



**Free Questions for JN0-351 by [braindumpscollection](#)**

**Shared by [Ryan](#) on 04-10-2023**

**For More Free Questions and Preparation Resources**

**[Check the Links on Last Page](#)**

# Question 1

---

**Question Type:** MultipleChoice

---

You are receiving multiple BGP routes from an upstream neighbor and only want to advertise a single summarized prefix to your internal OSPF neighbors. This route should only be advertised when you are receiving these BGP routes from this neighbor.

In this scenario, which type of route should you create?

## Options:

---

- A- aggregate route
- B- static route using the resolve feature
- C- generate route
- D- static route using qualified next hops

## Answer:

---

A

## Explanation:

---

In this scenario, you should create an aggregate route. Aggregate routes are used for advertising summarized network prefixes. They help minimize the number of routing tables in an IP network by consolidating selected multiple routes into a single route advertisement. This approach is in contrast to non-aggregation routing, in which every routing table contains a unique entry for each route.

Therefore, option A is correct. Options B, C, and D are not correct because:

**Static route using the resolve feature:** This type of route uses the resolve feature to install a static route in the routing table only if a specific condition is met. However, it does not provide the capability to summarize multiple routes into a single prefix.

**Generate route:** This type of route generates a route that is always present in the routing table and can be used to summarize routes. However, it does not have the capability to only advertise the route when specific BGP routes are being received from a neighbor.

**Static route using qualified next hops:** This type of route allows for the specification of multiple next-hop addresses for a static route. However, it does not provide the capability to summarize multiple routes into a single prefix.

## Question 2


---

**Question Type:** MultipleChoice

---

Exhibit

# Exhibit



```
[edit protocols]
user@R1# show
isis {
  interface ge-0/0/1 {
    level 1 disable;
  }
  interface lo0.0 {
    level 2 disable;
  }
}
```

```
[edit protocols]
user@R2# show
isis {
  interface ge-
    level 2 c
  }
  interface ge-
    level 2 c
  }
  interface lo0
    level 1 c
  }
}
```

```
[edit protocols]
user@R3# show
isis {
  interface ge-0/
    level 1 dis
  }
  interface lo0.0 {
    level 2 dis
  }
}
```

Referring to the exhibit, which two configuration changes must you apply for packets to reach from R1 to R3 using IS-IS? (Choose two.)

### Options:

---

- A- On R1, enable Level 1 on the ge-0/0/1 interface.
- B- On R3 disable Level 2 on the ge-0/0/4 interface.
- C- On R1, disable Level 2 on the ge-0/0/1 interface.
- D- On R3 enable Level 1 on the ge-0/0/4 interface

### Answer:

---

A, D

### Explanation:

---

A) On R1, enable Level 1 on the ge-0/0/1 interface. In IS-IS, both levels (Level 1 and Level 2) are enabled by default when you enable IS-IS on an interface<sup>1</sup>. Level 1 systems route within an area<sup>2</sup>. If the destination is outside an area, Level 1 systems route toward a Level 2 system<sup>2</sup>. Therefore, enabling Level 1 on the ge-0/0/1 interface on R1 would allow packets to reach from R1 to R3.

D) On R3 enable Level 1 on the ge-0/0/4 interface. Similarly, enabling Level 1 on the ge-0/0/4 interface on R3 would allow packets to reach from R1 to R3.

These explanations are based on the IS-IS configuration documents and learning resources available at Juniper Networks<sup>1</sup> and Cisco<sup>34</sup>.

## Question 3

---

**Question Type:** MultipleChoice

---

Which statement about aggregate routes is correct?

### Options:

---

- A-** Aggregate routes can only be used for static routing but not for dynamic routing protocols.
- B-** Aggregate routes are automatically generated for all of the subnets in a routing table.
- C-** Aggregate routes are always preferred over more specific routes, even when the specific routes have a better path.
- D-** Aggregate routes are used for advertising summarized network prefixes.

### Answer:

---

D

### Explanation:

---

Aggregate routes are used for advertising summarized network prefixes<sup>12</sup>. They help minimize the number of routing tables in an IP network by consolidating selected multiple routes into a single route advertisement<sup>1</sup>. This approach is in contrast to non-aggregation routing, in which every routing table contains a unique entry for each route<sup>1</sup>.

Therefore, option D is correct. Options A, B, and C are not correct because:

Aggregate routes can be used with both static routing and dynamic routing protocols<sup>1</sup>.

Aggregate routes are not automatically generated for all of the subnets in a routing table. They need to be manually configured<sup>1</sup>.

Aggregate routes are not always preferred over more specific routes. The route selection process in Junos OS considers several factors, including route preference and metric, before determining the active route<sup>1</sup>.

## Question 4

---

**Question Type:** MultipleChoice

---

Exhibit

## Exhibit

```
user@switch> show spanning-tree bridge
STP bridge parameters
Context ID : 0
Enabled protocol : RSTP
  Root ID : 4096.00:19:e2:55:36:1e
  Root cost : 40000
  Root port : ge-0/0/13.0
  Hello time : 2 seconds
  Maximum age : 20 seconds
  Forward delay : 15 seconds
  Message age : 2
  Number of topology changes : 2
  Time since last topology change : 72 seconds
Local parameters
  Bridge ID : 32768.00:19:e2:55:1d:30
  Extended system ID : 0
  Internal instance ID : 0
```



Referring to the exhibit, which statement is correct?

**Options:**

---

- A- The local device is using a bridge priority of 4k.
- B- The root bridge is using a bridge priority of 4k.
- C- The root bridge has not been elected for this RSTP topology.
- D- The local device is the root bridge for this RSTP topology.

**Answer:**

---

D

**Explanation:**

---

In a Rapid Spanning Tree Protocol (RSTP) topology, the root bridge is determined by the switch with the lowest bridge priority value<sup>12</sup>. If all switches have the same priority, then the root bridge is assigned to the switch whose MAC address's hex value is the lowest<sup>2</sup>. The default bridge priority value is 3276832. However, without the actual exhibit, it's difficult to definitively determine which device is the root bridge. But based on the options provided, if we assume that the local device has a lower bridge priority or a lower MAC address than other devices in the network, then it could be considered as the root bridge for this RSTP topology<sup>45</sup>.

## Question 5

---

### Question Type: MultipleChoice

---

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.)

#### Options:

---

- 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:

---

C, D

#### Explanation:

---

To isolate management traffic in a non-default routing-instance on Junos-based devices, you can use these system management-instance and set routing-instances mgmt\_junos commands<sup>12</sup>.

`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) instance<sup>1</sup>. 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 traffic<sup>1</sup>.

`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_junos`<sup>1</sup>.

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 traffic<sup>1</sup>.

## Question 6

---

**Question Type:** MultipleChoice

---

You implemented the MAC address limit feature with the shutdown action on all interfaces on your switch.

In this scenario, which statement is correct when a violation occurs?

## Options:

---

- A-** By default, you must manually clear the violation for the interface to send and receive traffic again.
- B-** By default, the violation will automatically be cleared after 300 seconds and the interface will resume sending and receiving traffic for all learned devices.
- C-** By default, devices that are learned before the violation occurs are still allowed to send and receive traffic through the specific interface.
- D-** By default, the interface will continue to send and receive traffic for all connected devices after a violation has occurred.

## Answer:

---

A

## Explanation:

---

When the MAC address limit feature with the shutdown action is implemented on a switch, if a violation occurs, the interface is disabled and a system log entry is generated<sup>1</sup>. If the switch has been configured with the port-error-disable statement, the disabled interface recovers automatically upon expiration of the specified disable timeout<sup>1</sup>. However, if the switch has not been configured for auto-recovery from port error disabled conditions, you must manually clear the violation by running the clear ethernet-switching port-error command for the interface to send and receive traffic again<sup>1</sup>. This explanation is based on the Enterprise Routing and Switching Specialist (JNCIS-ENT) documents and learning resources available at Juniper Networks<sup>1</sup>.

## Question 7

---

**Question Type:** MultipleChoice

---

You have DHCP snooping enabled but no entries are automatically created in the snooping database for an interface on your EX Series switch. What are two reasons for the problem? (Choose two.)

### Options:

---

- A- The device that is connected to the interface has performed a DHCPRELEASE.
- B- MAC limiting is enabled on the interface.
- C- The device that is connected to the interface has a static IP address.
- D- Dynamic ARP inspection is enabled on the interface.

### Answer:

---

B, C

### Explanation:

---

The DHCP snooping feature in Juniper Networks' EX Series switches works by building a binding database that maps the IP address, MAC address, lease time, binding type, VLAN number, and interface information<sup>1</sup>. This database is used to filter and validate DHCP messages from untrusted sources<sup>1</sup>.

However, there are certain conditions that could prevent entries from being automatically created in the snooping database for an interface:

**MAC limiting:** If MAC limiting is enabled on the interface, it could potentially interfere with the operation of DHCP snooping. MAC limiting restricts the number of MAC addresses that can be learned on a physical interface to prevent MAC flooding attacks<sup>1</sup>. This could inadvertently limit the number of DHCP clients that can be learned on an interface, thus preventing new entries from being added to the DHCP snooping database.

**Static IP address:** If the device connected to the interface is configured with a static IP address, it will not go through the DHCP process and therefore will not have an entry in the DHCP snooping database<sup>1</sup>. The DHCP snooping feature relies on monitoring DHCP messages to build its database<sup>1</sup>, so devices with static IP addresses that do not send DHCP messages will not have their information added.

Therefore, options B and C are correct. Options A and D are not correct because performing a DHCPRELEASE would simply remove an existing entry from the database<sup>1</sup>, and Dynamic ARP inspection (DAI) uses the information stored in the DHCP snooping binding database but does not prevent entries from being created<sup>1</sup>.

## Question 8

---

**Question Type:** MultipleChoice

---

You are an operator for a network running IS-IS. Two routers are failing to form an adjacency. What are two reasons for this problem? (Choose two.)

**Options:**

---

- 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:**

---

B, D

**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 adjacency.

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 adjacency<sup>1</sup>.

These explanations are based on the Enterprise Routing and Switching Specialist (JNCIS-ENT) documents and learning resources available at Juniper Networks<sup>23</sup>.



**To Get Premium Files for JN0-351 Visit**

**<https://www.p2pexams.com/products/jn0-351>**

**For More Free Questions Visit**

**<https://www.p2pexams.com/juniper/pdf/jn0-351>**

