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

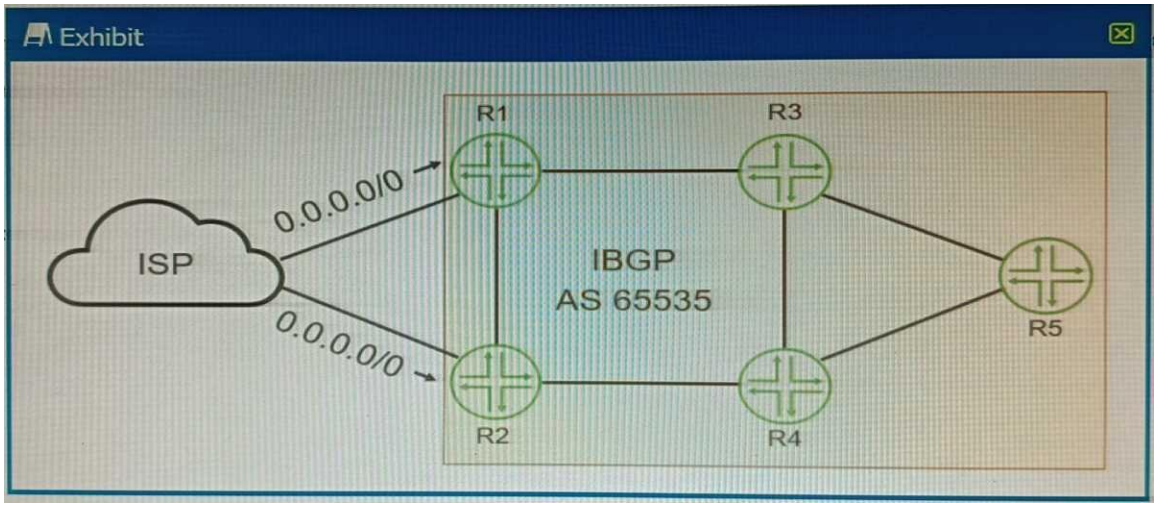

A. MED
B. next-hop
C. local preference
D. origin

Answer: C
Explanation:
The BGP attribute that you would use to forward all Internet traffic through the R1 device is the local 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: BD

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

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

Answer: C
Explanation:
The correct answer to your question is C. Option C. Here is why:
Option C shows the configuration of the chassis statement, which defines the properties of the router chassis, such as the number of aggregated Ethernet interfaces, the number of FPCs, and the number of PICs.
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 interfaces. 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 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:
BGP (Border Gateway Protocol) is a protocol designed to exchange routing and reachability information among autonomous systems (AS) on the internet. Option A is correct. When an EBGP router advertises routes to its neighbors, it appends its AS number to the AS_PATH attribute. This is a key mechanism in BGP to prevent routing loops. 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 it. This helps to prevent routing loops.
Option B is incorrect. EBGP routers do not accept routes that contain their own AS number in the AS_PATH. Instead, they drop such routes as part of the loop prevention mechanism. Option D is incorrect. While it's true that EBGP routers append their AS number when advertising routes, they do not prepend their AS number. The term "prepend" in BGP usually refers to a technique used to influence path selection by artificially lengthening the AS_PATH.

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.

Thank You for Trying Our Product

Lead2pass 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.lead2pass.com/all-products.html

10% Discount Coupon Code: **ASTR14**