Juniper

Exam Questions JN0-351

Enterprise Routing and Switching - Specialist (JNCIS-ENT)

NEW QUESTION 1
Exhibit.

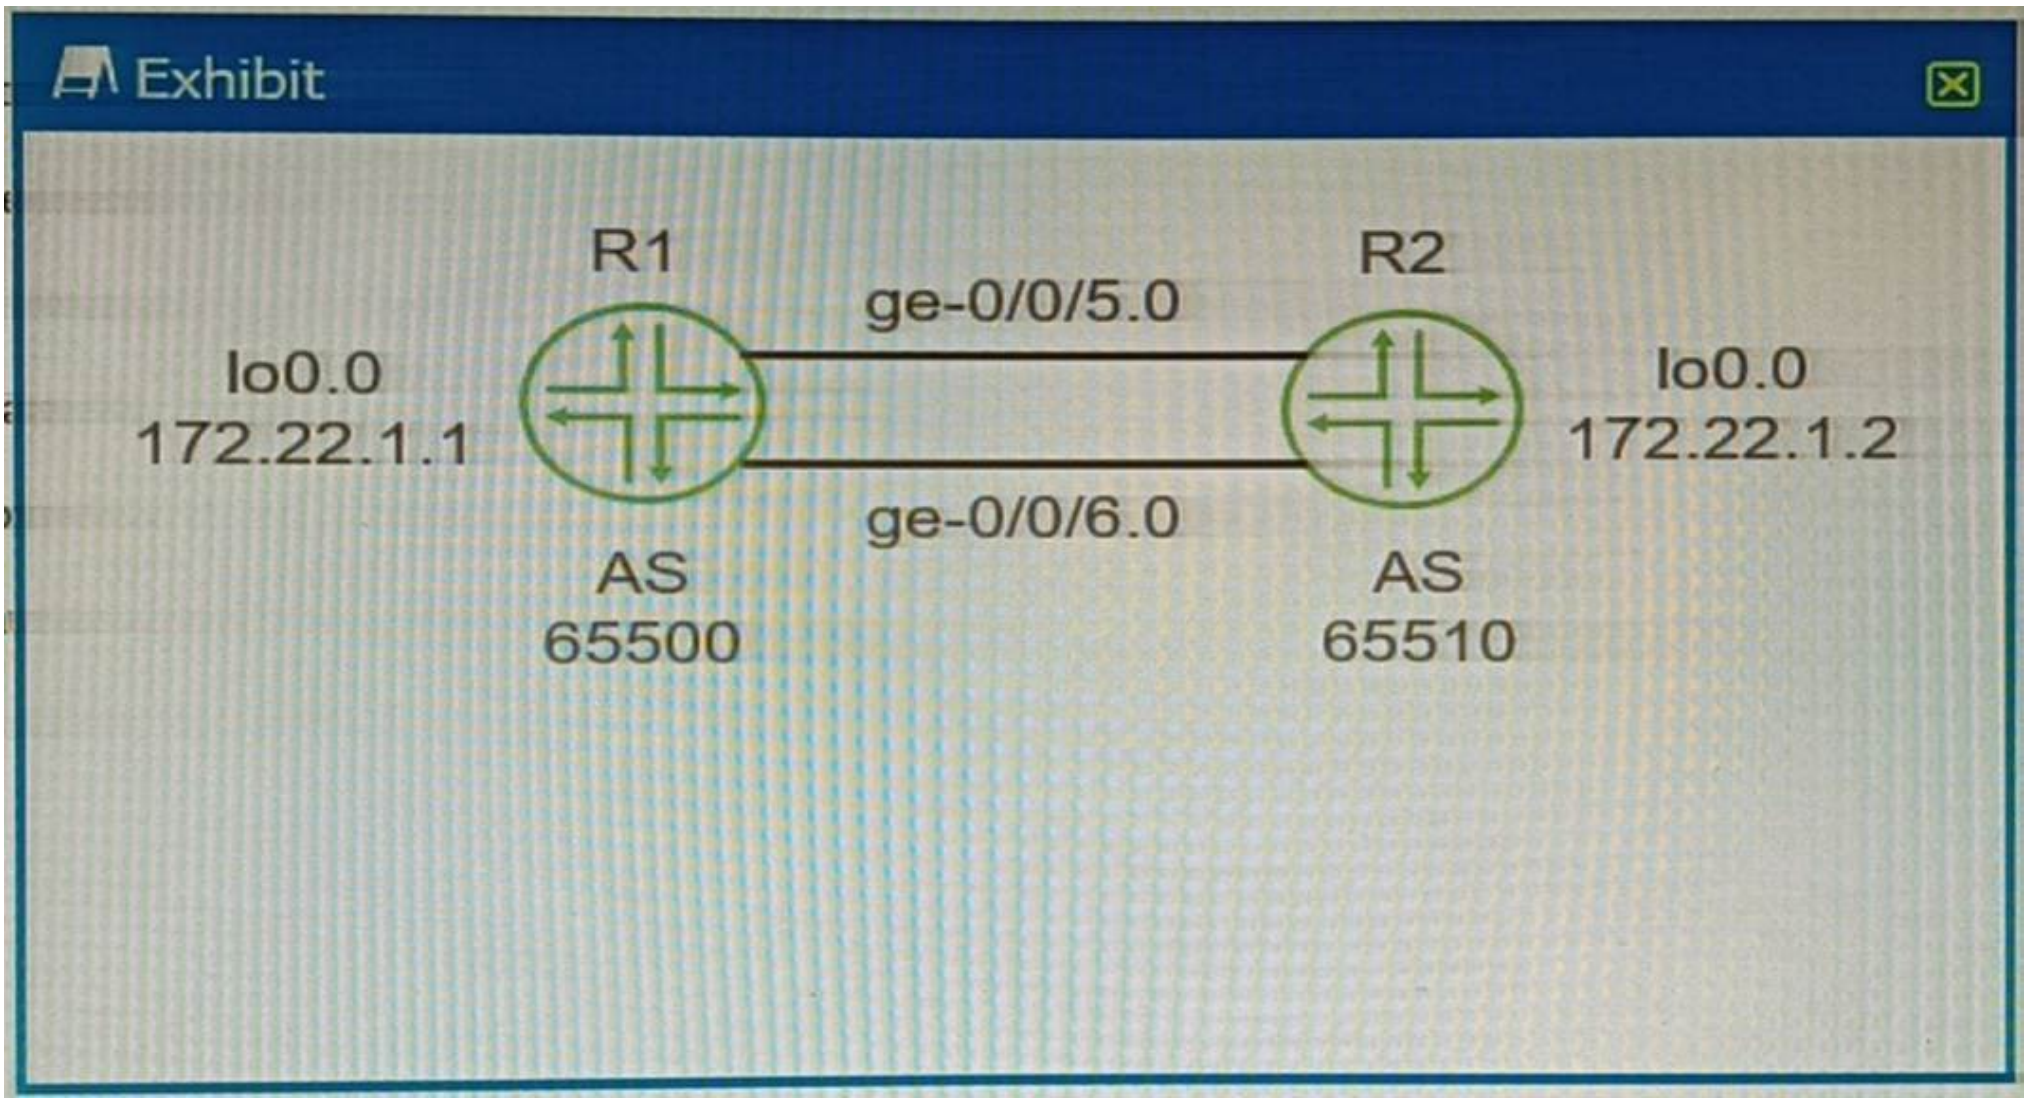

You want to enable redundancy for the EBGP peering between the two routers shown in the exhibit. Which three actions will you perform in this scenario? (Choose three.)

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 2
What is the default keepalive time for BGP?

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

Answer: B

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

NEW QUESTION 3
What is the maximum allowable MTU size for a default GRE tunnel without IPv4 traffic fragmentation?

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

Answer: D

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

NEW QUESTION 4
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 5
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 6
You are an operator for a network running 1S-IS. Two routers are failing to form an adjacency. What are two reasons for this problem? (Choose two.)

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

Answer: BD

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

NEW QUESTION 7
Which two statements are correct about generated routes? (Choose two.)

A. Generated routes require a contributing route.
B. Generated routes show a next hop in the routing table.
C. Generated routes appear in the routing table as static routes
D. Generated routes cannot be redistributed into dynamic routing protocols.

Answer: AB

Explanation:
? A is correct because generated routes require a contributing route. A contributing route is a route that matches the destination prefix of the generated route and has a valid next hop1. A generated route is only installed in the routing table if there is at least one contributing route available2. This ensures that the generated route is reachable and useful. If there is no contributing route, the generated route is not added to the routing table2.
? B is correct because generated routes show a next hop in the routing table. A generated route inherits the next hop of its primary contributing route, which is the most preferred route among all the contributing routes2. The next hop of the generated route can be either an IP address or an interface name, depending on the type of the contributing route2. The next hop of the generated route can also be modified by a routing policy3.

NEW QUESTION 8
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 9
An update to your organization's network security requirements document requires management traffic to be isolated in a non-default routing-instance. You want to implement
this requirement on your Junos-based devices.
Which two commands enable this behavior? (Choose two.)

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

Answer: CD

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

NEW QUESTION 10
You are concerned about spoofed MAC addresses on your LAN.
Which two Layer 2 security features should you enable to minimize this concern? (Choose two.)

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

Answer: AC

Explanation:
? A is correct because dynamic ARP inspection (DAI) is a Layer 2 security feature that prevents ARP spoofing attacks. ARP spoofing is a technique that allows an attacker to send fake ARP messages to associate a spoofed MAC address with a legitimate IP address. This can result in traffic redirection, man-in-the-middle attacks, or denial-of-service attacks. DAI validates ARP packets by checking the source MAC address and IP address against a trusted database, which is usually built by DHCP snooping1. DAI discards any ARP packets that do not match the database or have invalid formats1.
? C is correct because DHCP snooping is a Layer 2 security feature that prevents DHCP spoofing attacks. DHCP spoofing is a technique that allows an attacker to act as a rogue DHCP server and offer fake IP addresses and other network parameters to unsuspecting clients. This can result in traffic redirection, man-in- the-middle attacks, or denial-of-service attacks. DHCP snooping filters DHCP messages by classifying switch ports as trusted or untrusted. Trusted ports are allowed to send and receive any DHCP messages, while untrusted ports are allowed to send only DHCP requests and receive only valid DHCP replies from trusted ports2. DHCP snooping also builds a database of MAC addresses, IP addresses, lease times, and binding types for each client2.

NEW QUESTION 10
What is a purpose of using a spanning tree protocol?

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

Answer: B

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

NEW QUESTION 11
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 16
Exhibit

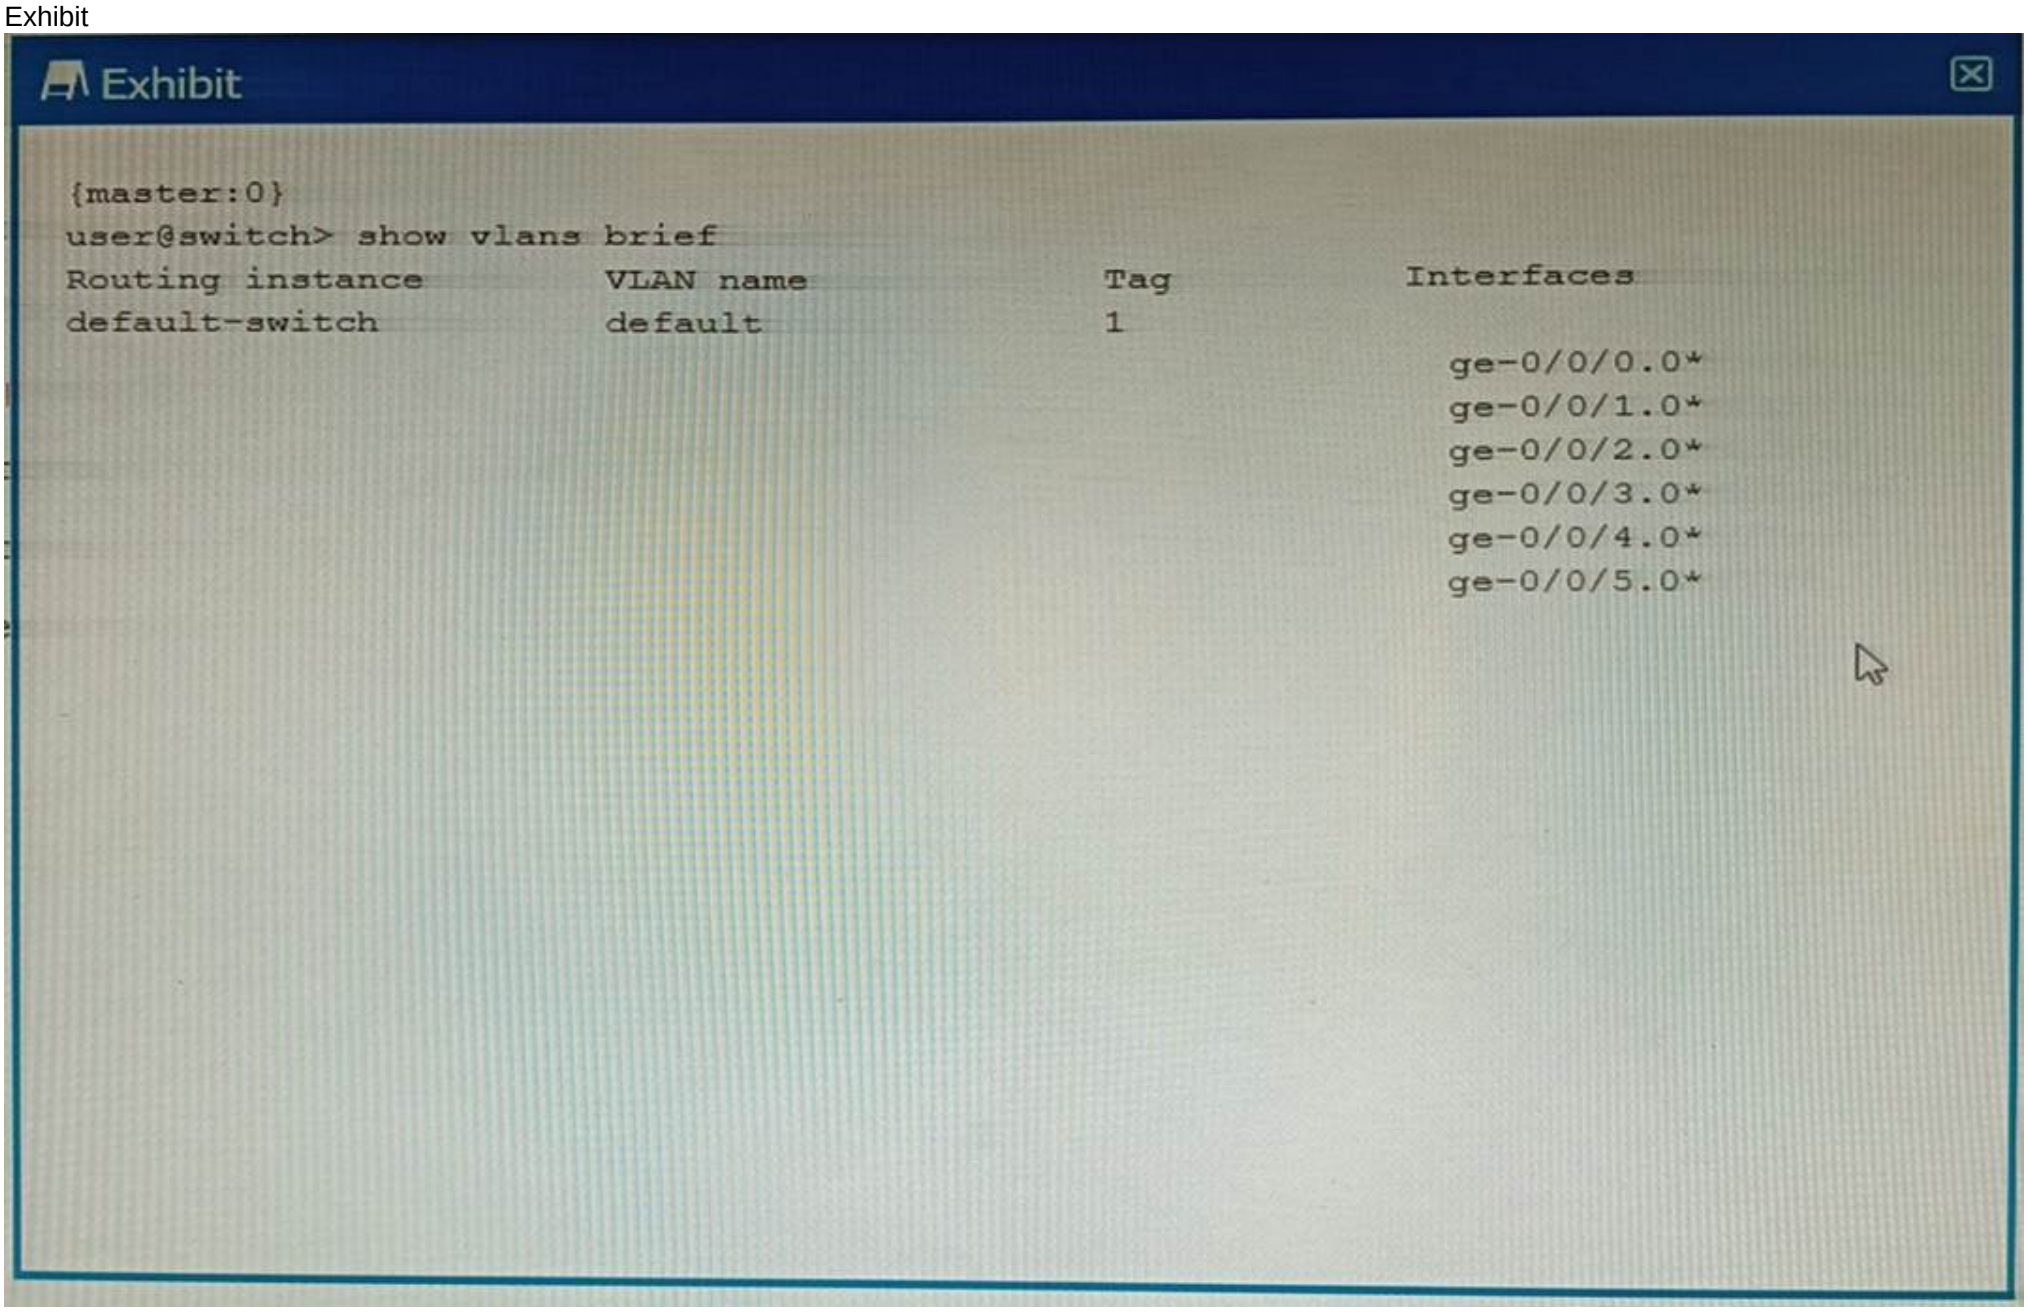

What does the * indicate in the output shown in the exhibit?

A. The switch ports have a router attached.
B. The interface is down.
C. The interface is active.

D. All interfaces have elected a root bridge.

Answer: C

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

NEW QUESTION 20
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 23
Which three protocols support BFD? (Choose three.)

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

Answer: BCD

Explanation:
BFD is a protocol that can be used to quickly detect failures in the forwarding path between two adjacent routers or switches. BFD can be integrated with various routing protocols and link aggregation protocols to provide faster convergence and fault recovery. According to the Juniper Networks documentation, the following protocols support BFD on Junos OS devices1:
? BGP: BFD can be used to monitor the connectivity between BGP peers and trigger
a session reset if a failure is detected. BFD can be configured for both internal and external BGP sessions, as well as for IPv4 and IPv6 address families2.
? OSPF: BFD can be used to monitor the connectivity between OSPF neighbors and
trigger a state change if a failure is detected. BFD can be configured for both OSPFv2 and OSPFv3 protocols, as well as for point-to-point and broadcast network types3.
? LACP: BFD can be used to monitor the connectivity between LACP members and
trigger a link state change if a failure is detected. BFD can be configured for both active and passive LACP modes, as well as for static and dynamic LAGs4.
Other protocols that support BFD on Junos OS devices are:
? IS-IS: BFD can be used to monitor the connectivity between IS-IS neighbors and trigger a state change if a failure is detected. BFD can be configured for both level 1 and level 2 IS-IS adjacencies, as well as for point-to-point and broadcast network types.
? RIP: BFD can be used to monitor the connectivity between RIP neighbors and trigger a route update if a failure is detected. BFD can be configured for both RIP version 1 and version 2 protocols, as well as for IPv4 and IPv6 address families.
? VRRP: BFD can be used to monitor the connectivity between VRRP routers and trigger a priority change if a failure is detected. BFD can be configured for both VRRP version 2 and version 3 protocols, as well as for IPv4 and IPv6 address families.
The protocols that do not support BFD on Junos OS devices are:
? RSTP: RSTP is a spanning tree protocol that provides loop prevention and rapid convergence in layer 2 networks. RSTP does not use BFD to detect link failures, but relies on its own hello mechanism that sends BPDU packets every 2 seconds by default.
? FTP: FTP is an application layer protocol that is used to transfer files between hosts over a TCP connection. FTP does not use BFD to detect connection failures, but relies on TCP??s own retransmission and timeout mechanisms.
References:
1: [Configuring Bidirectional Forwarding Detection] 2: [Configuring Bidirectional Forwarding Detection for BGP] 3: [Configuring Bidirectional Forwarding Detection for OSPF] 4: [Configuring Bidirectional Forwarding Detection for Link Aggregation Control Protocol] : [Configuring Bidirectional Forwarding Detection for IS-IS] : [Configuring Bidirectional Forwarding Detection for RIP] : [Configuring Bidirectional Forwarding Detection for VRRP] : [Understanding Rapid Spanning Tree Protocol] : [Understanding FTP]

NEW QUESTION 27
Exhibit

Exhibit

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

Answer: C

Explanation:
In BGP, the path selection process is based on a set of attributes1. The process starts by preferring the path with the highest weight, then the highest local preference, then the locally originated routes, and so on1. If all these attributes are the same, then it prefers the path with the shortest AS path1.
Referring to the exhibit, all four ISPs have the same weight, local preference, and origin1. However, ISP 4 has the shortest AS path1. Therefore, ISP 4 will be selected as the active path. So, option C is correct.

NEW QUESTION 32
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 35
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 36
.....

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