

JN0-662^{Q&As}

Service Provider Routing and Switching - Professional (JNCIP-SP)

Pass Juniper JN0-662 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.leads4pass.com/jn0-662.html>

100% Passing Guarantee
100% Money Back Assurance


Following Questions and Answers are all new published by Juniper
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers



QUESTION 1

Click the Exhibit button.



R5 must advertise Subnet A into IS-IS so that Subnet A and Subnet B can communicate. Subnet B must be able to forward traffic to Subnet A and towards Area 49.0001. However, R5 should not be able to route traffic from Subnet A to Area 49.0001.

Referring to the exhibit, how would you solve this problem?

- A. Configure Level 2 on all links in Area 49.0002.
- B. Configure the set protocols isis ignore-attached-bit parameter on R5.
- C. Configure the set protocols isis overload parameter on R6.
- D. Configure an export policy on R6 to reject all routes except Subnet B towards R5.

Correct Answer: B

QUESTION 2

How does Juniper Networks use selective BGP route damping to help stabilize the network?

- A. The damping decay half-life holds down peer links for a defined measure of time to prevent route table reconvergence.

- B. The damping max-suppress timer is set to the lowest possible value by default.
- C. The figure-of-merit value measures route stability by advertisement and withdrawal.
- D. Aggressive IGP hold timers prevent the routes from entering and exiting the BGP route selection process.

Correct Answer: C

QUESTION 3

Click the Exhibit button.

```
user@router> show route protocol bgp advertising-protocol bgp 172.17.10.49 10.16.0.20/30 extensive

inet.0: 64 destinations, 276 routes (63 active, 1 holddown, 0 hidden)
@ 10.16.0.20/30 (6 entries, 2 announced)
  BGP group ce type External
  Nexthop: Self
  AS path: [2856] 65200 ?

user@router> show protocols
  bgp {
    path-selection always-compare-med;
    log-updown;
    graceful-restart;
    group ce {
      type external;
      neighbor 172.17.10.49 {
        hold-time 180;
        cut-delay 0;
        damping;
        import L3vpn-standby;
        family inet {
          unicast {
            prefix-limit {
              maximum 200;
              teardown 80 idle-timeout forever;
            }
          }
        }
        authentication-key "CA0Ihrmf0I"; ## SECRET-DATA
        export L3vpn-ex;
        peer-as 65100;
        multipath;
      }
    }
  }
  multipath;
}
}
```

The route shown in the exhibit is being advertised to the EBGp peer and displays a next hop of itself. However, you do not have a next-hop self policy configured. What would cause this behavior?

- A. The IBGP peers have a next-hop self policy, which the router is exporting to the EBGp neighbors.
- B. The set protocols bgp path-selection as-path-ignore is not set and must be added so the next-hop attribute will propagate from the peer.
- C. The set protocols bgp accept-remote-next hop is not set and must be added so the next-hop attribute will propagate from the peer.

D. The next-hop attribute was modified by default when it was advertised to the EBGP peer, without applying a policy.

Correct Answer: D

QUESTION 4

Click the Exhibit button.

```
user@router> show bgp neighbor 10.222.222.3
Peer: 10.222.222.3+62377 AS 65511 Local: 10.222.222.4+179 AS 65511
  Type: Internal      State: Established      Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress AddressFamily Rib-group Refresh>
  Address families configured: inet-unicast inet-multicast inet-vpn-multicast inet6-vpn-unicast inst-vpn-
flow 12vpn-signaling
  Local Address: 10.222.222.4 Holdtime: 90 Preference: 170
  Number of flaps: 1
  Last flap event: RecvNotify
  Error: 'Cease' Sent: 0 Recv: 1
  Peer ID: 10.222.222.3 Local ID: 10.222.222.4 Active Holdtime: 90
  Keepalive Interval: 30 Peer index: 1
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast inet-multicast inet-vpn-multicast 12vpn inet6-vpn-
unicast inet-vpn-flow
  NLRI advertised by peer: inet-unicast inet-multicast inet-vpn-unicast inet-vpn-multicast 12vpn inet6-
vpn-unicast inet-vpn-flow
  NLRI for this session: inet-unicast inet-multicast inet-vpn-multicast 12vpn inet6-vpn-unicast inst-vpn-
flow
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-unicast inet-multicast inet-vpn-multicast 12vpn inet6-vpn-
unicast inet-vpn-flow
  NLRI of received end-of-rib markers: inet-unicast inet-multicast inet-vpn-multicast 12vpn inet6-vpn-
unicast inet-vpn-flow
  NLRI of all end-of-rib markers sent: inet-unicast inet-multicast inet-vpn-multicast 12vpn inet6-vpn-
unicast inet-vpn-flow
  Peer supports 4 byte AS extension (peer-as 65511)
  Peer does not support Addpath
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes: 0
    Received prefixes: 0
    Accepted prefixes: 0
    Suppressed due to damping: 0
    Advertised prefixes: 0
  Table bgp.invpnflow.0
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: not advertising
    Active prefixes: 0
    Received prefixes: 0
    Accepted prefixes: 0
    Suppressed due to damping: 0
```

The exhibit shows a BGP peering session for two PE routers. The BGP session is up, but IPv4 hosts in the Layer 3 VPN that use the BGP session are unable to communicate.

What is the problem?

- A. The BGP peer does not support the add-path feature.
- B. The BGP peer does not support the restarter functionality.

- C. The local BGP router does not support Layer 2 VPN and Layer 3 VPN NLRI address families at the same time.
- D. There is a mismatch in the supported NLRI address families between the BGP peers.

Correct Answer: D

QUESTION 5

You want to use IS-IS on a GRE interface where the underlying Layer 3 MTU is 1500.

Which statement is correct in this scenario?

- A. IS-IS can be used because every IS-IS interface must be capable of transmitting packets at least as large as 1476 bytes, and the GRE header is 24 bytes.
- B. IS-IS can be used, but the networking device directly attached to the circuit must be capable of fragmentation.
- C. IS-IS cannot be used, but the router can enable a GRE key that serves the same function as IS-IS.
- D. IS-IS cannot be used because the IS-IS hello is not allowed to be fragmented and has the DF bit set.

Correct Answer: B

[JN0-662 PDF Dumps](#)

[JN0-662 Practice Test](#)

[JN0-662 Study Guide](#)