



**Free Questions for JN0-351 by vceexamstest**

**Shared by Shepherd on 24-05-2024**

**For More Free Questions and Preparation Resources**

**Check the Links on Last Page**

# Question 1

---

## Question Type: MultipleChoice

---

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.)

### Options:

---

- 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:

---

A, C

### 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 2

---

**Question Type: MultipleChoice**

---

Exhibit

## 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.
2, 00:26:88:02:74:87	user	0		ucst	1305	3	ge-0/0/7.
2, 00:26:88:02:74:88	user	0		ucst	1306	3	ge-0/0/8.

Which command displays the output shown in the exhibit?

**Options:**

---

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

**Answer:**

---

B

**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 on<sup>2</sup>. 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 type<sup>1</sup>. 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 name<sup>3</sup>. It does not show the VLAN name or the forwarding state.

## Question 3

---

**Question Type:** MultipleChoice

---

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?

**Options:**

---

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

**Answer:**

---

D

**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 VLAN.

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 VLANs. This allows the switch to differentiate between voice and data traffic, ensuring that voice traffic can be treated with a higher priority. Therefore, option D is correct.

## Question 4

---

**Question Type:** MultipleChoice

---

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.)

### Options:

---

- 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.

### Answer:

---

C, D

### Explanation:

---

In BGP, routing policies are used to control the flow of routing information between BGP peers<sup>1</sup>.

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)<sup>1</sup>. The RIB-In is a database that stores all the routes that are received from all peers<sup>1</sup>.



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)<sup>1</sup>. The RIB-Local is a database that stores all the routes that the local router is using<sup>1</sup>.

Therefore, options C and D are correct.

## Question 5

---

**Question Type:** MultipleChoice

---

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.)

### Options:

---

- 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.

**Answer:**

---

A, B

**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 peer<sup>1</sup>.

B is 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 peers<sup>1</sup>.

## Question 6

---

**Question Type:** MultipleChoice

---

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

### Options:

---

- 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 bytes<sup>1</sup>. This is because GRE packets are formed by the addition of the original packets and the required GRE headers<sup>1</sup>. 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<sup>1</sup>. The most common IP MTU is 1500-bytes in length (Ethernet)<sup>1</sup>. 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<sup>1</sup>. 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<sup>1</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>**

