

KTest

高品質 より良いサービス



練習問題

<https://www.ktest.jp>

1年間無料アップデート

Exam : JN0-364

Title : Service Provider Routing
and Switching - Specialist
(JNCIS-SP)

Version : DEMO

1. You must ensure that your routing platform with redundant REs continues to forward packets, even if one RE fails.

Which technology would you use to accomplish this task?

- A. NSB
- B. LAG
- C. BFD
- D. GRES

Answer: D

Explanation:

For Juniper platforms equipped with dual Routing Engines (REs), the fundamental technology required to provide high availability during a hardware or software failure of the primary RE is Graceful Routing Engine Switchover (GRES).

According to Juniper Networks technical documentation, GRES allows the backup RE to stay in a "hot" standby state. When GRES is enabled, the primary RE synchronizes critical state information with the backup RE, specifically the chassis state and the interface state. This synchronization includes the Packet Forwarding Engine (PFE) configuration.

When the primary RE fails, the backup RE takes over immediately. Because the PFE (which resides on the line cards) was already synchronized and is not restarted during the switchover, the router continues to forward packets that are already in flight or part of established flows. This prevents a complete network outage during an RE failover.

Comparison with other options:

NSB (Non-Stop Bridging - Option A): Focuses specifically on maintaining Layer 2 protocol states (like STP) during a switchover.

LAG (Link Aggregation - Option B): Provides redundancy for physical links, not the control plane or the RE.

BFD (Bidirectional Forwarding Detection - Option C): Is a protocol used for rapid detection of link or neighbor failures; it does not protect the RE or maintain forwarding during an internal switchover.

It is important to note that while GRES maintains the forwarding state, it does not by itself maintain the routing protocol state (adjacencies). To keep OSPF or BGP sessions from dropping during the switchover, GRES must be paired with Non-Stop Active Routing (NSR). However, as the question focuses on the core requirement of continuing to forward packets, GRES is the foundational technology.

2. Which two statements regarding GRE and IP-IP tunnels are correct? (Choose two.)

- A. These tunnels add additional overhead to the packets that traverse them.
- B. These tunnels do not add any overhead to the packets that traverse them.
- C. These tunnels offer secure encryption mechanisms.
- D. These tunnels do not offer encryption mechanisms.

Answer: A D

Explanation:

In Juniper Networks Junos OS, Generic Routing Encapsulation (GRE) and IP-in-IP (IP-IP) are common tunneling mechanisms used to transport packets across a network by encapsulating them within another protocol. Understanding the header structure and the limitations of these protocols is essential for proper MTU (Maximum Transmission Unit) management and security design.

Overhead (Option A):

Both GRE and IP-IP tunnels operate by adding an additional IP header to the original packet. An IP-IP tunnel (Protocol 4) adds a 20-byte IPv4 header. A GRE tunnel (Protocol 47) adds the same 20-byte delivery IP header plus a minimum 4-byte GRE header (totaling 24 bytes, which can increase if keys or sequencing are used).

Because these headers are added to the payload, the total size of the packet increases. This "overhead" means that if the original packet was already at the MTU limit (e.g., 1500 bytes), the encapsulated packet will exceed it, potentially leading to fragmentation or the need to adjust the TCP MSS (Maximum Segment Size).

Encryption (Option D):

Crucially, according to Juniper Service Provider documentation, neither GRE nor IP-IP provides native encryption or data confidentiality. They are encapsulation protocols, not security protocols. The payload remains in cleartext and is visible to any device along the path. If security and encryption are required for data traversing these tunnels, they must be combined with IPsec (IP Security). While GRE is often used as the "carrier" for IPsec (to allow multicast or dynamic routing protocols which IPsec alone does not support), the GRE protocol itself remains an unencrypted delivery mechanism. Therefore, statements A and D accurately describe the architectural behavior of these tunnel types.

3. For two or more switches to participate in the same MSTP region, which parameter must match?

- A. Region name
- B. Extended system ID
- C. Root bridge priority
- D. Root bridge ID

Answer: A

Explanation:

Multiple Spanning Tree Protocol (MSTP), as defined in IEEE 802.1s and implemented in Juniper Networks Junos OS, allows for the grouping of VLANs into specific spanning tree instances. This provides significant scalability and load-balancing advantages over traditional STP or RSTP. To achieve this, switches must be grouped into logical "Regions."

According to Juniper documentation, for two or more switches to be considered part of the same MSTP Region, they must possess an identical MSTP Configuration Identifier. This identifier consists of three specific attributes that must match exactly across all participating switches:

MSTI Name (Region Name): A descriptive string (up to 32 characters) that identifies the region.

MSTI Revision Level: A numerical value (0–65535) used to track configuration changes.

VLAN-to-Instance Mapping: The specific table that defines which VLAN IDs are associated with which Multiple Spanning Tree Instances (MSTIs).

If even one of these parameters—such as the Region name (Option A)—differs, the switches will treat each other as being in separate regions. When switches are in different regions, they interact using the Common Spanning Tree (CST), effectively seeing the other region as a single "virtual bridge," which limits the granularity of traffic engineering.

The Extended system ID (Option B) is a component of the Bridge ID used to carry VLAN information in PVST+ but is not a region-matching requirement. Root bridge priority (Option C) and Root bridge ID (Option D) are variables used during the STP election process to determine the topology's root, but they do not define the boundaries of an MSTP region itself.

4. You are monitoring OSPF on a router and notice frequent state changes between Full and Down. Which condition would cause this behavior?

- A. physical interface flapping
- B. route preference mismatch
- C. area ID mismatch
- D. MTU mismatch

Answer: A

Explanation:

When troubleshooting OSPF in a service provider environment, distinguishing between "stuck" adjacencies and "flapping" adjacencies is the first step. A session that transitions frequently between Full and Down indicates that the relationship can be established successfully (meaning parameters match), but it cannot be maintained.

According to Juniper Networks documentation, the most common cause for a session to drop from Full to Down is the expiration of the Dead Interval. If a router does not receive a Hello packet within the Dead Interval (usually 40 seconds), it tears down the adjacency. A physical interface flapping (Option A) is the primary trigger for this. If the physical link or the underlying transport (like a leased line or a microwave link) goes down even momentarily, the OSPF process immediately detects the interface failure, flushes the neighbors, and moves the state to Down. As soon as the interface comes back up, the routers perform the Hello exchange and reach the Full state again, creating the flapping cycle.

Analysis of other options:

MTU Mismatch (Option D): This typically causes the adjacency to get "stuck" in the Exchange or ExStart state. The routers can exchange small Hello packets, but when they try to send larger Database Description (DBD) packets that exceed the MTU, the packets are dropped, preventing the session from ever reaching "Full."

Area ID Mismatch (Option C): This prevents the adjacency from even reaching the Init state; the routers will never form a neighbor relationship.

Route Preference (Option B): This affects which route is chosen for the forwarding table but has no impact on the OSPF neighbor state machine itself.

5. Which feature allows Junos OS to perform recursive lookups for static route next hops?

- A. resolve
- B. discard
- C. reject
- D. next-table

Answer: A

Explanation:

In standard routing, a static route is typically considered valid only if the specified next-hop IP address is directly reachable on a local subnet. However, in complex service provider designs, the next-hop might be a "distant" IP address that is reachable through another route (such as a BGP route or another static route). This process of looking up a next-hop within another routing entry is called recursive lookup.

In Junos OS, the resolve (Option A) parameter is explicitly used to enable this behavior for static routes. According to Juniper technical documentation, when you append the resolve keyword to a static route configuration, you are instructing the Routing Engine to search the routing table to find a path to that distant next-hop.

For example:

set routing-options static route 10.1.1.0/24 next-hop 192.168.100.1 resolve

If 192.168.100.1 is not on a local interface but is reachable via an OSPF route, the router will "resolve" the path and install the 10.1.1.0/24 route into the forwarding table using the OSPF path's exit interface.

Why other options are incorrect:

Discard (Option B) and Reject (Option C) are "next-hop types" used to drop traffic, either silently (discard) or by sending an ICMP unreachable message (reject).

Next-table (Option D) is used for Inter-VRF routing, where the router is told to look up the destination in a completely different routing instance (like a VRF table), which is a different architectural function than a recursive next-hop lookup within the same table.