



Juniper JN0-352 JNCIS-ENT Practice Questions

Shared by Strickland on 19-10-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: MultipleChoice

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

Options:

- 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 2

Question Type: MultipleChoice

Which IS-IS PDU should be used by a router to list all LSPs that it believes should exist in the

link-state database?

Options:

- A- CSNP
- B- PSNP
- C- LSP
- D- IIH

Answer:

A

Explanation:

The Complete Sequence Number PDU (CSNP) is the IS-IS mechanism used to summarize the entire contents of a router's link-state database, listing the LSP ID, sequence number, checksum, and remaining lifetime for every LSP the sending router currently believes should exist in the database for that level. On broadcast (LAN) circuits, the elected Designated Intermediate System periodically transmits CSNPs to all routers on the segment as a database-synchronization mechanism; any receiving router compares the summarized list against its own local database, and any discrepancy --- an LSP it is missing entirely, or one for which it holds an older sequence number --- triggers a targeted request for the specific, updated LSP content. On point-to-point circuits, a single CSNP is exchanged once when the adjacency first comes up, serving the identical synchronization purpose without needing periodic repetition. The Partial Sequence Number PDU (PSNP), by contrast, is used to acknowledge receipt of specific LSPs or to explicitly request specific LSPs identified as missing or outdated after a CSNP comparison, making it a partial, targeted PDU rather than a complete database listing. The LSP itself is the PDU that actually carries the detailed topology and reachability information being advertised, rather than a database summary. The IS-IS Hello PDU (IIH) is used purely for neighbor discovery and adjacency maintenance and carries no link-state database summary information whatsoever. Reference topics: Junos Enterprise Routing -- IS-IS, CSNP and Link-State Database Synchronization.

Question 3

Question Type: MultipleChoice

Which IS-IS TLV should be used to advertise information about the router's neighbors and link metrics?

Options:

- A- TLV 22
- B- TLV 2
- C- TLV 132
- D- TLV 129

Answer:

A


Explanation:

TLV 22, the Extended IS Reachability TLV defined in RFC 5305, is the modern and Junos-preferred mechanism for advertising a router's adjacent neighbors together with the metric associated with each link. It carries the seven-octet system ID and pseudonode number of each neighbor along with a 24-bit wide metric field, and supports an extensible set of sub-TLVs used for traffic-engineering data such as interface and neighbor IP addresses, bandwidth, and administrative groups. The original IS Reachability TLV, TLV 2, defined in the base ISO 10589 standard, carries similar neighbor and metric information but is constrained to a 6-bit narrow metric with a maximum value of 63, which is inadequate for expressing meaningful cost differentiation in modern, high-capacity enterprise and service-provider topologies; consequently, Junos deployments are commonly configured for wide-metrics-only, favoring TLV 22 over the legacy TLV 2 whenever traffic-engineering-capable or simply more granular metric representation is desired. TLV 132, the IP Interface Address TLV, only advertises the IPv4 addresses configured on the originating router's interfaces and carries no neighbor-adjacency or metric information. TLV 129, the Protocols Supported TLV, merely lists which network-layer protocols (such as IPv4 or IPv6) the originating router is capable of routing, and likewise has nothing to do with neighbor reachability or link cost. Reference topics: Junos Enterprise Routing -- IS-IS, Extended IS Reachability and Wide Metrics.

Question 4

Question Type: MultipleChoice

[Exhibit]



Click the Exhibit button.

You configure the aggregate route 172.16.0.0/16 on router R1. The routing table currently contains active routes for 172.16.10.0/24 and 172.16.99.0/24. No other more-specific routes exist.

Referring to the exhibit, which statement describes what R1 would do with traffic destined to 172.16.200.5?

Options:

- A- R1 forwards the packet to the default route.
- B- R1 drops the packet and deactivates the aggregate route.
- C- R1 drops the packet and sends an ICMP unreachable message.
- D- R1 forwards the packet using the next hop of the active contributing route.

Answer:

C

Explanation:

The destination address 172.16.200.5 falls within the 172.16.0.0/16 aggregate's summarized range but does not fall within either of the two active, more-specific contributing routes actually present in the table --- 172.16.10.0/24 and 172.16.99.0/24 --- meaning no route more specific than the /16 itself exists to cover it. Under Junos's longest-match forwarding logic, the lookup for 172.16.200.5 therefore resolves to the aggregate route itself, and since aggregate routes are installed by default with a reject next hop rather than any real forwarding path, the router drops

the packet and simultaneously generates an ICMP destination-unreachable message back to the originating source, explicitly signaling that no valid path exists for that specific destination even though a covering summary is being advertised. This reject behavior is precisely why aggregate routes are valuable for reducing the number of routes advertised upstream while still providing clear, immediate feedback for traffic aimed at address space within the summary that has no genuine underlying route, rather than silently black-holing it or misdirecting it toward an unrelated contributing route's next hop. The aggregate route itself remains active and installed throughout this process --- a gap in contributing coverage does not deactivate the aggregate, since the aggregate's activation depends only on at least one contributing route being active, which is satisfied here by both existing /24 blocks. There is also no default route present in this scenario to fall back upon. Reference topics: Junos Enterprise Routing -- Protocol Independent Routing, Aggregate Route Reject Behavior for Uncovered Address Space.



Question 5

Question Type: MultipleChoice

[Exhibit]

```
user@R1> show route advertising-protocol bgp 192.168.1.8
XYZ.inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
  Prefix                Nexthop          MED      Lclpref   AS path
* 172.16.1.0/24         Self              100        100       I
* 192.168.150.0/24      Self              100        100       I
* 192.168.250.0/24      Self              100        100       I
```

Click the Exhibit button.

Which three statements about this BGP session are correct? (Choose three.)



Options:

- A- The export routing policy has been processed.
- B- The routes are active.
- C- The routes are stored in the default routing instance.
- D- The routes originate from an external AS.
- E- R1 modifies the BGP next hop.

Answer:

A, B, E

Explanation:

The `show route advertising-protocol bgp` command displays exactly what would be transmitted to a specific BGP neighbor after all configured export routing policy has already been evaluated and applied, meaning any prefixes shown in this output have successfully passed through the export policy chain --- confirming the export routing policy has been processed. The command header explicitly states '7 destinations, 7 routes (7 active, ...)' and every listed prefix is preceded by an asterisk, Junos's standard marker for an active route, confirming all seven displayed routes are active in the local routing table. The routing table name shown, XYZ.inet.0, indicates these routes reside in a named routing instance called XYZ rather than the default master routing instance, whose table would instead be labeled simply inet.0; this rules out the statement claiming the routes belong to the default instance. The AS path column shows only the origin code 'I' (IGP) with no AS numbers prepended ahead of it, meaning the AS path is empty, which indicates these prefixes were originated locally within R1's own AS (through redistribution or local origination) rather than having been learned from and passed through an external AS, ruling out that statement as well. Finally, the Nexthop column explicitly shows 'Self' for every route rather than an actual peer or learned address, confirming R1 is rewriting the next-hop attribute to its own address before advertising --- a deliberate next-hop-self policy action typically applied toward IBGP peers. Reference topics: Junos Enterprise Routing -- BGP, Verifying Advertised Routes and Export Policy Effects.

Question 6

Question Type: MultipleChoice

Refer to Exhibit:

```
[edit interfaces ge-0/0/0]
user@R1# show
  unit 0 {
    family inet {
      address 10.10.10.1/24 {
        vrrp-group 5 {
          virtual-address 10.100.100.1;
          priority 150;
        }
      }
    }
  }
```

R1 is not responding to ICMP requests sent to the VIP address.

Referring to the exhibit, which command option would you configure on the virtual address to correct this problem?

Options:

- A- track
- B- accept-data
- C- advertise-interval
- D- no-preempt

Answer:

B

Explanation:

In this configuration, R1's physical interface address (10.10.10.1) is different from the VRRP virtual IP address (10.100.100.1), meaning R1 is operating as a non-owner VRRP master --- it did not originate the virtual address as one of its own real interface addresses. By default, Junos VRRP masters that do not own the virtual address will forward transit traffic destined through that address but will not process or respond to traffic addressed directly to the VIP itself, such as ICMP echo requests, Telnet, or SSH sessions aimed at 10.100.100.1. This is deliberate default behavior intended to prevent ambiguity about which physical device is answering on behalf of a shared virtual identity. To allow the master to accept and respond to packets whose destination is the virtual address itself, the accept-data statement must be explicitly configured under the vrrp-group hierarchy. Once enabled, the elected master will process locally destined traffic sent to the VIP, resolving exactly the symptom described. The track option is used for interface or route-based priority adjustment to influence mastership, advertise-interval tunes the frequency of VRRP hello advertisements, and no-preempt prevents a higher-priority router from reclaiming mastership once a lower-priority router has taken over --- none of these affect whether the master processes traffic destined to the VIP. Reference topics: Junos Enterprise Routing -- High Availability, VRRP accept-data Behavior.

Question 7

Question Type: MultipleChoice

You are configuring a GRE tunnel between Router A (192.168.1.1) and Router B (192.168.2.1) to connect two isolated networks. The tunnel interfaces (gr-0/0/0) need to support IPv4 traffic.

In this scenario, which configuration ensures the tunnel endpoints are correctly defined and reachable?

Options:

- A- set interfaces gr-0/0/0 unit 0 tunnel source 192.168.1.1 destination 192.168.2.1

- B- set interfaces gr-0/0/0 unit 0 family inet6
- C- set interfaces gr-0/0/0 unit 0 family inet address 192.168.1.1
- D- set interfaces gr-0/0/0 unit 0 tunnel destination 192.168.2.1

Answer:

A

Explanation:

A functioning GRE tunnel interface in Junos requires, at minimum, both a tunnel source address and a tunnel destination address configured together under the tunnel statement of the gr-logical unit, since these two values together define the outer, delivery-layer IP header that will encapsulate every passenger packet sent through the tunnel; the source identifies which of the router's own addresses will originate the encapsulated traffic, and the destination identifies the remote tunnel endpoint the encapsulated traffic must be routed toward across the underlying IP network. The first option correctly supplies both required values in a single, complete statement --- source 192.168.1.1 (Router A's own address) and destination 192.168.2.1 (Router B's address) --- fully defining the tunnel's endpoints and ensuring the encapsulation and route-resolution process can succeed. The fourth option supplies only the destination and omits the source entirely, which is an incomplete configuration; without an explicit source, Junos cannot always reliably determine which local address to use for encapsulation, particularly on multi-address or multi-interface routers, and the tunnel configuration is left non-deterministic or non-functional as shown. The second option enables the IPv6 protocol family on the tunnel's logical unit, which is irrelevant to a scenario explicitly requiring IPv4 passenger traffic support. The third option attempts to assign a family inet address directly using one of the underlying tunnel endpoint addresses, which conflates the tunnel's own transport addressing with logical-unit protocol addressing and does not by itself establish the required source/destination tunnel definition. Reference topics: Junos Enterprise Routing -- Tunneling, Configuring GRE Tunnel Source and Destination.

Question 8

Question Type: MultipleChoice

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?

Options:

- 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.

Question 9

Question Type: MultipleChoice

You enable the point-to-point parameter on your OSPF interface configuration.

In this scenario, what are two benefits of this action? (Choose two.)

Options:

- A- It enables OSPF adjacencies to establish quicker.
- B- It enables the advertisement of a connected OSPF network.
- C- It enables an OSPF router to win the designated router election on a broadcast link.
- D- It reduces the size of the OSPF link-state database.

Answer:

A, D

Explanation:

Configuring interface-type p2p under an OSPF interface tells Junos to treat that interface as a point-to-point network type regardless of the underlying physical medium, which removes the interface from the standard broadcast-network state machine entirely. On a true broadcast network, two routers must first complete a Designated Router and Backup Designated Router election, wait through the associated timers, and only then proceed through 2-Way, ExStart,

Exchange, and Loading before reaching Full --- a multi-step process. A point-to-point interface skips DR/BDR election altogether and proceeds directly from 2-Way into ExStart, so adjacencies form measurably faster, which is the first correct benefit. Because there is no DR on a point-to-point interface, Junos also does not generate a Type 2 Network LSA to describe the segment; the connected subnet is instead advertised as a simple stub-network entry inside the Router LSA itself, eliminating an entire LSA type from the area's link-state database and reducing its overall size, which is the second correct benefit. Point-to-point configuration is entirely incompatible with DR election, so a router configured this way cannot 'win' a DR election --- there is no DR to win. Advertisement of the connected network happens on any operational OSPF interface irrespective of network type, so that option describes normal OSPF behavior rather than a benefit specific to the point-to-point setting. Reference topics: Junos Enterprise Routing -- OSPF Network Types, Configuring OSPF Interface Type.



Question 10

Question Type: MultipleChoice

Two Juniper Networks EX Series Switches are connected with four 1-GbE links configured as a single link aggregation group (LAG). One of the physical member links experiences a failure.

By default, what impact will this failure have on traffic forwarding between the switches?

Options:

- A- The LAG continues forwarding traffic across the remaining member links.
- B- All traffic will be dropped because the LAG goes administratively down.
- C- Only broadcast traffic continues to be forwarded across the LAG.
- D- The LAG renegotiates and temporarily disables all member links.

Answer:

A

Explanation:

A link aggregation group bundles multiple physical Ethernet links into a single logical interface (an ae interface) specifically to provide both increased aggregate bandwidth and link-level redundancy. By default, the aeX interface remains operationally up as long as at least one member link within the bundle remains active and passing traffic; the LAG does not require every member link to be functional to continue operating. When one of the four 1-GbE member links in this scenario fails, LACP (if in use) or the underlying link-status monitoring detects the failure, removes that specific link from the active forwarding set, and redistributes traffic hashing across the three remaining healthy member links automatically and without operator

intervention. End devices and upper-layer protocols continue to see a single, continuously up ae interface throughout this process, experiencing only a proportional reduction in available aggregate bandwidth rather than any interruption in connectivity. This default behavior is precisely why LAG is a foundational building block for resilient EX Series switch interconnects in enterprise and data center designs, since it allows individual optic, cable, or transceiver failures to be absorbed transparently. None of the distractor behaviors --- full administrative shutdown of the bundle, forwarding limited to broadcast traffic only, or a temporary disabling of all member links during renegotiation --- reflect actual Junos LAG behavior, which is designed explicitly to avoid any such disruption. Reference topics: Junos Enterprise Switching -- Layer 2 Switching, Link Aggregation Group Resiliency.

Question 11

Question Type: MultipleChoice

Which statement describes the additional overhead added when a packet is encapsulated in a GRE tunnel?

Options:

- A- GRE adds 12 bytes of overhead to each packet.
- B- GRE adds 32 bytes of overhead to each packet.
- C- GRE adds 24 bytes of overhead to each packet.
- D- GRE adds 20 bytes of overhead to each packet.

Answer:

C

Explanation:

A GRE-encapsulated packet consists of the original passenger packet wrapped inside a GRE header and a new outer IP delivery header, and the combined size of these two additions is what constitutes GRE's per-packet overhead. The base GRE header, as defined in RFC 2784 and used in its minimal, non-keyed, non-checksummed form (the default and most common deployment), is exactly 4 bytes in length, containing the flags/version field and the protocol type field identifying the encapsulated payload's ether type. Layered around that GRE header is a completely new outer IPv4 header, which itself consumes a standard 20 bytes, containing the tunnel source and destination addresses used for delivery across the underlying network. Adding these two components together --- 4 bytes for the GRE header plus 20 bytes for the new outer IPv4 header --- yields a total of exactly 24 bytes of overhead imposed on every encapsulated packet, which is precisely why Junos automatically sets the default protocol MTU on gr- tunnel interfaces to 1476 bytes (1500 minus 24) to prevent oversized, post-encapsulation

packets from requiring fragmentation on a standard 1500-byte Ethernet path. The 20-byte figure alone describes only the new IP header in isolation and omits the GRE header itself, while 12 and 32 bytes do not correspond to any standard, non-keyed GRE encapsulation configuration and are included purely as plausible-looking distractors. Reference topics: Junos Enterprise Routing -- Tunneling, GRE Header Overhead and MTU Planning.



To Get Premium Files for JN0-352 Visit

<https://www.p2pexams.com/products/jn0-352>

For More Free Questions Visit

<https://www.p2pexams.com/juniper/pdf/jn0-352>

20%
DISCOUNT

P2P
exams