



Vendor: Juniper

Exam Code: JN0-351

Exam Name: Enterprise Routing and Switching, Specialist
(JNCIS-ENT)

Version: DEMO

QUESTION 1

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 port.

C is correct because the interface is configured as a trunk port. By default, all trunk ports on the switch are trusted for DHCP snooping. 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 port.

QUESTION 2

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



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

Answer: C

Explanation:

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

local 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 3

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

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

Answer: AC

Explanation:

A is correct because RSTP alternate ports block traffic while receiving superior BPDUs from a neighboring switch. An alternate port is a backup port for a root port, which means it receives better BPDUs from another bridge than the current root port. However, an alternate port does not forward any traffic, as it is in a discarding state. 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 port.

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 ID. However, an alternate port receives a higher cost BPDU than the root port, otherwise it would be the root port itself. Therefore, an alternate port provides an alternate higher cost path to the root bridge than the root port.

QUESTION 4

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 prevention. There are four categories of BGP attributes:

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

Well-known discretionary: Supported by all BGP implementations, and are optionally included in BGP updates.

Optional transitive: May not be supported by all implementations of BGP.

Optional non-transitive: May not be supported by all implementations of BGP. The well-known mandatory attributes must be supported by all BGP implementations and must be included in every update. These include the AS path and next hop attributes.

QUESTION 5

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

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

Answer: BCD

Explanation:

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

The discarding state is associated with the backup, alternate, and disabled roles. 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 state. Disabled ports are also in the discarding state.

QUESTION 6

Two routers share the same highest priority and start time. In this situation, what is evaluated next when determining the designated router?

- A. The router with the lowest router ID become the DR.
- B. The router with the highest router ID becomes the DR
- C. The routers perform another DR election.
- D. The router with the highest MAC address become the DR

Answer: B

Explanation:

According to the OSPF protocol, the designated router (DR) is the router that acts as the focal point for exchanging routing information on a multi-access network segment, such as a 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 7

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

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

Answer: BC

Explanation:

Redundant trunk groups are designed to provide link redundancy. If the primary link fails, the secondary link will automatically take over to ensure continued connectivity.

In redundant trunk groups, while regular data traffic is sent over the primary link, Layer 2 control traffic (such as STP, LACP, etc.) is allowed on the secondary link to maintain network stability and protocol operations.

<https://www.juniper.net/documentation/us/en/software/junos/multicast-l2/topics/topic-map/redundant-trunk-groups.html>

QUESTION 8

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

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

Answer: B

Explanation:

Device count should be more than 1.

QUESTION 9

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

Answer: AC

Explanation:

When an EBGp router advertises a route to a neighbor in a different AS, it appends its own AS number to the AS_PATH attribute. This helps other routers recognize the path the route has taken and prevents routing loops.

If an EBGp router receives a route advertisement that already contains its own AS number in the AS_PATH attribute, it will drop the route. This prevents the router from accepting a route that would create a routing loop.

QUESTION 10

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).

An ISO NET address consists of three parts: an area ID, a system ID, and a selector. 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 management.

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 paths. 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 11

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 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 12

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 hop. A generated route is only installed in the routing table if there is at least one contributing route available. 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 table.

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 routes. The next hop of the generated route can be either an IP address or an interface name, depending on the type of the contributing route. The next hop of the generated route can also be modified by a routing policy.

QUESTION 13

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 degradation. A broadcast storm can occur when there are loops in the network topology, meaning that there are multiple paths between two devices.

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 devices.

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 congestion.

QUESTION 14

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: BC

Explanation:

GRE tunnels can be created on all Junos devices. GRE is a tunneling protocol that encapsulates a wide variety of network layer protocols inside virtual point-to-point links.

IP-IP tunnels can also be created on all Junos devices. IP-IP tunneling is used to encapsulate IP packets within IP packets, providing a simple way to tunnel IP traffic.

QUESTION 15

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:

Redundant trunk groups are designed to provide link redundancy. If the primary link fails, the secondary link will automatically take over to ensure continued connectivity.

In redundant trunk groups, while regular data traffic is sent over the primary link, Layer 2 control traffic (such as STP, LACP, etc.) is allowed on the secondary link to maintain network stability and protocol operations.

<https://www.juniper.net/documentation/us/en/software/junos/multicast-l2/topics/topic-map/redundant-trunk-groups.html>

Thank You for Trying Our Product

Braindump2go Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.braindump2go.com/all-products.html>



Microsoft



ORACLE



JUNIPER
NETWORKS



EMC²
where information lives

10% Discount Coupon Code: ASTR14