



*Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt. ----- Guaranteed.*



JN0-351 Dumps  
JN0-351 Braindumps  
JN0-351 Real Questions  
JN0-351 Practice Test  
JN0-351 Actual Questions



**Juniper**


# JN0-351

*Enterprise Routing and Switching, Specialist (JNCIS-ENT)*



## Question: 216

Exhibit.



The exhibit shows a terminal window titled "Exhibit" with a blue header bar. The terminal output displays the command `user@R1> show route receive-protocol bgp 10.36.1.4` and its result. The output indicates 33 destinations and 57 routes (33 active, 0 holddown, 0 hidden). A table of routes follows, with columns for Prefix, Nexthop, MED, Lc1pref, and AS path. The routes listed are:

Prefix	Nexthop	MED	Lc1pref	AS path
* 10.30.100.8/32	10.36.1.4			65401 65520 I
* 10.30.100.9/32	10.36.1.4			65401 65521 I
* 10.30.189.0/30	10.36.1.4			65401 65521 I
10.32.1.0/30	10.36.1.4			65401 I
* 10.32.2.0/30	10.36.1.4			65401 I
* 10.32.12.0/30	10.36.1.4			65401 I
* 10.52.100.2/32	10.36.1.4			65401 I

You want to verify prefix information being sent from 10.36.1.4.

Which two statements are correct about the output shown in the exhibit? (Choose two.)

- A. The routes displayed have traversed one or more autonomous systems.
- B. The output shows routes that were received prior to the application of any BGP import policies.
- C. The output shows routes that are active and rejected by an import policy.
- D. The routes displayed are being learned from an I BGP peer.

**Answer: A,B**

Explanation:

The output shown in the exhibit is the result of the command `show ip bgp neighbor 10.36.1.4 received-routes`, which displays all received routes (both accepted and rejected) from the specified neighbor.

Option A is correct, because the routes displayed have traversed one or more autonomous systems. This can be seen from the AS\_PATH attribute, which shows the sequence of AS numbers that the route has passed through. For example, the route 10.0.0.0/8 has an AS\_PATH of 65001 65002, which means that it has traversed AS 65001 and AS 65002 before reaching the local router.

Option B is correct, because the output shows routes that were received prior to the application of any BGP import policies. This can be seen from the fact that some routes have a status code of `R`, which means that they are rejected by an import policy. The `received-routes` keyword shows the routes coming from a given neighbor before the inbound policy has been applied. To see the routes after the inbound policy has been applied, the `routes` keyword should be used instead.

Option C is incorrect, because the output does not show routes that are active and rejected by an import policy. The status code of `R` means that the route is rejected by an import policy, but it does not mean that it is active. The status code of `B` means that the route is active and selected as the best path. None of the routes in the output have both `R` and `B` status codes.

Option D is incorrect, because the routes displayed are not being learned from an IBGP peer. An IBGP peer is a BGP neighbor that belongs to the same AS as the local router. The output shows that the neighbor 10.36.1.4 has a remote AS of 65001, which is different from the local AS of 65002. Therefore, the neighbor is an EBGP peer, not an IBGP peer.

### Question: 217

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 seconds<sup>1</sup>. The keepalive time is the interval at which BGP sends keepalive messages to maintain the connection with its peer<sup>1</sup>. If the keepalive message is not received within the hold time, the connection is considered lost<sup>1</sup>. By default, the hold time is three times the keepalive time, which is 180 seconds<sup>1</sup>.

### Question: 218

Which two statements are correct about tunnels? (Choose two.)

- A. BFD cannot be used to monitor tunnels.
- B. Tunnel endpoints must have a valid route to the remote tunnel endpoint.
- C. IP-IP tunnels are stateful.
- D. Tunnels add additional overhead to packet size.

### Answer: A,B,D

Explanation:

A tunnel is a connection between two computer networks, in which data is sent from one network to another through an encrypted link. Tunnels are commonly used to secure data communications between two networks or to connect two networks that use different protocols.

Option B is correct, because tunnel endpoints must have a valid route to the remote tunnel endpoint. A tunnel endpoint is the device that initiates or terminates a tunnel connection. For a tunnel to be established, both endpoints must be able to reach each other over the underlying network. This means that they must have a valid route to the IP address of the remote endpoint<sup>1</sup>.

Option D is correct, because tunnels add additional overhead to packet size. Tunnels work by encapsulating packets: wrapping packets inside of other packets. This means that the original packet becomes the payload of the surrounding packet, and the surrounding packet has its own header and trailer. The header and trailer of the surrounding packet add extra bytes to the packet size, which is called overhead. Overhead can reduce the efficiency and performance of a network, as it consumes more bandwidth and processing power<sup>2</sup>.

Option A is incorrect, because BFD can be used to monitor tunnels. 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. BFD can also be used to monitor the connectivity of tunnels, such as GRE, IPsec, or MPLS.

Option C is incorrect, because IP-IP tunnels are stateless. IP-IP tunnels are a type of tunnels that use IP as both the encapsulating and encapsulated protocol. IP-IP tunnels are simple and easy to configure, but they do not provide any security or authentication features. IP-IP tunnels are stateless, which means that they do not keep track of the state or status of the tunnel connection. Stateless tunnels do not require any signaling or negotiation between the endpoints, but they also do not provide any error detection or recovery mechanisms.

Reference: 1: What is Tunneling? | Tunneling in Networking 2: What Is Tunnel In Networking, Its Types, And Its Benefits? : [Configuring Bidirectional Forwarding Detection] : [IP-IP Tunneling]

## Question: 219

Which statement is correct about IP-IP tunnels?

- A. IP-IP tunnels only support encapsulating IP traffic.
- B. IP-IP tunnels only support encapsulating non-IP traffic.
- C. The TTL in the inner packet is decremented during transit to the tunnel endpoint.
- D. There are 24 bytes of overhead with IP-IP encapsulation.

## Answer: A

Explanation:

IP-IP tunnels are a type of tunnels that use IP as both the encapsulating and encapsulated protocol. IP-IP tunnels are simple and easy to configure, but they do not provide any security or authentication features. IP-IP tunnels only support encapsulating IP traffic, which means that the payload of the inner packet must be an IP packet. IP-IP tunnels cannot encapsulate non-IP traffic, such as Ethernet frames or MPLS labels<sup>1</sup>.

Option A is correct, because IP-IP tunnels only support encapsulating IP traffic. Option B is incorrect, because IP-IP tunnels only support encapsulating non-IP traffic. Option C is incorrect, because the TTL in the inner packet is not decremented during transit to the tunnel endpoint. The TTL in the outer packet is decremented by each router along the path, but the TTL in the inner packet is preserved

until it reaches the tunnel endpoint<sup>2</sup>. Option D is incorrect, because there are 20 bytes of overhead with IP-IP encapsulation. The overhead consists of the header of the outer packet, which has a fixed size of 20 bytes for IPv4<sup>3</sup>.

Reference:

1: IP-IP Tunneling 2: What is tunneling? | Tunneling in networking 3: IPv4 - Header

### Question: 220

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: A,D

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 adjacencies<sup>1</sup>. If the MTU is less than this, IS-IS adjacencies will not be established<sup>1</sup>.

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-IS<sup>1</sup>. If the loopback interface is not included in IS-IS, it could prevent IS-IS adjacencies from being established<sup>1</sup>. Therefore, options A and D are correct.

### Question: 221

You are asked to create a new firewall filter to evaluate Layer 3 traffic that is being sent between VLANs. In this scenario, which two statements are correct? (Choose two.)

- A. You should create a family Ethernet-switching firewall filter with the appropriate match criteria and actions.
- B. You should apply the firewall filter to the appropriate VLA
- C. You should create a family inet firewall filter with the appropriate match criteria and actions.
- D. You should apply the firewall filter to the appropriate IRB interface.

### Answer: A,C

Explanation:

A firewall filter is a configuration that defines the rules that determine whether to forward or discard packets at specific processing points in the packet flow. A firewall filter can also modify the attributes of the packets, such as priority, marking, or logging. A firewall filter can be applied to various interfaces, protocols, or routing instances on a Juniper device<sup>1</sup>.



A firewall filter has a family attribute, which specifies the type of traffic that the filter can evaluate. The family attribute can be one of the following: inet, inet6, mpls, vpls, iso, or ethernet-switching<sup>2</sup>. The family inet firewall filter is used to evaluate IPv4 traffic, which is the most common type of Layer 3 traffic on a network.

To create a family inet firewall filter, you need to specify the appropriate match criteria and actions for each term in the filter. The match criteria can include various fields in the IPv4 header, such as source address, destination address, protocol, port number, or DSCP value. The actions can include accept, discard, reject, count, log, policer, or next term<sup>3</sup>.

To apply a firewall filter to Layer 3 traffic that is being sent between VLANs, you need to apply the filter to the appropriate IRB interface. An IRB interface is an integrated routing and bridging interface that provides Layer 3 functionality for a VLAN on a Juniper device. An IRB interface has an IP address that acts as the default gateway for the hosts in the VLAN. An IRB interface can also participate in routing protocols and forward packets to other VLANs or networks<sup>4</sup>.

Therefore, option C is correct, because you should create a family inet firewall filter with the appropriate match criteria and actions. Option D is correct, because you should apply the firewall filter to the appropriate IRB interface.

Option A is incorrect, because you should not create a family ethernet-switching firewall filter with the appropriate match criteria and actions. A family ethernet-switching firewall filter is used to evaluate Layer 2 traffic on a Juniper device. A family ethernet-switching firewall filter can only match on MAC addresses or VLAN IDs, not on IP addresses or protocols<sup>5</sup>.

Option B is incorrect, because you should not apply the firewall filter to the appropriate VLAN. A VLAN is a logical grouping of hosts that share the same broadcast domain on a Layer 2 network. A VLAN does not have an IP address or routing capability. A firewall filter cannot be applied directly to a VLAN; it must be applied to an interface that belongs to or connects to the VLAN<sup>6</sup>.

Reference: 1: Firewall Filters Overview 2: Configuring Firewall Filters 3: Configuring Firewall Filter Match

Conditions and Actions 4: Understanding Integrated Routing and Bridging Interfaces 5: Configuring

Ethernet-Switching Firewall Filters 6: Understanding VLANs

## Question: 222

Exhibit

```
user@host# show
protocols {
  oam {
    gre-tunnel {
      interface gr-1/1/10.1 {
        keepalive-time 10;
        hold-time 10;
      }
    }
  }
  lldp {
    interface all;
  }
}
```

You have configured a GRE tunnel. To reduce the risk of dropping traffic, you have configured a keepalive OAM probe to monitor the state of the tunnel; however, traffic drops are still occurring.

Referring to the exhibit, what is the problem?

- A. For GRE tunnels, the OAM protocol requires that the BFD protocols also be used.
- B. The "event link-adjacency-loss" option must be set.
- C. LLDP needs to be removed from the gr-1/1/10.1 interface.
- D. The hold-time value must be two times the keepalive-time value

**Answer: D**

Explanation:

A keepalive OAM probe is a mechanism that can be used to monitor the state of a GRE tunnel and detect any failures in the tunnel path. A keepalive OAM probe consists of sending periodic packets from one end of the tunnel to the other and expecting a reply. If no reply is received within a specified time, the tunnel is considered down and the line protocol of the tunnel interface is changed to down1.

To configure a keepalive OAM probe for a GRE tunnel, you need to specify two parameters: the keepalive-time and the hold-time. The keepalive-time is the interval between each keepalive packet sent by the local router. The hold-time is the maximum time that the local router waits for a reply from the remote router before declaring the tunnel down2.

According to the Juniper Networks documentation, the hold-time value must be two times the keepalive-time value for a GRE tunnel2. This is because the hold-time value must account for both the round-trip time of the keepalive packet and the processing time of the remote router. If the hold-time value is too small, it may cause false positives and unnecessary tunnel flaps.

In the exhibit, the configuration shows that the keepalive-time is set to 10 seconds and the hold-time

is set to 15 seconds for the gr-1/1/10.1 interface. This means that the local router will send a keepalive packet every 10 seconds and will wait for 15 seconds for a reply from the remote router. However, this hold-time value is not two times the keepalive-time value, which violates the recommended configuration. This may cause traffic drops if the remote router takes longer than 15 seconds to reply.

Therefore, option D is correct, because the hold-time value must be two times the keepalive-time value for a GRE tunnel. Option A is incorrect, because BFD is not required for GRE tunnels; BFD is another protocol that can be used to monitor tunnels, but it is not compatible with GRE keepalives<sup>3</sup>. Option B is incorrect, because the `event link-adjacency-loss` option is not related to GRE tunnels; it is an option that can be used to trigger an action when a link goes down<sup>4</sup>. Option C is incorrect, because LLDP does not need to be removed from the gr-1/1/10.1 interface; LLDP is a protocol that can be used to discover neighboring devices and their capabilities, but it does not interfere with GRE tunnels<sup>5</sup>.

Reference:

1: Configuring Keepalive Time and Hold time for a GRE Tunnel Interface 2: keepalive | Junos OS |

Juniper Networks 3: Configuring Bidirectional Forwarding Detection 4: event link-adjacency-loss |

Junos OS | Juniper Networks 5: Understanding Link Layer Discovery Protocol

## Question: 223

Exhibit



```

user@R1> show bgp neighbor
Peer: 10.32.1.2+63645 AS 65401 Local: 10.32.1.1+179 AS 65400
  Description: EBGp peering to 10.32.1.2
  Group: IPCLOS_eBGP          Routing-Instance: master
  Forwarding routing-instance: master
  Type: External      State: Established      Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Export: [ IPCLOS_BGP_EXP ] Import: [ IPCLOS_BGP_IMP ]
  Options: <Preference PeerAS Multipath LocalAS Refresh>
  Options: <VpnApplyExport MtuDiscovery MultipathAs BfdEnabled>
  Holdtime: 90 Preference: 170 Local AS: 65400 Local System AS: 0
  Number of flaps: 0
  Peer ID: 10.52.100.2      Local ID: 10.52.100.1      Active Holdtime: 90
  Keepalive Interval: 30    Group index: 0      Peer index: 0      SNMP
index: 0
  I/O Session Thread: bgpio-0 State: Enabled
  BFD: enabled, up
  Local Interface: ge-0/0/1.0
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  Restart flag received from the peer: Notification
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer does not support LLGR Restarter functionality
  Peer supports 4 byte AS extension (peer-as 65401)
  Peer does not support Addpath
  Table inet.0 Bit: 20000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:          6
    Received prefixes:        9
    Accepted prefixes:        9
    Suppressed due to damping: 0
    Advertised prefixes:      22
  Last traffic (seconds): Received 22 Sent 10 Checked 69617
  Input messages: Total 2568 Updates 4 Refreshes 0 Octets 48991
  Output messages: Total 2572 Updates 8 Refreshes 0 Octets 49362
  Output Queue[1]: 0 (inet.0, inet-unicast)

```

You are a network operator troubleshooting BGP connectivity.

Which two statements are correct about the output shown in the exhibit? (Choose two.)

- A. Peer 10.32.1.2 is configured for AS 63645.
- B. The BGP session is not established.
- C. The R1 is configured for AS 65400.
- D. The routers are exchanging IPv4 routes.

**Answer: A,B,C**

Explanation:

Option B suggests that the BGP session is not established. This is correct because in the output, the state of the BGP session is shown as `Idle`. In BGP, an `Idle` state means that the BGP session is not currently established.

Option C suggests that R1 is configured for AS 65400. This is also correct because in the output, it's shown that the local AS number is 654001. The local AS number represents the Autonomous System (AS) number of the router on which you're checking the BGP session1.

### Question: 224

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.

### Question: 17

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.)

- 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

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 peer1.

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: 225

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.)

- 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: A,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: 226

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?

- 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<sup>12</sup>.

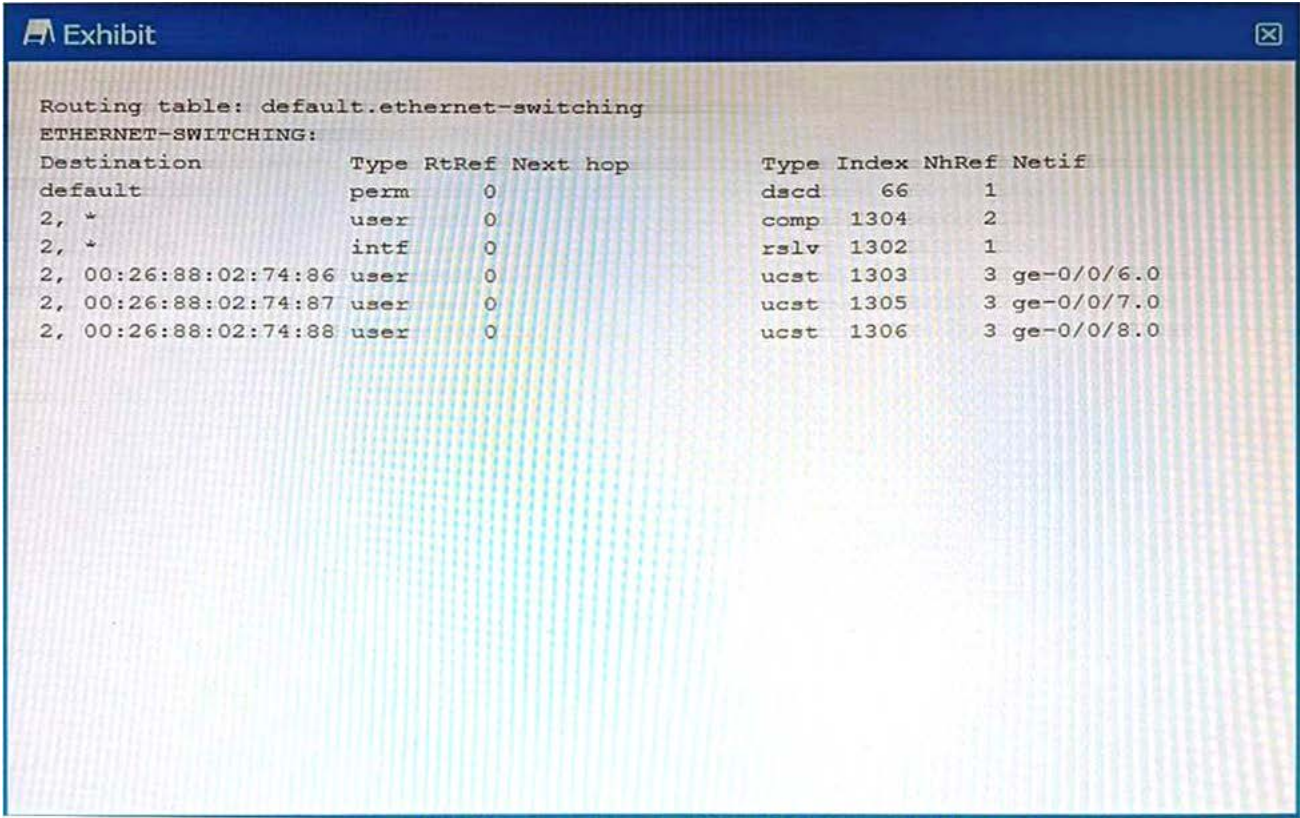
The voice VLAN feature in EX-series switches enables access ports to accept both data (untagged) and voice (tagged) traffic and separate that traffic into different VLANs<sup>12</sup>. This allows the switch to differentiate between voice and data



traffic, ensuring that voice traffic can be treated with a higher priority<sup>12</sup>. Therefore, option D is correct.

**Question: 227**

Exhibit



Which command displays the output shown in the exhibit?

- 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 interface<sup>1</sup>.

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 name<sup>1</sup>.

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.



# SAMPLE QUESTIONS

*These questions are for demo purpose only. **Full version** is up to date and contains actual questions and answers.*

*Killexams.com is an online platform that offers a wide range of services related to certification exam preparation. The platform provides actual questions, exam dumps, and practice tests to help individuals prepare for various certification exams with confidence. Here are some key features and services offered by Killexams.com:*

**Actual Exam Questions:** *Killexams.com provides actual exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these actual questions, candidates can familiarize themselves with the content and format of the real exam.*

**Exam Dumps:** *Killexams.com offers exam dumps in PDF format. These dumps contain a comprehensive collection of questions and answers that cover the exam topics. By using these dumps, candidates can enhance their knowledge and improve their chances of success in the certification exam.*

**Practice Tests:** *Killexams.com provides practice tests through their desktop VCE exam simulator and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice tests cover a wide range of questions and enable candidates to identify their strengths and weaknesses.*

**Guaranteed Success:** *Killexams.com offers a success guarantee with their exam dumps. They claim that by using their materials, candidates will pass their exams on the first attempt or they will refund the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exams.*

**Updated Content:** *Killexams.com regularly updates its question bank and exam dumps to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.*

**Technical Support:** *Killexams.com provides free 24x7 technical support to assist candidates with any queries or issues they may encounter while using their services. Their certified experts are available to provide guidance and help candidates throughout their exam preparation journey.*

For More exams visit <https://killexams.com/vendors-exam-list>  
Kill your exam at First Attempt....Guaranteed!