

Juniper JN0-352


Enterprise Routing and Switching, Specialist

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

- 1. Up to Date products, reliable and verified.**
- 2. Questions and Answers in PDF Format.**



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/jn0-352>

Latest Version: 6.0

Question: 1

A GRE tunnel is experiencing fragmentation issues. You confirm that the tunnel is up and is functioning correctly. You confirm that hosts are sending 1500-byte packets. In this scenario, which statement is correct?

- A. You should use the clear-dont-fragment parameter to allow fragmentation.
- B. You should lower the physical interface's MTU to 1476 bytes to avoid fragmentation.
- C. You should enable BFD to more rapidly identify tunnel link failures.
- D. You should use ToS copying to reduce header size.

Answer: B

Explanation:

GRE encapsulation adds a fixed 24 bytes of overhead to every packet it carries — 4 bytes for the GRE header itself and 20 bytes for the new outer IPv4 delivery header. When a host transmits a full 1500-byte Ethernet payload into a gr- interface, the resulting encapsulated packet becomes 1524 bytes, which exceeds the physical interface's standard 1500-byte MTU and forces fragmentation or drops if the don't-fragment bit is set. The correct remediation is to reduce the effective MTU seen by end hosts so that encapsulated packets never exceed the physical link's transmission limit. Junos automatically applies this logic to gr- logical interfaces, which default to a 1476-byte protocol MTU (1500 minus 24), but when the underlying physical interface or a manually configured path is involved, administrators must explicitly size the interface MTU to 1476 bytes to prevent post-encapsulation oversize packets. Allowing fragmentation with clear-dont-fragment is a workaround that increases CPU load and can degrade performance rather than solving the root cause, BFD addresses link-failure detection rather than MTU sizing, and ToS copying affects only the type-of-service byte, not packet length. Reference topics: Junos Enterprise Routing – Tunneling, GRE Encapsulation and MTU Considerations; Junos OS Configuring GRE Tunnel Interfaces.

Question: 2

How would you view the metric assigned to a route in OSPF? (Choose two.)

- A. Use the show route protocol ospf command.
- B. Use the show ospf route intra command.
- C. Use the show ospf database extensive command.
- D. Use the show ospf interface command.

Answer: A, B

Explanation:

The metric that OSPF assigns to a destination is visible from two complementary vantage points in Junos. The `show route protocol ospf` command displays the main routing table filtered to OSPF-learned prefixes, and each entry shows the computed cost alongside the next hop, exactly as it was installed after SPF calculation. The `show ospf route` command (which accepts filters such as intra-area, inter-area, and extern) presents the OSPF-specific routing table, organized by route type, and explicitly lists the metric column for every intra-area, inter-area, and external route the local router has calculated. Together these two commands give both the RIB-level and the protocol-level view of route cost. By contrast, `show ospf database` extensive dumps the raw link-state advertisements, where metrics appear only as link-level values buried inside Router or Network LSAs rather than as a resolved route cost, so it is not the direct tool for viewing a route's metric. `show ospf interface` reports interface operational state, area, and DR/BDR information but does not present a cost or metric field at all in its standard output. Candidates should be comfortable distinguishing the RIB-oriented and protocol-table-oriented verification commands, since JNCIS-ENT scenarios frequently test whether a candidate reaches for the correct `show` command layer during troubleshooting. Reference topics: Junos Enterprise Routing – OSPF Operation and Verification, Monitoring OSPF.

Question: 3

You have enabled MACsec on two directly connected Ethernet devices but MACsec is not working.

Which two actions will solve this problem? (Choose two.)

- A. Configure the switch interface as a trunk port.
- B. Configure the switch interface as an access port.
- C. Verify the CAK and CKN match on both sides of the link.
- D. Verify the MTU on both sides of the link.

Answer: C, D

Explanation:

MACsec (IEEE 802.1AE) secures a point-to-point Ethernet link independently of how the logical unit above it is configured for VLAN tagging, so neither trunk mode nor access mode has any bearing on whether the MACsec Key Agreement (MKA) protocol can establish a secure channel. What does determine success is the pre-shared key material: when static CAK security mode is used, the Connectivity Association Key (CAK) and the Connectivity Association Key Name (CKN) must be configured identically on both ends of the link. If they do not match, MKA never completes the exchange, no secure channel is created, and all traffic on that interface is dropped rather than merely left unencrypted. Equally important, and frequently overlooked, is that MACsec adds up to 32 bytes of overhead per frame for the security tag (SecTAG) and integrity check value (ICV); Juniper's official guidance is to ensure the difference between the

interface's physical MTU and the protocol MTU is large enough to absorb this overhead, or frames will be silently discarded once encryption is active, producing symptoms that look identical to a failed session. Both the CAK/CKN and MTU checks are documented first-line troubleshooting steps for MACsec on EX Series switches. Reference topics: Junos Enterprise Switching – Layer 2 Security, Configuring and Troubleshooting MACsec on EX Series Switches.

Question: 4

What is the purpose of the native VLAN feature on the Juniper Networks EX Series Switches?

- A. It enables an access port to accept and forward untagged traffic without a tag assigned.
- B. It assigns a specific VLAN tag to all access ports that do not have a specific VLAN assigned.
- C. It enables a trunk port to accept and forward untagged traffic.
- D. It restricts traffic on a trunk port to traffic that matches the configured native-vlan-id.

Answer: C

Explanation:

A native VLAN is a trunk-port concept, not an access-port concept. On an 802.1Q trunk, every frame is normally expected to carry a VLAN tag identifying its membership; the native VLAN is the single exception that Junos permits an administrator to define using the native-vlan-id statement under the trunk interface's family ethernet-switching hierarchy. Any untagged frame arriving on that trunk port is automatically classified into the configured native VLAN, and conversely, outbound frames belonging to the native VLAN are transmitted untagged rather than with an 802.1Q header. This behavior exists chiefly for interoperability with legacy or third-party devices that either cannot generate 802.1Q tags or intentionally send management or default traffic untagged across an otherwise tagged trunk. Access ports, by definition, only ever carry a single VLAN's traffic and never receive tagged frames in normal operation, so the native VLAN mechanism has no relevance there — which eliminates the access-port-oriented answer choices. The native VLAN also does not restrict a trunk to a single VLAN; the trunk continues to carry all configured tagged VLANs simultaneously, with the native VLAN simply being the designated untagged exception. Misconfiguring or mismatching native VLAN IDs between two trunk peers is a classic Layer 2 troubleshooting scenario on the exam. Reference topics: Junos Enterprise Switching – VLANs, Configuring the Native VLAN on a Trunk Interface.

Question: 5

Refer to Exhibit:

```

filter TRAFFIC_MONITOR {
    term track-icmp {
        from {
            protocol icmp;
        }
        then {
            count icmp-counter;
            log;
        }
    }
    term allow-tcp-established {
        from {
            tcp-established;
        }
        then {
            count tcp-counter;
            accept;
        }
    }
    term block-rest {
        then {
            count blocked-counter;
            discard;
        }
    }
}

```

You have applied a firewall filter as an ingress filter on the ge-0/0/0 interface of your Juniper Networks EX Series Switch as shown in the exhibit. An external host sends an ICMP ping packet to the router.

Which statement is correct about how this packet is processed?

- A. The icmp-counter and blocked-counter are both incremented and the packet is dropped.
- B. The icmp-counter is incremented and the packet is dropped.
- C. The icmp-counter is incremented and the packet is accepted.
- D. The icmp-counter and tcp-counter are both incremented and the packet is accepted.

Answer: A

Explanation:

Junos firewall filters evaluate terms sequentially from top to bottom, and a term is only terminal for a packet if its then clause contains an explicit terminating action such as accept, discard, reject, or next term without a match. The count and log actions used in track-icmp are non-terminating; they update statistics and generate a log entry but do not stop evaluation, so a matching ICMP packet falls through to the next term after being counted. The second term, allow-tcp-established, matches only TCP packets carrying the ACK or RST flags as defined by the tcp-established keyword; an ICMP echo request does not satisfy protocol tcp, so this term is skipped entirely and does not accept the packet. Evaluation therefore reaches the implicit catch-all term block-rest, which has no from conditions and matches every remaining packet, incrementing blocked-counter and discarding it. The net effect for the described ICMP ping is that icmp-counter is incremented by the first term, the packet is neither accepted nor rejected by the second term, and blocked-counter is incremented immediately before the packet is silently discarded by the final term. This layered-term behavior, where counting is independent of the accept/discard decision, is a frequently tested nuance of Junos firewall filter processing logic. Reference topics: Junos Enterprise Switching – Layer 2 Security and Firewall Filters, Filter Term Evaluation Order.

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>