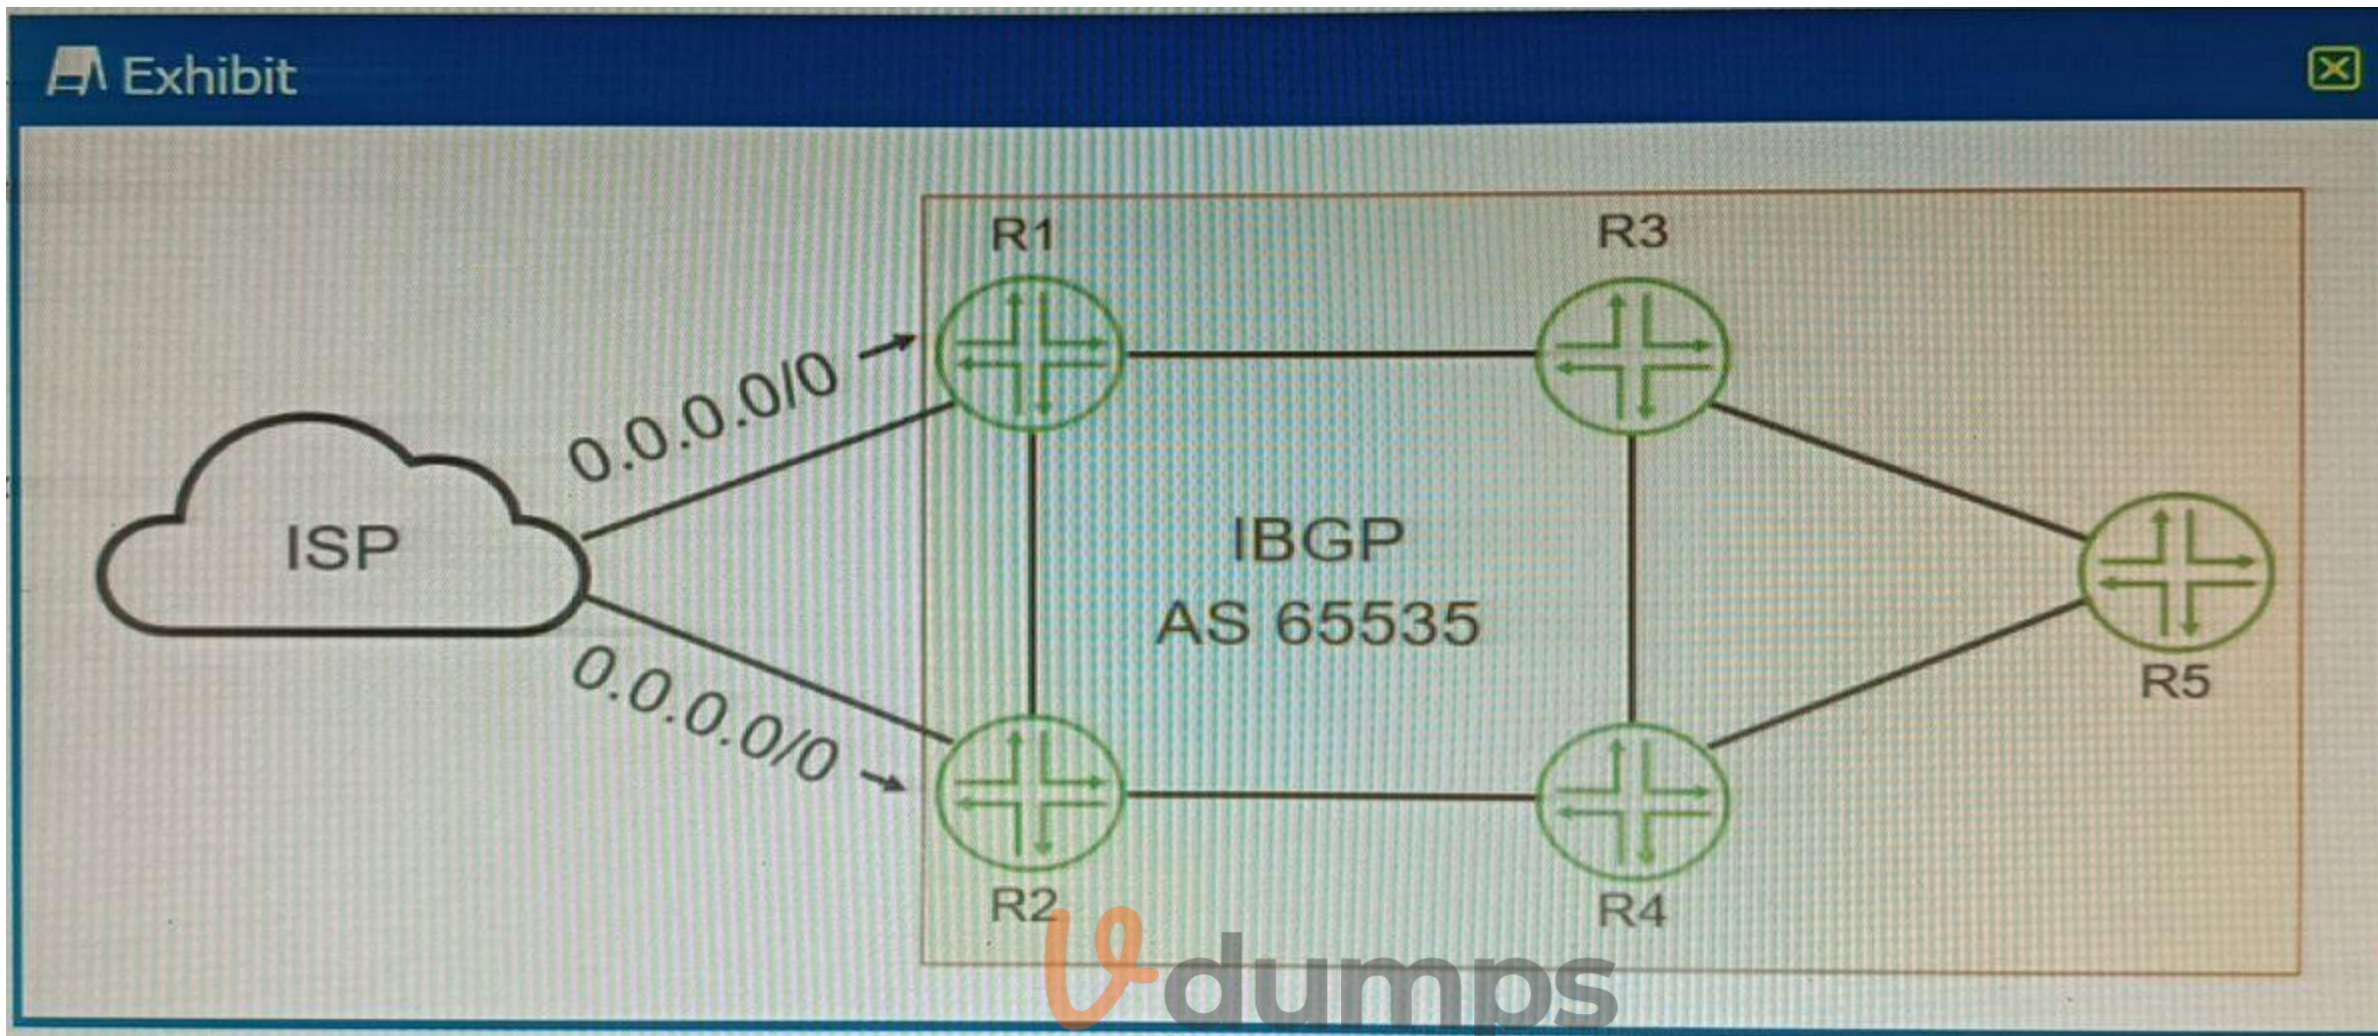
Section:**Explanation:**

A is correct because the interface is configured as a disabled port. A disabled port does not forward any traffic, including DHCP packets. Therefore, DHCP snooping cannot learn any MAC addresses or lease information from a disabled port1.

C is correct because the interface is configured as a trunk port. By default, all trunk ports on the switch are trusted for DHCP snooping2. This means that DHCP snooping does not inspect or filter any DHCP packets received on a trunk port. Therefore, DHCP snooping does not add any entries to the snooping database for a trunk port2.

QUESTION 14

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

Correct Answer: C

Section:

Explanation:

The BGP attribute that you would use to forward all Internet traffic through the R1 device is the local preference.

The local preference is an attribute that is used within an autonomous system (AS) and exchanged between iBGP routers. It is used to select an exit point from the AS. The path with the highest local preference is preferred. By setting a higher local preference for the routes received from R1, you can make R1 the preferred exit point for all Internet traffic.

QUESTION 15

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.

Correct Answer: A, C

Section:

Explanation:

A is correct because RSTP alternate ports block traffic while receiving superior BPDUs from a neighboring switch. An alternate port is a backup port for a root port, which means it receives better BPDUs from another bridge than the current root port¹. However, an alternate port does not forward any traffic, as it is in a discarding state². It only listens to BPDUs and waits for the root port to fail. If the root port fails, the alternate port can immediately transition to a forwarding state and become the new root port¹.

C is correct because RSTP alternate ports provide an alternate higher cost path to the root bridge. An alternate port is selected based on the same criteria as the root port, which are the lowest bridge ID, the lowest path cost, the lowest sender port ID, and the lowest receiver port ID³. However, an alternate port receives a higher cost BPDU than the root port, otherwise it would be the root port itself¹. Therefore, an alternate port provides an alternate higher cost path to the root bridge than the root port.

QUESTION 16

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

Correct Answer: A, C

Section:

Explanation:

BGP attributes are properties that BGP uses for route advertisement, path selection, and loop prevention¹. There are four categories of BGP attributes^{2,3}:

Well-known mandatory: Must be recognized by all BGP routers, present in all BGP updates, and passed on to other BGP routers^{2,3}.

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 update^{2,3}. These include the AS path and next hop attributes^{2,3}. Therefore, options A and C are correct.

QUESTION 17

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.

Correct Answer: B

Section:

Explanation:

The exhibit shows the configuration of BFD liveness detection for BGP at the global level, which applies to all BGP neighbors by default¹. However, this configuration does not specify the session mode, which determines whether BFD uses single-hop or multihop mode to communicate with a neighbor².

For single-hop BGP neighbors, which are directly connected on the same subnet, the session mode can be either automatic or single-hop. For multihop BGP neighbors, which are not directly connected and require multiple hops to reach, the session mode must be multihop².

Since your BGP neighbors are in different countries, they are likely to be multihop neighbors. Therefore, you need to configure the session mode as multihop for each neighbor individually at the [edit protocols bgp group group-name neighbor address bfd-liveness-detection] hierarchy level². For example:

```
protocols { bgp { group usa { neighbor 192.0.2.1 { bfd-liveness-detection { session-mode multihop; } } } group france { neighbor 198.51.100.1 { bfd-liveness-detection { session-mode multihop; } } } }
```

If you do not configure the session mode for multihop neighbors, BFD will use the default mode of automatic, which will try to use single-hop mode and fail to establish a BFD session with the remote neighbor².

This will prevent BGP from using BFD to detect liveness and failover.

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

QUESTION 18

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.

Correct Answer: A, D

Section:

Explanation:

On Juniper switches, the default VLAN is set to a VLAN ID of 1 by default¹². This means that all interfaces on the switch are members of VLAN 1 until they are specifically assigned to another VLAN¹². Therefore, option A is correct.

The default VLAN ID can be changed¹². This allows network administrators to configure the switch to use a different VLAN as the default, if necessary¹². Therefore, option D is correct.

QUESTION 19

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.

Correct Answer: B, C

Section:

Explanation:

B is correct because duplex and speed settings are required to match on both participating devices. According to the Juniper Networks documentation¹, all the interfaces in a LAG must have the same speed and be in full-duplex mode. This ensures that the LAG can operate as a single logical link without any performance or compatibility issues.

C is correct because member links are not required to be contiguous ports. According to the Juniper Networks documentation², you can group any Ethernet interfaces on a switch into a LAG, regardless of their physical location or slot number. This provides flexibility and scalability for configuring LAGs on switches.

QUESTION 20

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*

Vdumps

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.

Correct Answer: C

Section:

Explanation:

The exhibit shows the output of the command `show vlans brief`, which displays brief information about VLANs and their associated interfaces¹.

The output has four columns: Routing instance, VLAN name, Interfaces, and Tagging.

The * symbol indicates that the interface is active, meaning that it is up and forwarding traffic¹. This can be verified by the command `show interfaces terse`, which displays the status of the interfaces².

QUESTION 21

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

Which statement is correct in this scenario?

- A. At least three routers are required for a DR/BDR election
- B. Router IDs must match for an adjacency to form.
- C. R2 will be the BDR.
- D. R1 will be the BDR.

Correct Answer: D

Section:

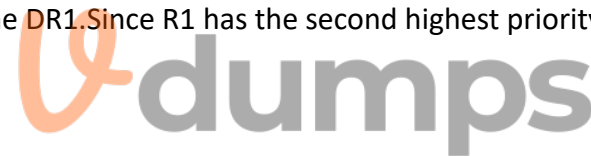
Explanation:

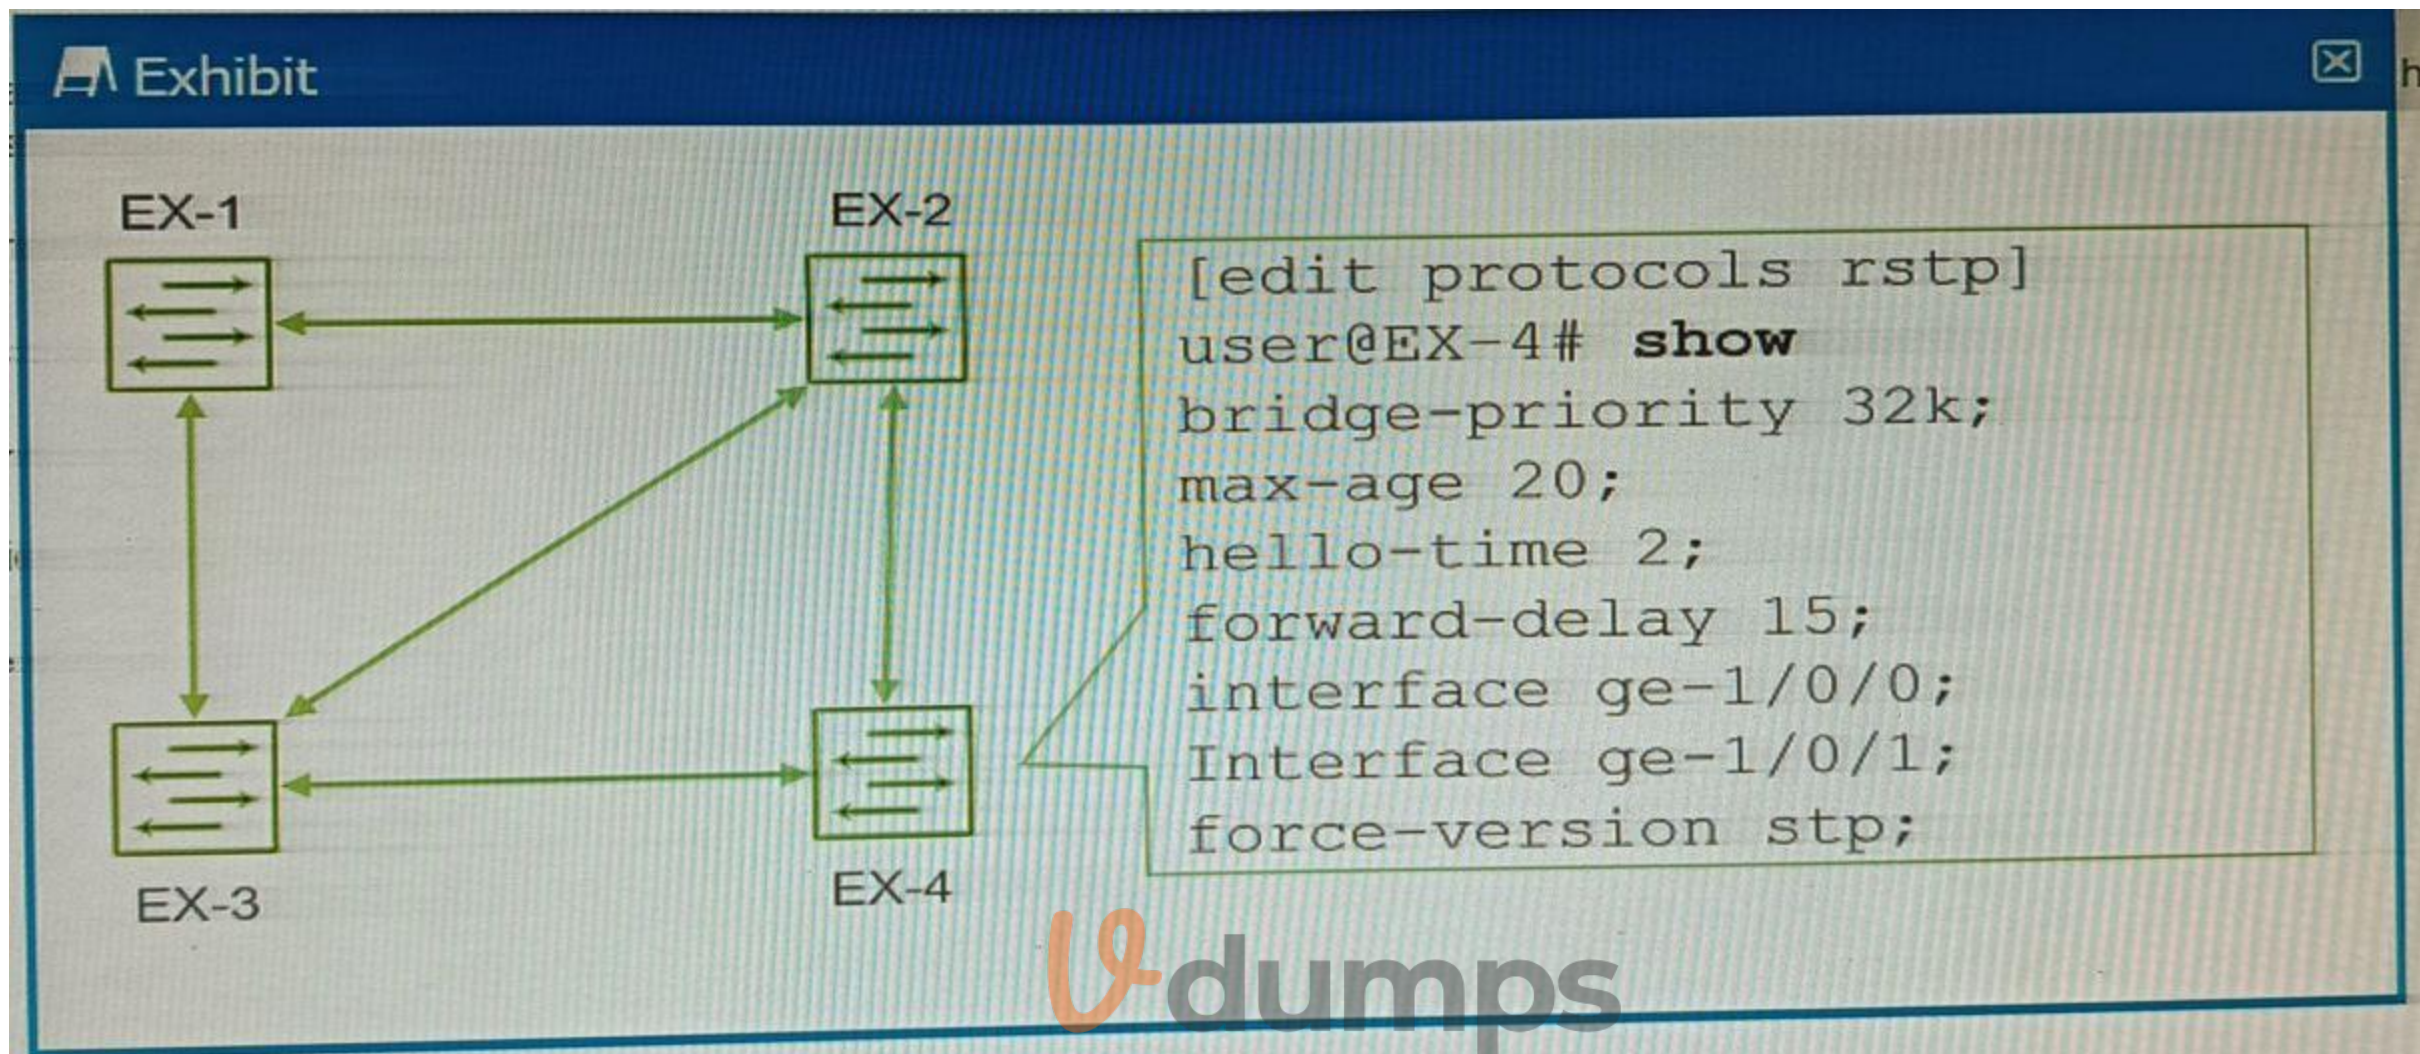
In OSPF, the Designated Router (DR) and Backup Designated Router (BDR) are elected based on the priority of the routers¹. The router with the highest priority becomes the DR, and the router with the second highest priority becomes the BDR¹. If there is a tie in priority, then the router with the highest Router ID is chosen¹.

In this scenario, R2 has a higher priority (64) than R1 (32), so R2 will become the DR¹. Since R1 has the second highest priority, it will become the BDR¹. Therefore, option D is correct.

QUESTION 22

Exhibit.





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

In this scenario, which action would solve the delay in RSTP convergence?

- A. The hello-time must be increased.
- B. The force-version must be removed.
- C. The bridge priority for EX-4 must be set at 4000.
- D. The max-age must be increased to 20

Correct Answer: B

Section:

Explanation:

The exhibit shows the configuration of RSTP on EX-4, which has the command `force-version stp`. This command forces the switch to use the legacy STP protocol instead of RSTP, even though the switch supports RSTP. This means that EX-4 will not be able to take advantage of the faster convergence and enhanced features of RSTP, such as edge ports, link type, and proposal/agreement sequence. The other switches in the network are likely to be running RSTP, as it is the default protocol for EX Series switches. Therefore, there will be a compatibility issue between EX-4 and the other switches, which will result in longer convergence times and suboptimal performance. The switch will also generate a warning message that says "Warning: STP version mismatch with neighbor" when it receives a BPDU from a RSTP neighbor.

To solve this problem, the `force-version` command must be removed from EX-4, so that it can run RSTP natively and interoperate with the other switches in the network. This will enable faster convergence and better stability for the network topology. To remove the command, you can use the `delete protocols rstp force-version` command in configuration mode.

QUESTION 23

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.

Correct Answer: A, D

Section:

Explanation:

In Rapid Spanning Tree Protocol (RSTP), there are several port roles that determine the behavior of the port in the spanning tree.

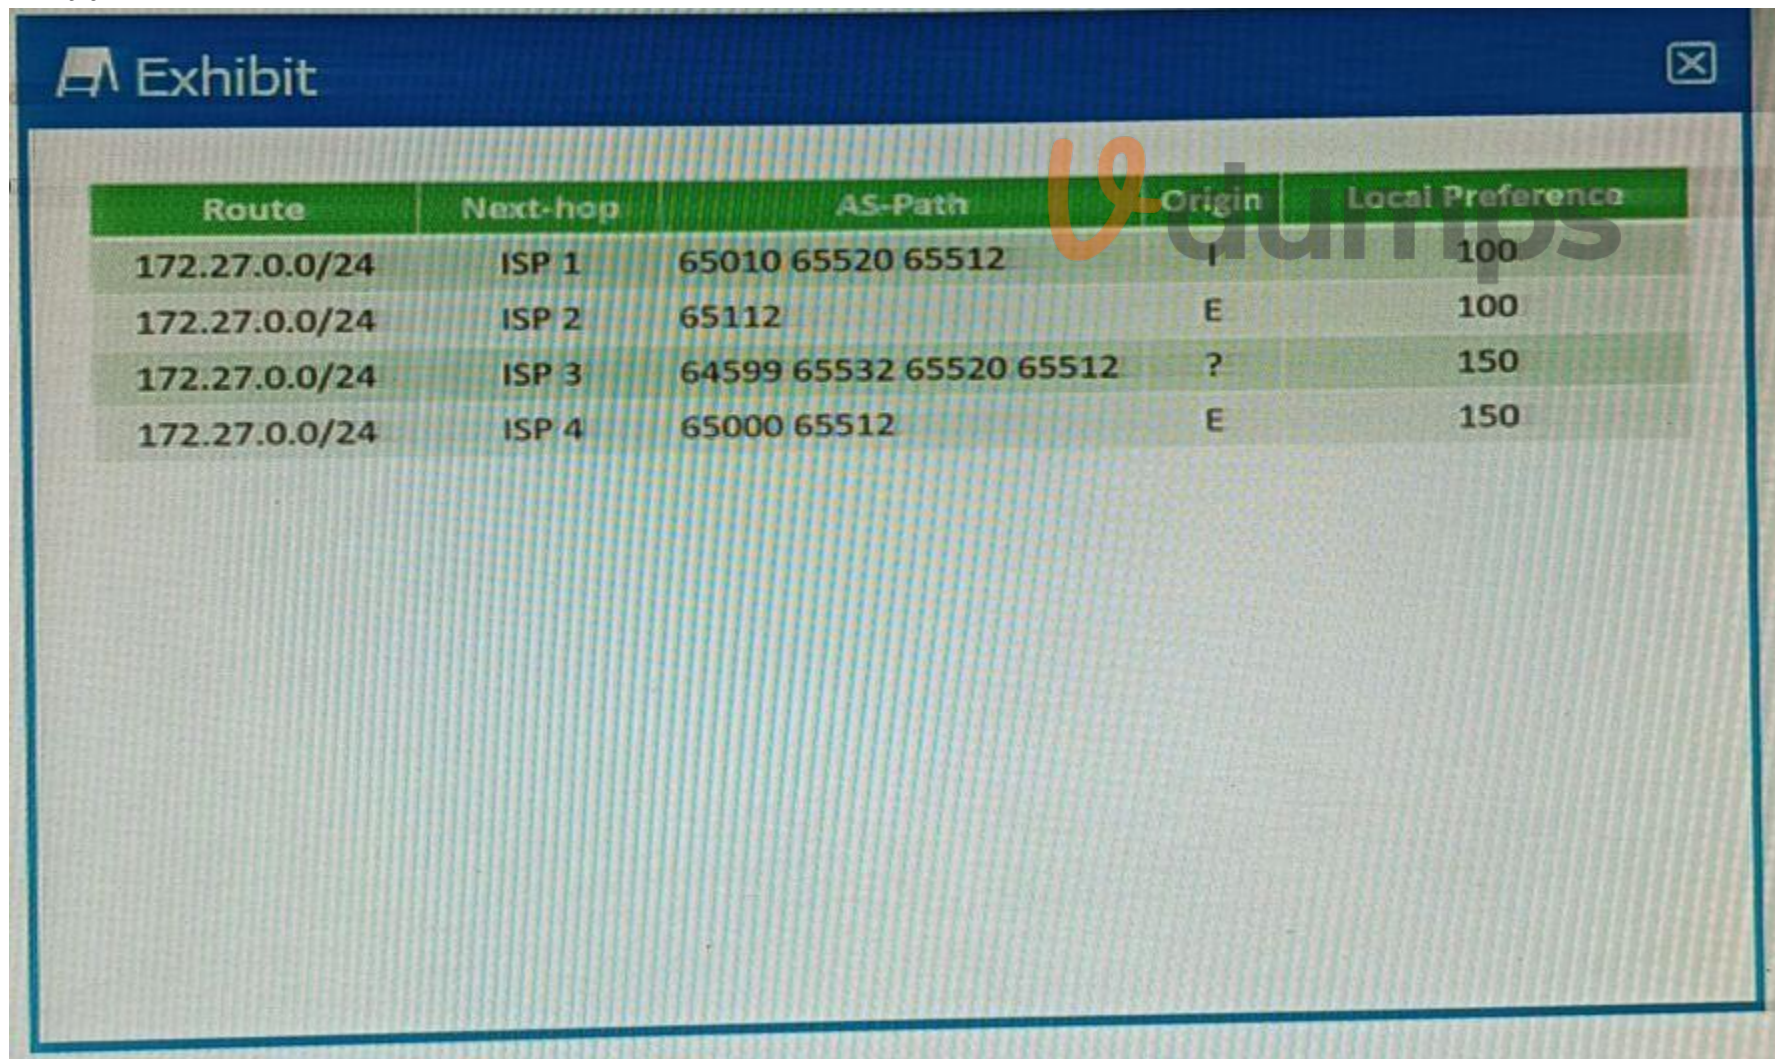
Option A suggests that the designated port forwards data to the downstream network segment or device. This is correct because the designated port is the port on a network segment that has the best path to the root bridge. It's responsible for forwarding frames towards the root bridge and sending configuration messages into its segment.

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 bridge. It's used to forward traffic towards the root bridge.

Therefore, options A and D are correct.

QUESTION 24

Exhibit



The exhibit shows a BGP route table with the following columns: Route, Next-hop, AS-Path, Origin, and Local Preference. The table lists four entries for the route 172.27.0.0/24, each received from a different ISP.

Route	Next-hop	AS-Path	Origin	Local Preference
172.27.0.0/24	ISP 1	65010 65520 65512	I	100
172.27.0.0/24	ISP 2	65112	E	100
172.27.0.0/24	ISP 3	64599 65532 65520 65512	?	150
172.27.0.0/24	ISP 4	65000 65512	E	150

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

- A. ISP1
- B. ISP 3
- C. ISP 4
- D. ISP 2

Correct Answer: C

Section:

Explanation:

In BGP, the path selection process is based on a set of attributes. The process starts by preferring the path with the highest weight, then the highest local preference, then the locally originated routes, and so on. If all these attributes are the same, then it prefers the path with the shortest AS path.

Referring to the exhibit, all four ISPs have the same weight, local preference, and origin. However, ISP 4 has the shortest AS path. Therefore, ISP 4 will be selected as the active path. So, option C is correct.

QUESTION 25

Exhibit.



Why is this OSPF adjacency remaining in this state?

- A. A subnet mask mismatch exists between the OSPF neighbors.
- B. An MTU mismatch exists between the OSPF neighbors.
- C. A hello interval mismatch exists between the OSPF neighbors.
- D. An area ID mismatch exists between the OSPF neighbors

Correct Answer: B

Section:

Explanation:

The exhibit shows the output of the command `show ospf neighbor`, which displays information about the OSPF neighbors on a router.

The output shows that the OSPF neighbor with the address 172.26.1.1 and the interface ge-0/0/3.0 is in the Exstart state.

The Exstart state is the fourth state in the OSPF neighbor formation process, after Down, Init, and 2-Way states. In this state, the OSPF neighbors establish a master-slave relationship and exchange database description (DBD) packets, which contain summaries of their link-state databases.

The most common reason for OSPF neighbors to be stuck in the Exstart state is an MTU mismatch between the interfaces. MTU stands for maximum transmission unit, which is the largest size of a packet that can be transmitted on a network segment. If the MTU values of two OSPF neighbors are different, they may not be able to exchange DBD packets successfully, as some packets may be dropped or fragmented due to their size exceeding the MTU limit.

To solve this problem, you need to ensure that the MTU values of both OSPF neighbors are the same or compatible. You can use the command `show interface` to display the MTU value of an interface. You can also

use the command ping with the do-not-fragment option to test the MTU size between two routers. You can change the MTU value of an interface by using the command set interfaces interface-name mtu mtu-value in configuration mode.

QUESTION 26

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

Correct Answer: C, D, E

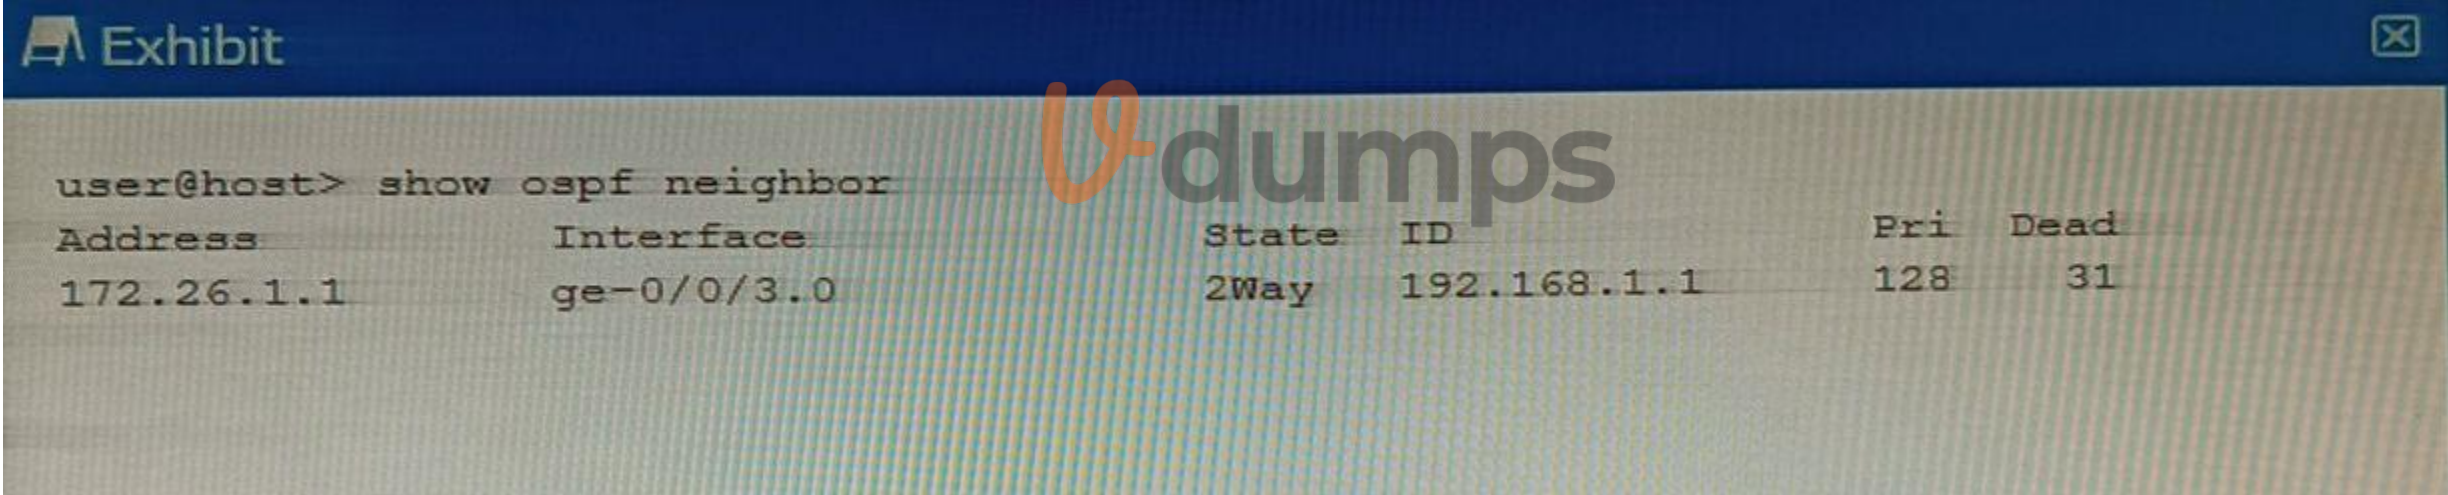
Section:

Explanation:

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

QUESTION 27

Refer to the 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.

Correct Answer: B

Section:

Explanation:

In OSPF, the state of the neighbor relationship is determined by the exchange of OSPF packets between routers. The state "2Way" as shown in the exhibit indicates that bi-directional communication has been established between the two OSPF routers. 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 network. These neighbors are often referred to as 'DRothers'. Therefore, option B is correct.

QUESTION 28

You are concerned about spoofed MAC addresses on your LAN.

Which two Layer 2 security features should you enable to minimize this concern? (Choose two.)

- A. dynamic ARP inspection
- B. IP source guard
- C. DHCP snooping
- D. static ARP

Correct Answer: A, C

Section:

Explanation:

A is correct because dynamic ARP inspection (DAI) is a Layer 2 security feature that prevents ARP spoofing attacks. ARP spoofing is a technique that allows an attacker to send fake ARP messages to associate a spoofed MAC address with a legitimate IP address. This can result in traffic redirection, man-in-the-middle attacks, or denial-of-service attacks. DAI validates ARP packets by checking the source MAC address and IP address against a trusted database, which is usually built by DHCP snooping¹. DAI discards any ARP packets that do not match the database or have invalid formats¹.

C is correct because DHCP snooping is a Layer 2 security feature that prevents DHCP spoofing attacks. DHCP spoofing is a technique that allows an attacker to act as a rogue DHCP server and offer fake IP addresses and other network parameters to unsuspecting clients. This can result in traffic redirection, man-in-the-middle attacks, or denial-of-service attacks. DHCP snooping filters DHCP messages by classifying switch ports as trusted or untrusted. Trusted ports are allowed to send and receive any DHCP messages, while untrusted ports are allowed to send only DHCP requests and receive only valid DHCP replies from trusted ports². DHCP snooping also builds a database of MAC addresses, IP addresses, lease times, and binding types for each client².

QUESTION 29

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

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



Correct Answer: B, C, D

Section:

Explanation:

In Rapid Spanning Tree Protocol (RSTP), there are several port roles that determine the behavior of the port in the spanning tree¹²³. The roles include root, designated, alternate, backup, and disabled¹²³.

The discarding state is associated with the backup, alternate, and disabled roles¹²³. In a stable topology with consistent port roles throughout the network, RSTP ensures that every root port and designated port immediately transition to the forwarding state while all alternate and backup ports are always in the discarding state². Disabled ports are also in the discarding state³.

Therefore, options B, C, and D are correct.

QUESTION 30

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

Correct Answer: B

Section:

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 LAN. The DR election process is based on the following criteria, in order of precedence:

The router with the highest OSPF priority becomes the DR. The default priority is 1, and a priority of 0 means the router will not participate in the election.

If there is a tie in priority, the router with the highest router ID becomes the DR. The router ID is a 32-bit number that uniquely identifies a router in an OSPF domain. It can be manually configured or automatically derived from the highest IP address of a loopback interface or a physical interface.

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

In your scenario, two routers share the same highest priority and start time. This means that they have equal chances of becoming the DR based on the first and third criteria. Therefore, the second criterion will be used to break the tie, which is the router ID. The router with the highest router ID will become the DR, and the other router will become the backup designated router (BDR), which is ready to take over the role of DR if it fails.

QUESTION 31

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.

Correct Answer: B, D

Section:

Explanation:

Redundant Trunk Groups (RTGs) on EX Series switches provide a simple solution for network recovery when a trunk port on a switch goes down. They are configured on the access switch and contain two links: a primary or active link, and a secondary link. 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 convergence.

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 switches. This implies that RTGs must be connected to the same aggregation switch.

QUESTION 32

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

Correct Answer: C

Section:

Explanation:

The correct answer to your question is C. Option C. Here is why:

Option C shows the configuration of the chassis statement, which defines the properties of the router chassis, such as the number of aggregated Ethernet interfaces, the number of FPCs, and the number of PICs¹.

To enable aggregated Ethernet interfaces on a router, you need to specify the aggregated-devices statement under the chassis statement and set the ethernet parameter to the desired number of interfaces². For example, to enable two aggregated Ethernet interfaces, you can use the following configuration:

```
chassis { aggregated-devices { ethernet { device-count 2; } } }
```

Option C shows this configuration with the device-count set to 2, which will enable two aggregated Ethernet interfaces on the router. The other options do not show this configuration and will not enable any aggregated Ethernet interfaces on the router.

Therefore, option C is the correct answer to your question.

QUESTION 33

Which two statements about BGP facilitate the prevention of routing loops between two autonomous systems? (Choose two.)

- A. EBGP routers will append their AS number when advertising routes to their neighbors.
- B. EBGP routers will only accept routes that contain their own AS number in the AS_PATH.
- C. EBGP routers will drop routes that contain their own AS number in the AS_PATH
- D. EBGP routers will prepend their AS number when advertising routes to their neighbors

Correct Answer: A, C

Section:

Explanation:

BGP (Border Gateway Protocol) is a protocol designed to exchange routing and reachability information among autonomous systems (AS) on the internet¹.

Option A is correct. When an EBGP router advertises routes to its neighbors, it appends its AS number to the AS_PATH attribute¹. This is a key mechanism in BGP to prevent routing loops¹.

Option C is correct. BGP has a built-in loop prevention mechanism whereby if a BGP router detects its own AS in the AS_PATH attribute, it will drop the prefix and will not continue to advertise it². This helps to prevent routing loops².

Option B is incorrect. EBGP routers do not accept routes that contain their own AS number in the AS_PATH². Instead, they drop such routes as part of the loop prevention mechanism².

Option D is incorrect. While it's true that EBGP routers append their AS number when advertising routes, they do not prepend their AS number¹. The term "prepend" in BGP usually refers to a technique used to influence path selection by artificially lengthening the AS_PATH³.

QUESTION 34

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.



Correct Answer: B

Section:

Explanation:

An ISO NET address is a type of network address used by the IS-IS routing protocol. It identifies a point of connection to the network, such as a router interface, and is also called a Network Service Access Point (NSAP)¹.

An ISO NET address consists of three parts: an area ID, a system ID, and a selector². The area ID identifies the IS-IS area to which the device belongs. The system ID uniquely identifies the device within the area. The selector identifies a specific service or function on the device, such as routing or management².

An ISO NET address must be unique for each device in the network, because it is used by IS-IS to establish adjacencies, exchange routing information, and compute shortest paths². If two devices have the same ISO NET address, they will not be able to communicate with each other or with other devices in the network. Therefore, it is important to assign different ISO NET addresses to each device in the network.

QUESTION 35

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

Correct Answer: D

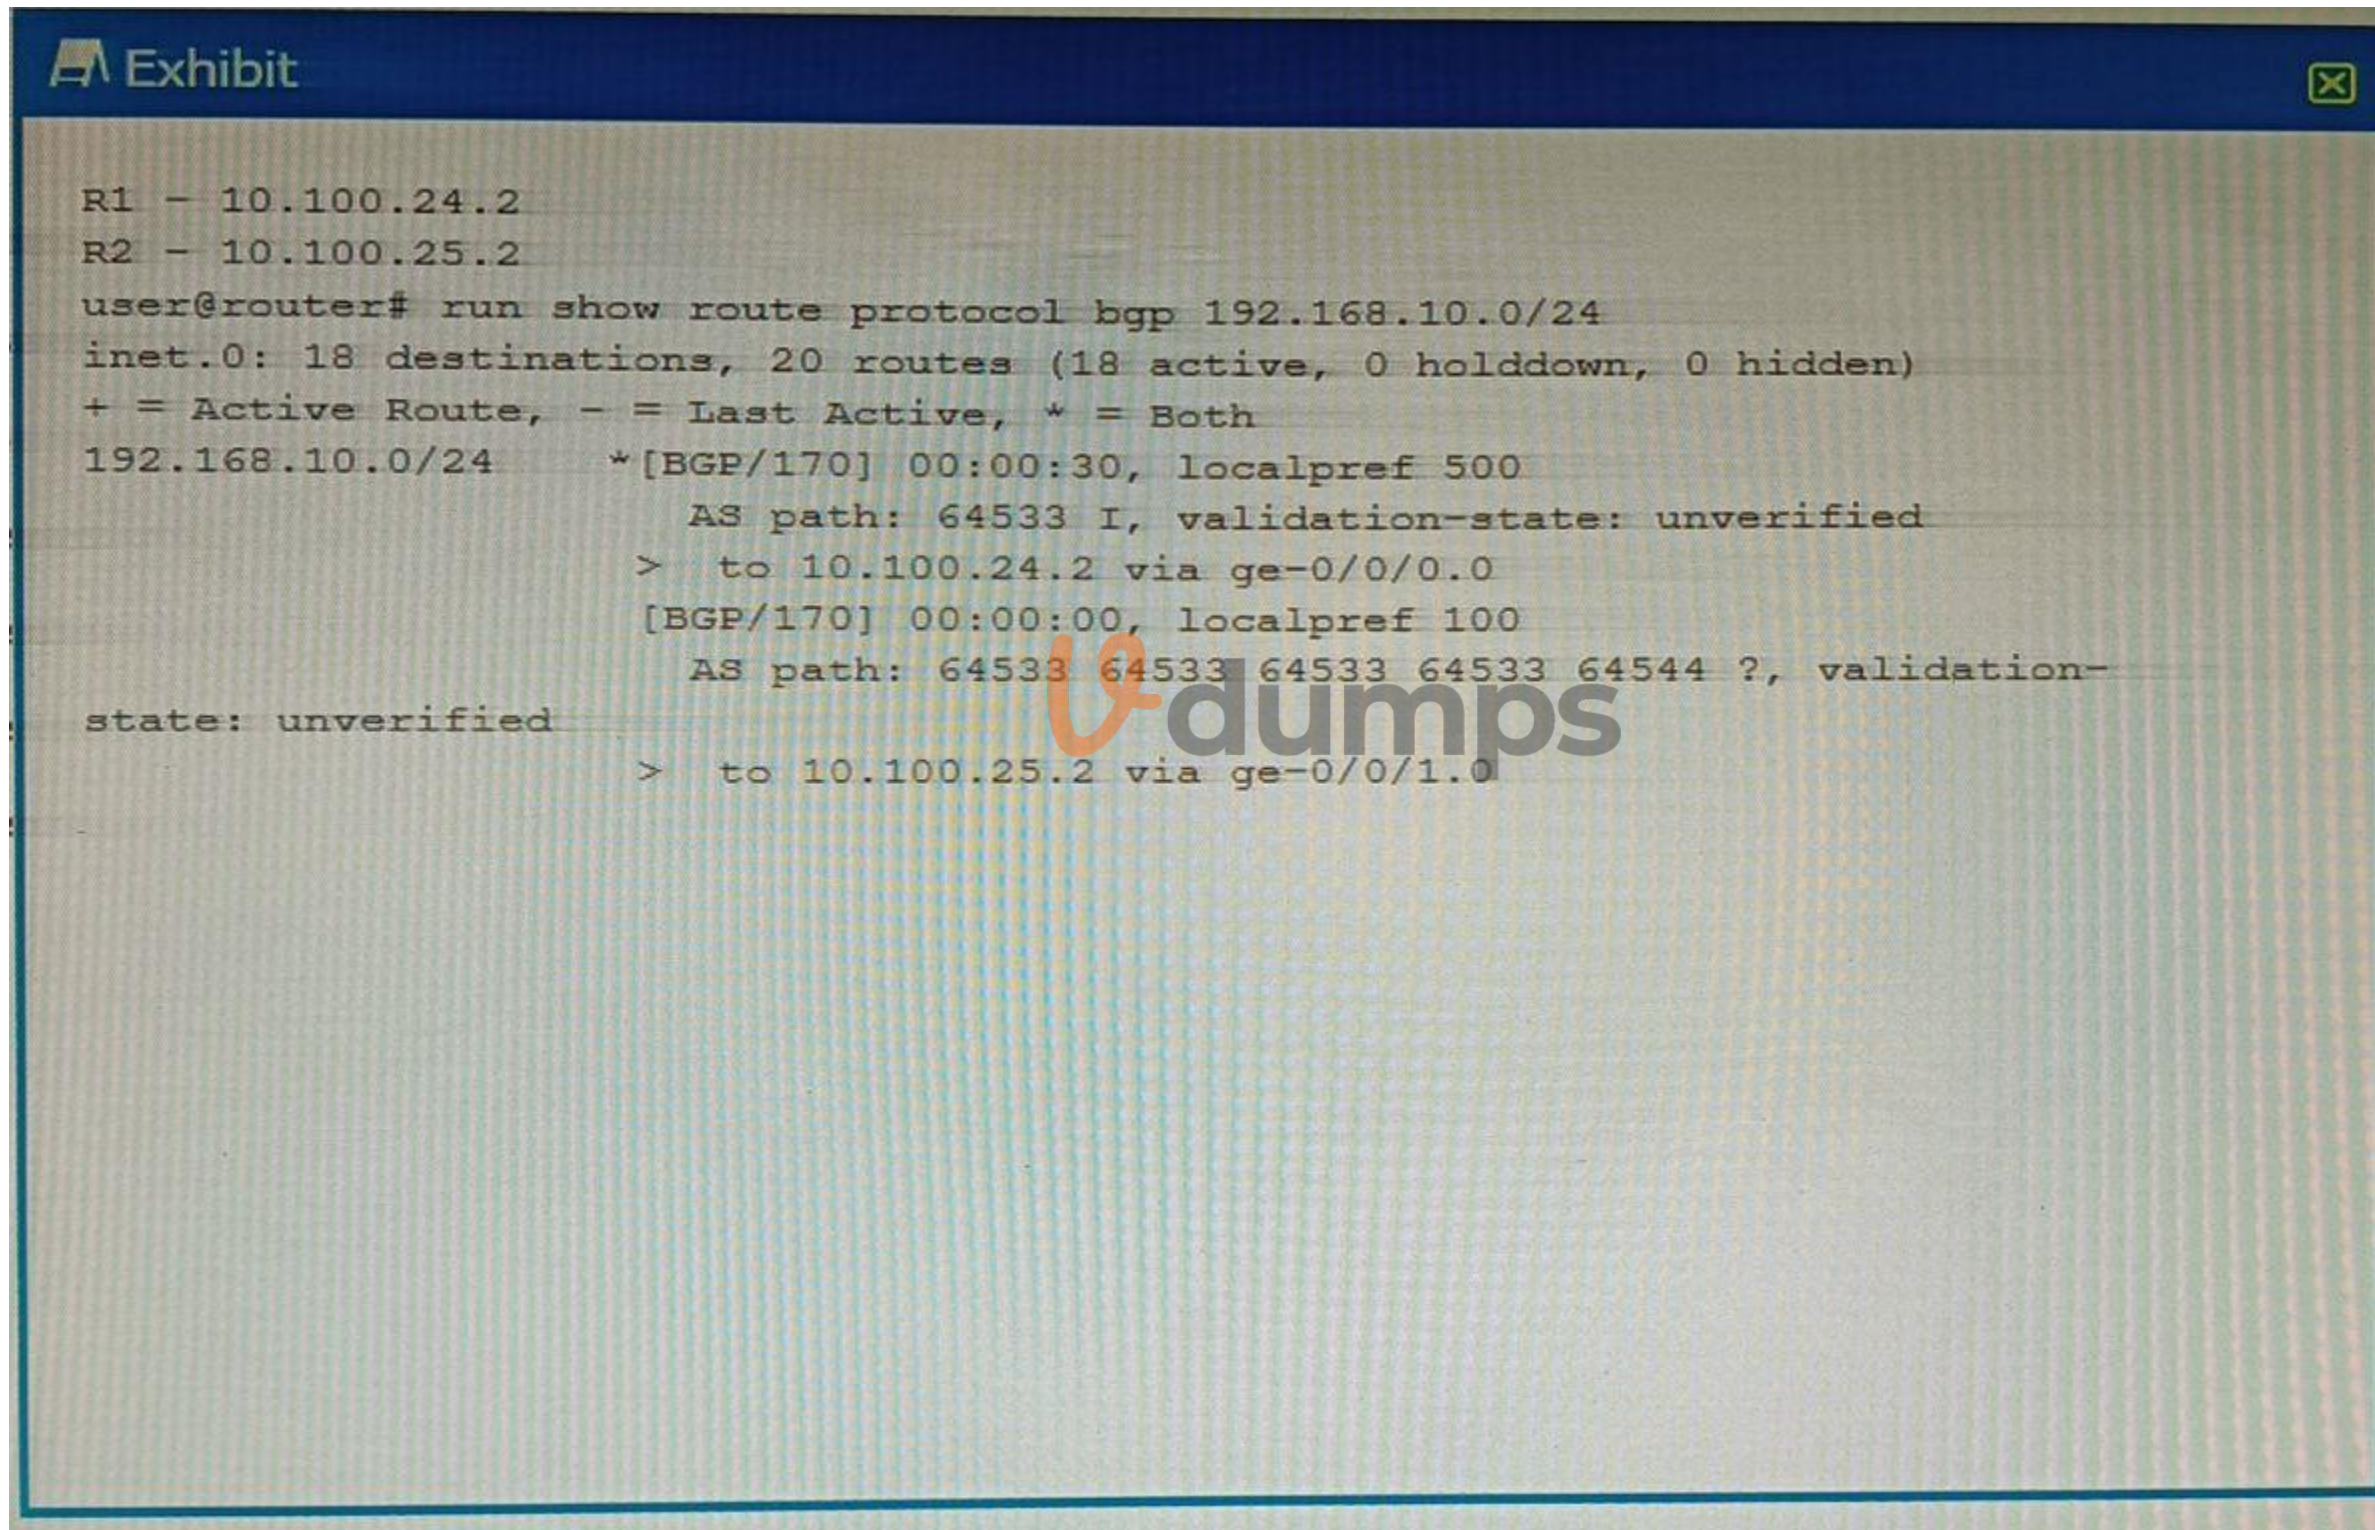
Section:

Explanation:

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

QUESTION 36

Exhibit



```
R1 - 10.100.24.2
R2 - 10.100.25.2
user@router# run show route protocol bgp 192.168.10.0/24
inet.0: 18 destinations, 20 routes (18 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
192.168.10.0/24    *[BGP/170] 00:00:30, localpref 500
                  AS path: 64533 I, validation-state: unverified
                  > to 10.100.24.2 via ge-0/0/0.0
                  [BGP/170] 00:00:00, localpref 100
                  AS path: 64533 64533 64533 64533 64544 ?, validation-
state: unverified
                  > to 10.100.25.2 via ge-0/0/1.0
```

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.

Correct Answer: C

Section:

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.

