

Enterprise Routing and Switching, Specialist (JNCIS-ENT)

Juniper JN0-352

Version Demo

Total Demo Questions: 8

Total Premium Questions: 80

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Layer 2 Switching	6
Topic 2, VLANs	4
Topic 3, Spanning Tree Protocol	8
Topic 4, Layer 2 Security	8
Topic 5, Layer 3 Routing	9
Topic 6, OSPF	13
Topic 7, IS-IS	8
Topic 8, BGP	14
Topic 9, Tunnels	6
Topic 10, High Availability	4
Total	80

QUESTION NO: 1

A distribution switch must remain the root bridge for the spanning-tree topology. A downstream interface receives superior BPDUs from an incorrectly connected switch.

Which spanning-tree protection feature prevents the local interface from becoming a root port through that downstream switch?

- A. BPDU protection
- B. Root protection
- C. Loop protection
- D. Storm control

ANSWER: B

Explanation:

Root protection prevents a port from becoming a root port as a result of receiving superior BPDUs. When a protected interface receives a BPDU that would cause the local switch to select a new root bridge or root port through that interface, the port is placed into a root-inconsistent blocking state. This protects the intended root-bridge placement in the topology.

BPDU protection is intended for edge ports and disables or blocks an interface that receives any BPDU. Loop protection addresses loss of BPDUs rather than reception of superior BPDUs.

QUESTION NO: 2

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:

Verify the CAK and CKN match on both sides of the link. is required because MACsec peers using static connectivity-association keys must authenticate through MKA with the same Connectivity Association Key and Connectivity Association Key Name. These values identify and secure the connectivity association. When the key material is not identical at each endpoint, the peers cannot successfully establish the MACsec secure channel, so protected traffic cannot pass. Confirming the configured CAK and CKN on both directly connected interfaces is therefore a fundamental remediation step.

Verify the MTU on both sides of the link. is also necessary because MACsec encapsulation adds Ethernet-frame overhead, including the security tag and integrity-check information. The interfaces must have sufficient physical MTU available to carry the protected frame while preserving the configured protocol payload size. A mismatched or insufficient MTU can cause encrypted frames to be dropped even after MACsec is configured, making end-to-end connectivity appear broken. Juniper's MACsec documentation describes the MACsec configuration model and the relevant interface requirements; see the [Juniper MACsec documentation](#) and the [Junos MACsec CLI reference](#).

QUESTION NO: 3

[Exhibit]



Click the Exhibit button.

You want the RSTP primary root path from switch C to traverse switch B.

Referring to the exhibit, which solution will accomplish this task?

- A. Set switch B priority to 1.
- B. Set the metrics from switch C directly to switch A higher than the combined metrics of switch C through switch B to switch A.**
- C. Set the metrics from switch C directly to switch A lower than the combined metrics of switch C through switch B to switch A.
- D. Set switch B priority to 0.

ANSWER: B

Explanation:

Set the metrics from switch C directly to switch A higher than the combined metrics of switch C through switch B to switch A. is correct because RSTP selects a nonroot switch's root port based on the best received spanning-tree information, including the lowest root-path cost to the elected root bridge. With switch A established as the root bridge, switch C compares the cumulative cost of its direct path to A with the cumulative cost of the path that reaches A through switch B. To make C use B as its primary path toward the root, the total root-path cost through B must be lower than the cost of the direct C-to-A path. Raising the port cost on C's direct link to A, or otherwise ensuring that direct path's resulting metric exceeds the C-to-B-to-A total, causes the port from C toward B to be selected as C's root port. The direct C-to-A port then remains available as an alternate RSTP path and can transition rapidly if the preferred path fails. Junos supports configuring spanning-tree port cost to influence this path selection behavior; the configured cost is included in the root-path-cost

calculation. See Juniper's [Spanning Tree Protocol configuration documentation](#) and the [Junos STP cost statement reference](#).

QUESTION NO: 4

[Exhibit]

```
user@switch> show configuration interfaces ae0
aggregated-ether-options {
    lacp {
        active;
    }
}
unit 0 {
    family ethernet-switching {
        interface-mode trunk;
        vlan {
            members all;
        }
    }
}
user@switch> show interfaces ae0.0
error: device ae0 not found
```

Click the Exhibit button.

You are troubleshooting an aggregated Ethernet interface on a Juniper Networks EX Series Switch. The aggregated Ethernet interface does not appear in the output shown in the exhibit.

Which action will solve this problem?

- A. Configure the number of aggregated Ethernet devices under the [edit chassis] hierarchy.
- B. Install the aggregated Ethernet interface license on the switch.
- C. Restart the switch to enable the aggregated Ethernet switching process.
- D. Configure LACP passive-mode.

ANSWER: A

Explanation:

Configure the number of aggregated Ethernet devices under the [edit chassis] hierarchy. On Junos OS, an ae logical aggregate interface is available only after the switch has reserved aggregated-Ethernet device instances with the `set chassis aggregated-devices ethernet device-count` configuration statement. For example, configuring `device-count 1` enables the first aggregate device, ae0; a higher value reserves additional instances sequentially. This is a platform-level prerequisite rather than a per-interface LACP setting. Once the appropriate device count is committed, the corresponding ae interface can be configured under `interfaces`, member physical interfaces can be assigned using `ether-options 802.3ad`, and LACP parameters can be applied if dynamic link aggregation is required. Juniper documents the device-count setting as the mechanism used to specify the number of aggregated Ethernet interfaces supported and configured on the device. See the Junos [aggregated-devices chassis statement reference](#) and Juniper's [Aggregated Ethernet Interfaces documentation](#).

QUESTION NO: 5

Which two characteristics describe the default behavior of aggregate routes? (Choose two.)

- A. They advertise all contributing routes individually by default.
- B. They hide internal routing instability from external peers.**
- C. They use a default next hop of discard.
- D. They use a default next hop of reject.**

ANSWER: B D

Explanation:

They hide internal routing instability from external peers. Aggregate routing creates a less-specific summary prefix from contributing, more-specific routes. The aggregate remains active when at least one contributing route is active, so changes affecting individual contributors do not necessarily cause the summary route to be withdrawn or changed. Advertising the stable aggregate instead of exposing each individual route reduces routing-information churn seen by peers outside the summarized portion of the network. This is a key operational benefit of route aggregation: it provides a simpler and more stable routing view to downstream devices while allowing the network to retain more-specific forwarding information internally.

They use a default next hop of reject. In Junos, an aggregate route is installed with a reject next hop by default. This protects against traffic addressed to the aggregate prefix when no active, more-specific contributing route can forward that traffic. A reject next hop discards the packet and returns an ICMP unreachable message. Junos also supports explicitly configuring a discard next hop for an aggregate route when silent packet dropping is desired, but discard is not the default behavior. See Juniper's [Configuring Route Aggregation](#) documentation and the [aggregate route statement reference](#).

QUESTION NO: 6

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

- 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:

GRE adds 24 bytes of overhead to each packet. This value applies to standard GRE carried over IPv4, as normally assumed for Junos GRE tunnel MTU calculations. GRE encapsulation preserves the original packet as the passenger payload, then prepends a minimal GRE header and an outer IPv4 delivery header. The base GRE header is 4 bytes: it includes the flags/version field and the protocol-type field that identifies the encapsulated payload protocol. The new outer IPv4 header is 20 bytes when no IPv4 options are present. Therefore, the encapsulating headers total 24 bytes per packet (4 + 20 bytes). This overhead matters when setting the tunnel MTU, because a 1,500-byte packet sent into a GRE tunnel would become 1,524 bytes after encapsulation. To avoid fragmentation across a conventional 1,500-byte underlying path, the usable payload MTU must account for these 24 bytes. GRE's base header format is defined by RFC 2784, while Junos documentation describes GRE tunnel interfaces and their encapsulation behavior. See [RFC 2784: Generic Routing Encapsulation](#) and [Juniper Networks: Tunnel Interfaces](#).

QUESTION NO: 7

You configure an OSPF area range of 10.10.0.0/16 on an ABR for routes learned from area 0.0.0.10. Several active intra-area routes within that range exist in the area.

Which two results occur when the range is advertised to the other OSPF areas? (Choose two.)

- A. The ABR advertises the 10.10.0.0/16 summary prefix into the other areas.
- B. The ABR suppresses covered, more-specific interarea summary advertisements into the other areas.
- C. The ABR summarizes all external Type 5 LSAs within 10.10.0.0/16.
- D. The ABR removes the covered routes from the source area's link-state database.
- E. The ABR prevents routers in the source area from calculating paths to the covered subnets.

ANSWER: A B

Explanation:

An OSPF area range configured on an ABR summarizes eligible routes from one area when they are advertised into other areas. The ABR advertises the configured summary prefix and suppresses the individual, more-specific interarea summary advertisements that are covered by that range.

Area range configuration does not summarize external Type 5 LSAs originated by an ASBR, and it does not alter the intra-area link-state database or SPF calculation within the source area.

QUESTION NO: 8

Refer to Exhibit:

```
user@R3> show isis interface
IS-IS interface database:
Interface          L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
ge-0/0/1.0         1   0x1 Point to Point    Disabled        10/10
ge-0/0/2.0         1   0x1 Point to Point    Disabled        10/10
ge-0/0/4.0         3   0x1 Point to Point    Point to Point  10/10
ge-0/0/5.0         3   0x1 Point to Point    Point to Point  10/10
lo0.0              3   0x1 Passive          Passive         0/0
```

Click the Exhibit button.

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

- A. A designated router is elected for each IS-IS interface.
- B. R3 cannot establish an L2 adjacency on the ge-0/0/2.0 interface.
- C. R3 uses its loopback interface to establish an L2 IS-IS adjacency.
- D. R3 advertises its loopback interface IP address.

ANSWER: B D

Explanation:

R3 cannot establish an L2 adjacency on the ge-0/0/2.0 interface because the interface output shows an IS-IS level value of 1. In the `show isis interface` display, this identifies the circuit as enabled for Level 1 only; the Level 2 DR field is also shown as disabled. Because Level 2 is not active on that interface, R3 does not send or process Level 2 IS-IS hellos there and therefore cannot form a Level 2 neighbor relationship over that circuit. Junos documents the interface level and designated-router fields reported by this operational command at [show isis interface](#).

R3 advertises its loopback interface IP address because the loopback is configured as a passive IS-IS interface. A passive interface does not participate in IS-IS adjacency establishment, but Junos still advertises the interface's eligible address prefix in the router's link-state information. This is the standard design for loopback addresses: the prefix is reachable through IS-IS while no unnecessary hello traffic or adjacency is attempted on the loopback. Junos configuration documentation describes passive IS-IS interfaces and their advertisement behavior at [passive \(IS-IS\)](#).