Juniper

Exam Questions JN0-364

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

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 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 dualRouting Engines (REs), the fundamental technology required to provide high availability during a hardware or software failure of the primary RE isGraceful 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 thechassis stateand theinterface 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 routercontinues to forward packetsthat 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 theforwardingstate, it does not by itself maintain therouting protocolstate (adjacencies). To keep OSPF or BGP sessions from dropping during the switchover, GRES must be paired withNon-Stop Active Routing (NSR). However, as the question focuses on the core requirement of continuing to forward packets,GRESis the foundational technology.

NEW QUESTION 2
You are asked to configure interfaces on Juniper devices to support dual VLAN tags. In this scenario, which two interface statements would accomplish this task? (Choose two.)

A. flexible-vlan-tagging
B. gigether-options
C. vlan-tagging
D. stacked-vlan-tagging

Answer: AD

Explanation:
To supportdual VLAN tagging(often referred to as Q-in-Q or 802.1ad), a Juniper interface must be configured to process more than one 802.1Q header. In Junos OS, this is handled at the physical interface level ([edit interfaces]).
According to Juniper Service Provider documents, two primary configuration statements enable this capability:
stacked-vlan-tagging (Option D):This is the traditional command used to enable an interface to accept frames with two VLAN tags. When this is enabled, the router expects an outer "service" tag and an inner "customer" tag. This is specifically used in provider edge scenarios where a service provider is tunneling multiple customer VLANs.
flexible-vlan-tagging (Option A):This is a more modern and versatile command. It allows the interface to support a mix of different encapsulation types across different logical units. For example, with flexible-vlan-tagging, you can have one logical unit (unit 10) doing standard single-tagging and another logical unit (unit 20) doing dual-tagging (vlan-tags outer X inner Y). This is the preferred method on newer hardware (like the MX Series) because it provides the highest level of configuration flexibility.
Vlan-tagging (Option C)only enables the interface to support a single 802.1Q tag, andgigether-options (Option B)contains physical-layer settings like auto-negotiation or flow control, which do not influence VLAN encapsulation. Therefore, A and D are the correct mechanisms for enabling dual-tag support.

NEW QUESTION 3
You are asked to configure a new network environment that will be based on IPv6 and use OSPF. In this scenario, which two statements correctly identify configuration task considerations? (Choose two.)

A. Participating interfaces must be configured with both IPv4 and IPv6 protocol families and addresses.
B. The router ID used must be based on a 128-bit identifier value.
C. The router ID used must be based on a 32-bit identifier value.
D. Participating interfaces are only required to be configured with the IPv6 protocol family and address.

Answer: CD

Explanation:
When transitioning to an IPv6 environment usingOSPFv3(the version of OSPF designed for IPv6), there are significant architectural differences compared to OSPFv2 (IPv4). According to Juniper Networks technical documentation, OSPFv3 was redesigned to be more protocol-agnostic.
Router ID (Option C):
Despite OSPFv3 routing IPv6 (which uses 128-bit addresses), the OSPFRouter IDremains a32-bit valueformatted like an IPv4 address (e.g., 1.1.1.1). This is a common point of confusion. In a pure IPv6 environment where no IPv4 addresses are configured on any interfaces, a Juniper router cannot automatically derive a Router ID. Therefore, the administrator must manually configure a 32-bit Router ID under [edit routing-options] for the OSPFv3 process to initialize.
Interface Configuration (Option D):
OSPFv3 runs directly over the IPv6 link-local scope. Unlike OSPFv2, it does not require an IPv4 address to function. Therefore, interfaces areonly required to be configured with family inet6(Option D). You do not need "dual-stack" (both IPv4 and IPv6) functionality just to run OSPFv3. The protocol uses the link-local address (fe80::/10) of the interface for neighbor adjacencies and as the next hop for routing updates. This separation allows OSPFv3 to carry multiple "address families" (both IPv4 and IPv6 unicast) if needed, but the base requirement for an IPv6-only network is simply the family inet6 configuration.

NEW QUESTION 4
Which two protocols would be used for dynamic routing in IPv6 environments? (Choose two.)

A. IGMP
B. IS-IS
C. OSPFv2
D. BGP

Answer: BD

Explanation:
The transition to IPv6 requires routing protocols that are capable of carrying 128-bit address information. Juniper Networks Junos OS supports several "IPv6-ready" protocols for dynamic routing.
* 1. IS-IS (Option B):
As discussed in previous questions,IS-ISis inherently extensible due to its use ofTLVs (Type, Length, Value). To support IPv6, the protocol did not need a major rewrite; instead, new TLVs (such as TLV 236 for IPv6 reachability and TLV 232 for IPv6 interface addresses) were added. A single IS-IS process in Junos can simultaneously carry both IPv4 and IPv6 routing information, making it a highly efficient choice for "dual-stack" service provider backbones.
* 2. BGP (Option D):
BGP was updated to support multiple protocols throughMultiprotocol Extensions (MP-BGP), defined in RFC 4760. By usingAddress Family Identifiers (AFI)andSubsequent Address Family Identifiers (SAFI), a single BGP session can exchange NLRI (Network Layer Reachability Information) for IPv4 unicast, IPv6 unicast, and even VPNv4/VPNv6 routes. In Junos, this is configured under the family inet6 unicast hierarchy within the BGP protocols configuration.
Why other options are incorrect:
IGMP (Option A):This is a management protocol for IPv4 multicast (Internet Group Management Protocol). Its IPv6 equivalent isMLD (Multicast Listener Discovery).
OSPFv2 (Option C):OSPF version 2 is strictly for IPv4. To run OSPF in an IPv6 environment,OSPFv3must be used, as it was specifically redesigned to handle the IPv6 address space and link-local communication.

NEW QUESTION 5
Exhibit:

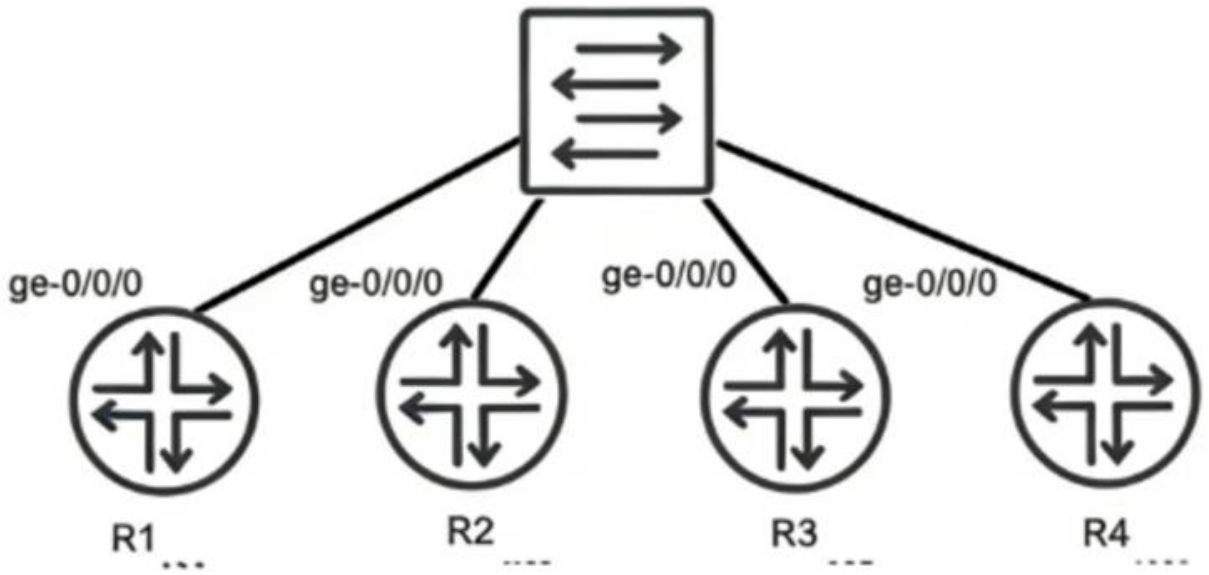


```
user@R1> show configuration routing-options
router-id 192.168.1.1;

user@R1> show configuration protocols ospf
area 0.0.0.0 {
    interface ge-0/0/0.0 {
        priority 200;
    }
}
```

```
user@R1> show configuration routing-options
router-id 192.168.1.3;

user@R1> show configuration protocols ospf
area 0.0.0.0 {
    interface ge-0/0/0.0 {
        priority 50;
    }
}
```

```
user@R1> show configuration routing-options
router-id 192.168.1.2;

user@R1> show configuration protocols ospf
area 0.0.0.0 {
    interface ge-0/0/0.0 {
        priority 100;
    }
}
```

```
user@R1> show configuration routing-options
router-id 192.168.1.4;

user@R1> show configuration protocols ospf
area 0.0.0.0 {
    interface ge-0/0/0.0 {
        priority 90;
    }
}
```

Referring to the exhibit, you have configured R1, R2, R3, and R4 to be a part of OSPF area 0 and you have connected them to a broadcast segment. Assuming all four routers come online within one minute of each other, which router becomes the DR and which router becomes the BDR?

A. R4 is the DR and R1 is the BDR
B. R1 is the DR and R4 is the BDR
C. R4 is the DR and R3 is the BDR
D. R1 is the DR and R2 is the BDR

Answer: D

Explanation:
In OSPF networks, when multiple routers are connected to a shared multi-access broadcast segment (like an Ethernet switch), they undergo an election process to select aDesignated Router (DR)and aBackup Designated Router (BDR). This mechanism is essential for reducing the number of adjacencies and limiting the volume of Link State Advertisement (LSA) flooding on the segment.
The OSPF election process follows a strict hierarchy based on the following criteria:
Interface Priority:The router with the highest OSPF interface priority is elected as the DR. The router with the second-highest priority becomes the BDR. In Junos, the default priority is 128, but it can be manually configured between 0 and 255.
Router ID:If there is a tie in priority, the router with the numerically highest Router ID (RID) wins the election.
Analyzing the configuration provided in the exhibit:
R1:Priority 200, Router-ID 192.168.1.1
R2:Priority 100, Router-ID 192.168.1.2
R3:Priority 50, Router-ID 192.168.1.3
R4:Priority 90, Router-ID 192.168.1.4
Comparing the priority values,R1 has the highest priority (200)and therefore becomes theDR. The next highest priority value among the remaining routers is100, which belongs to R2, making it theBDR. Although R4 has a higher Router ID than R2, the priority value is evaluated first and takes precedence.
Since all routers came online within a short window (one minute), they participate in the same election cycle, ensuring the configured priorities dictate the outcome rather than "first-come, first-served" preemption behavior common in OSPF once a DR is already established.

NEW QUESTION 6
Which statement about RSVP-signaled LSPs is correct?

A. CSPF is not required for LSPs using admin-groups.
B. CSPF is used to calculate the path for a traffic-engineered LSP.
C. The paths used by LSPs are always calculated using the SRGB.
D. The paths used by LSPs are always calculated using the TED.

Answer: B

Explanation:
In a Juniper Networks environment,Resource Reservation Protocol (RSVP)is a signaling protocol used to establish Label-Switched Paths (LSPs). While RSVP handles the actual signaling (requesting labels and reserving bandwidth along a path), it does not inherently know which path to take. This is where Constrained Shortest Path First (CSPF)comes into play.
CSPF is an advanced version of the Dijkstra algorithm used specifically for traffic engineering. Unlike the standard SPF used by IGPs, which only considers the shortest metric, CSPF takes into account multiple constraints such as available bandwidth, link coloring (administrative groups), and explicit hop requirements.
According to Juniper technical documentation, when an LSP is configured, the Ingress router uses CSPF to calculate a loop-free path that satisfies all these constraints before RSVP begins signaling. This is why statement B is the correct description of the operational flow.
Statement D is a common distractor. While CSPF uses the Traffic Engineering Database (TED)to perform its calculations, the path is not "calculated by the TED" itself; the TED is merely the repository of link-state information (provided by OSPF or IS-IS extensions). Statement C refers to Segment Routing Global Block (SRGB), which is relevant to Segment Routing (SR-TE), not standard RSVP-signaled LSPs. Finally, statement A is incorrect because admin-groups (link coloring) are actually one of the primary constraints that requireCSPF to determine a valid path.

NEW QUESTION 7
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, astatic routeis 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 calledrecursive lookup.
In Junos OS, theresolve (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)andReject (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 forInter-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.

NEW QUESTION 8
You are configuring BGP on a Juniper router to peer with an external provider. After committing the configuration, the BGP session remains in the Idle state. Which configuration issue would prevent the BGP session from progressing beyond the Idle state?

A. The peer IP address is unreachable.
B. The local AS number is higher than the peer??s AS number.
C. The peer is configured with a different router ID.
D. The BGP group type is set to internal instead of external.

Answer: A

Explanation:
In the BGP finite state machine, the Idle state is the "stop" or "start" point of the protocol. When a session is stuck in Idle, it means the BGP process is either administratively disabled or, more commonly, is unable to initiate the underlying TCP connection required for BGP.

According to Juniper Networks Service Provider documentation, the most common reason for a BGP session to remain in Idle is a lack of routing reach ability. For BGP to move to the Connect state, the Junos kernel must have a route to the IP address specified in the neighbor statement. If thepeer IP address is unreachable (Option A)—meaning there is no route in inet.0 (via OSPF, IS-IS, or static)—the router cannot initiate the TCP three-way handshake on port 179. Consequently, the state machine will never progress.
Analysis of incorrect options:
Option B:BGP does not care if the local AS is higher or lower than the peer's; it only cares if they match the configuration. AS numbers are identifiers, not priorities.
Option C:A mismatched Router ID does not prevent a session from leaving the Idle state. It would typically cause the session to reach the Open Confirm state, and then fail with a "Notification" message due to a collision or identification error.
Option D:While a mismatched group type(internal vs. external) will cause the session to fail, it usually fails during the Open message exchange (Open Sent state) because the AS numbers provided will not match the expected peer type (IBGP vs. EBGP).
Only the lack of a path to the neighbor (reach ability) keeps the session at the very beginning of the process: the Idle state.

NEW QUESTION 9
You are configuring LDP in a service provider network. After enabling LDP on core interfaces, you notice that labels are being advertised for every loopback IPv4 address that is in your IGP. Which label distribution mode is being used in this scenario?

A. conservative retention
B. ordered control
C. downstream unsolicited
D. downstream on demand

Answer: C

Explanation:
In the context of theLabel Distribution Protocol (LDP), the method by which a router advertises labels to its neighbors is defined by itsLabel Advertisement Mode. According to Juniper Networks documentation and industry standards (RFC 5036), there are two primary modes:Downstream Unsolicited (DU)andDownstream on Demand (DoD).
InDownstream Unsolicited (DU)mode, which is the default behavior for Junos OS and most service provider implementations, an LSR (Label Switching Router) does not wait for a specific request from its neighbors. Instead, as soon as the LSR learns a prefix through its Interior Gateway Protocol (IGP) and establishes an LDP session, it automatically generates a label for that prefix and advertises it to all of its LDP peers. This explains the scenario where labels appear for every loopback address in the IGP as soon as LDP is enabled. DU mode is highly efficient for fast convergence because the labels are already present in the neighbors' databases before they are even needed for traffic forwarding.
By contrast,Downstream on Demand (DoD)requires a router to explicitly request a label for a specific prefix from its next-hop neighbor.Ordered Control (Option B)andIndependent Controlrefer to thetimingof label creation (whether a router waits for the next-hop to provide a label before creating its own), whileConservative Retention (Option A)refers to how a router stores labels it receives but doesn't currently use for forwarding. In the Junos default environment, LDP utilizesDownstream Unsolicitedadvertisement combined withOrdered ControlandLiberal Retentionto ensure a robust and rapidly converging MPLS control plane.

NEW QUESTION 10
Exhibit:

```
user@R1> show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ge-0/0/0.0            R2               3  Up                25
ge-0/0/1.0            R6               2  Up                25
```

Referring to the exhibit, why is the ge-0/0/0.0 interface shown as belonging to Level 3?

A. This interface is configured as a point-to-point interface, that uses Level 3 as shorthand for both Level 1 and Level 2.
B. This interface is configured as a broadcast interface that has three adjacencies with other routers on the shared LAN.
C. This interface connects to a super spine.
D. This interface is configured as a broadcast interface, that uses Level 3 as shorthand for both Level 1 and Level 2.

Answer: A

Explanation:
In theIS-IS (Intermediate System to Intermediate System)protocol as implemented in Junos OS, the output of operational commands uses specific numerical representations to denote the hierarchy levels of a neighbor adjacency. Understanding these values is crucial for troubleshooting peering relationships in a multi-level IS-IS network.
According to Juniper Networks technical documentation, the show isis adjacency command displays the status of the neighbors. The "L" column indicates the level of the adjacency:
Level 1:Indicates the adjacency is strictly for intra-area routing.
Level 2:Indicates the adjacency is strictly for backbone/inter-area routing.
Level 3:This is ashorthand representationused by Junos to indicate that a single adjacency has been established forboth Level 1 and Level 2 simultaneously.
The critical distinction in this question lies in the interface type. On abroadcast interface(such as standard Ethernet), IS-IS typically establishes and maintains separate adjacencies for Level 1 and Level 2. In the CLI output for a broadcast link, you would generally see two separate lines for the same neighbor—one for Level 1 and one for Level 2.
However, on apoint-to-point (P2P)interface, IS-IS can negotiate both levels within a single adjacency. When this occurs, Junos consolidates the output into a single entry and usesLevel 3to signify that the adjacency is functional for both levels. Since the exhibit shows ge-0/0/0.0 as Level 3, it confirms that the link is configured with a point-to-point encapsulation (either natively or via the interface-type p2p command) and is acting as a Level 1/2 adjacency.
Option B is incorrect as the number "3" refers to protocol levels, not the count of neighbors. Option C is a reference to data center architectures that does not influence IS-IS level nomenclature. Option D is incorrect because, as noted, broadcast interfaces display these levels separately rather than using the Level 3 shorthand.

NEW QUESTION 10
A BGP router receives two routes to the same prefix. One route has a higher local preference, while the other has a shorter AS path. In this scenario, which route would be selected?

A. The route with the shorter AS path.
B. The route with the higher local preference.
C. The route with the lower origin code.
D. The route with the lowest MED value.

Answer: B

Explanation:
TheBGP path selection algorithmis a deterministic process used by Juniper routers to select the single "best" path from the BGP table to be placed into the routing table (inet.0). This algorithm follows a specific, hierarchical set of rules. According to Juniper Networks technical documentation, the router evaluates attributes in a fixed order, and once a tie is broken at a specific step, the remaining steps are ignored.
The order of the primary BGP attributes in Junos OS is as follows:
Highest Local Preference:This is the first attribute evaluated after the basic check for a reachable next hop. Local preference is used within an Autonomous System (AS) to prioritize one exit point over another.
Shortest AS_PATH:If the local preference is equal, the router then evaluates the length of the AS_PATH attribute.
Lowest Origin Code:(IGP < EGP < Incomplete).
Lowest Multi-Exit Discriminator (MED).
In this specific scenario, the router compares a path with ahigher local preferenceagainst a path with a shorter AS path. Because theLocal Preferencecheck occurs at Step 1 and theAS_PATHcheck occurs later at Step 2, the router will select the path with the higher local preference immediately. The length of the AS path becomes irrelevant in this comparison because the tie was already broken by the local preference value. This allows network administrators to override the default "shortest path" logic of BGP to prefer specific providers or links based on business requirements.

NEW QUESTION 11
Exhibit:

Exhibit

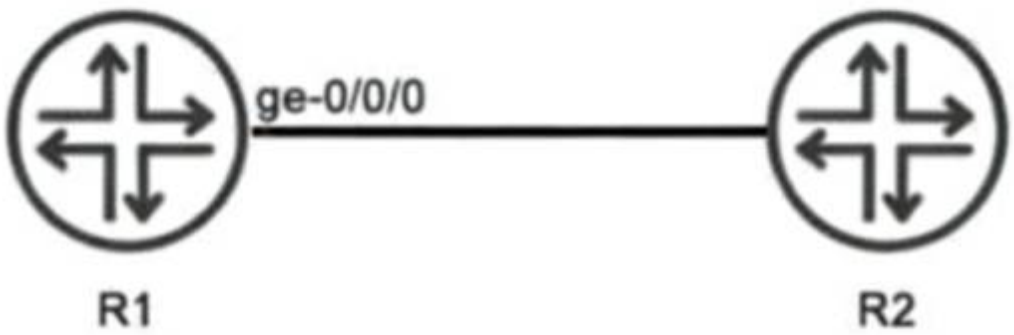


```
user@R1> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
ge-0/0/0.0            2   0x1 Disabled        Point to Point      100/100
lo0.0                 2   0x1 Passive         Passive               0/0
```

Referring to the exhibit, R1 and R2 are configured to run IS-IS. The IS-IS adjacency between R1 and R2 is up. What does the output of the show isis interface command tell you about R1?

A. R1 is not configured to use wide metrics.
B. R1 only forms a Level 2 adjacency with R2.
C. R1 advertises a Level 1 metric of 100 and a Level 2 metric of 100 toward R2 in its link-state PDU.
D. R1 sends Level 1 hello PDUs to R2.

Answer: B

Explanation:
In theIS-IS (Intermediate System to Intermediate System)protocol as implemented in Junos OS, routers can operate at two hierarchical levels:Level 1 (L1)for intra-area routing andLevel 2 (L2)for inter-area backbone routing. By default, a Juniper router and its interfaces are configured to act asLevel 1/2, meaning they will attempt to form adjacencies at both levels simultaneously.
According to Juniper Networks technical documentation, the show isis interface command provides a granular view of how the protocol is interacting with specific local links. In the provided exhibit, we must examine the L (Level)column and theDR (Designated Router)status columns to understand R1's operational state.
Level Configuration:Under the L column for both the physical interface ge-0/0/0.0 and the loopback lo0.0, the value is strictly2. This indicates that these interfaces have been explicitly configured to operate only at Level 2.
Adjacency Capabilities:For the interface ge-0/0/0.0, the Level 1 DR field is marked asDisabled. This confirms that R1 is not participating in Level 1 operations on this link; it will not transmit Level 1 Hello PDUs, nor will it listen for them. Consequently, R1 is incapable of forming a Level 1 adjacency with R2 on this segment.
Metric Implications:The exhibit shows an L1/L2 Metric of100/100. In Junos, "narrow" metrics (the default) are limited to a maximum value of 63 per interface. A metric of 100 indicates thatwide metrics (wide-metrics-only) have been enabled. Therefore, option A is incorrect because the routerisusing wide metrics.
Since the prompt states the adjacency is "up," and the interface is restricted to Level 2, we can conclude that R1 only forms a Level 2 adjacency with R2 (Option B). Even though an L1 metric of 100 is displayed in the table as a configured value, it is not actually "advertised" in a Link-State PDU because the Level 1 protocol is disabled on that interface.

NEW QUESTION 12
Which IS-IS packet type will establish and maintain neighbor relationships?

A. link-state PDU
B. hello PDU
C. partial sequence number PDU
D. update PDU

Answer: B

Explanation:
In theIS-IS (Intermediate System to Intermediate System)protocol, communication between routers is performed usingProtocol Data Units (PDUs). To discover neighbors and maintain adjacencies, IS-IS relies on theHello PDU (IIH - IS-IS Hello).
According to Juniper Networks technical documentation, when IS-IS is enabled on an interface, the router begins transmitting Hello PDUs to a multi-destination address (multicast). These PDUs contain essential information such as the router'sSystem ID, its configuredArea Addresses, and itsLevel capability(Level 1, Level 2, or both). For two routers to become neighbors, they must exchange these Hello PDUs and agree on specific parameters, such as the MTU of the link and the hello/hold timers.
Once an adjacency is established, the Hello PDU serves as a "keepalive" mechanism. If a router stops receiving Hello PDUs from a neighbor for a duration exceeding theHolding Time, it assumes the neighbor is down and flushes the associated Link-State PDUs (LSPs) from its database.
To clarify the other options:
Link-State PDU (Option A):These are used to distribute actual topology and reachability information, not to form adjacencies.
Partial Sequence Number PDU (Option C):PSNPs are used on point-to-point links to acknowledge the receipt of LSPs or to request missing LSPs.
Update PDU (Option D):This is not a standard IS-IS term; in IS-IS, updates are handled via the flooding of LSPs.

NEW QUESTION 17
Exhibit:

```
user@R1> show configuration protocols mpls
label-switched-path to-r3 {
    to 192.168.100.3;
}
interface ge-0/0/0.0;
user@R1> show configuration protocols ospf
area 0.0.0.0 {
    interface ge-0/0/0.0;
    interface lo0.0;
}
user@R1> show route 192.168.100.3
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
192.168.100.3/32    *[OSPF/10] 00:05:39, metric 2
                     > to 172.16.1.2 via ge-0/0/0.0
user@R1> show mpls lsp detail
Ingress LSP: 1 sessions
192.168.100.3
From: 192.168.100.1, State: Dn, ActiveRoute: 0, LSPname: to-r3
ActivePath: (none)
LSPtype: Static Configured, Penultimate hop popping
LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                    State: Dn
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Will be enqueued for recomputation in 27 second(s).
   17 Sep 14 20:29:00.840 CSPF: could not determine self
Total 1 displayed, Up 0, Down 1
Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
Transit LSP: 0 sessions
```

```
    Total 0 displayed, Up 0, Down 0
    user@R1> show configuration interfaces
    ge-0/0/0 {
     unit 0 {
         family inet {
             address 172.16.1.1/24;
         }
         family mpls;
     }
}
fxp0 {
    unit 0 {
        family inet {
            address 10.0.1.11/24;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
```

```
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 10.0.1.11/24;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.100.1/32;
            }
        }
    }
```

You have configured an MPLS LSP to 192.168.100.3. However, the LSP is in the down state. Referring to the exhibit, which two actions would solve this problem? (Choose two.)

A. Issue the set routing-options rib inet.3 static route 192.168.100.1 command and commit.
B. Issue the set protocols mpls label-switched-path to-r3 no-cspf command and commit.
C. Issue the set interfaces lo0 family mpls command on router R1 and commit.
D. Issue the set protocols ospf traffic-engineering command and commit.

Answer: BD

Explanation:
In a Juniper Networks environment, establishing a functionalMultiprotocol Label Switching (MPLS)Label- Switched Path (LSP) requires synchronized control plane operations. According to Juniper technical documentation, the most common reason for an LSP to remain in the "Down" state at the ingress router is a failure of theConstrained Shortest Path First (CSPF)algorithm during the path computation phase.
The provided exhibit for routerR1reveals a critical error in the show mpls lsp detail output: "CSPF: could not determine self". This specific error indicates that the CSPF process is unable to find its own local router ID within theTraffic Engineering Database (TED). For CSPF to build a valid TED, the underlying Interior Gateway Protocol (IGP), such as OSPF, must be configured to flood opaque link-state advertisements (Type 10 LSAs) that carry traffic engineering attributes. As seen in the OSPF configuration, traffic engineering is not enabled. Therefore, issuing theset protocols ospf traffic-engineeringcommand (Option D) will allow R1 to populate the TED with its own local information and that of its neighbors, enabling CSPF to calculate a valid path.
Alternatively, an administrator can choose to bypass the requirement for a TED entirely by disabling CSPF on the specific LSP. By issuing theset protocols mpls label-switched-path to-r3 no-cspfcommand (Option B), the router will stop attempting to perform a constrained path calculation. Instead, the signaling protocol (RSVP) will rely on the standardinet.0routing table to determine the hop-by-hop path to the egress destination (192.168.100.3), allowing the LSP to establish without traffic engineering constraints.
Regarding the other options, whilefamily mplsis required on all transit interfaces, the ingress loopback interface (lo0) generally does not require it for standard LSP signaling unless it's used as a transit hop. Furthermore, adding a static route toinet.3(Option A) is used for next-hop resolution of BGP routes over LSPs but does not assist in the signaling or establishment of the LSP itself.

NEW QUESTION 18

Exhibit:

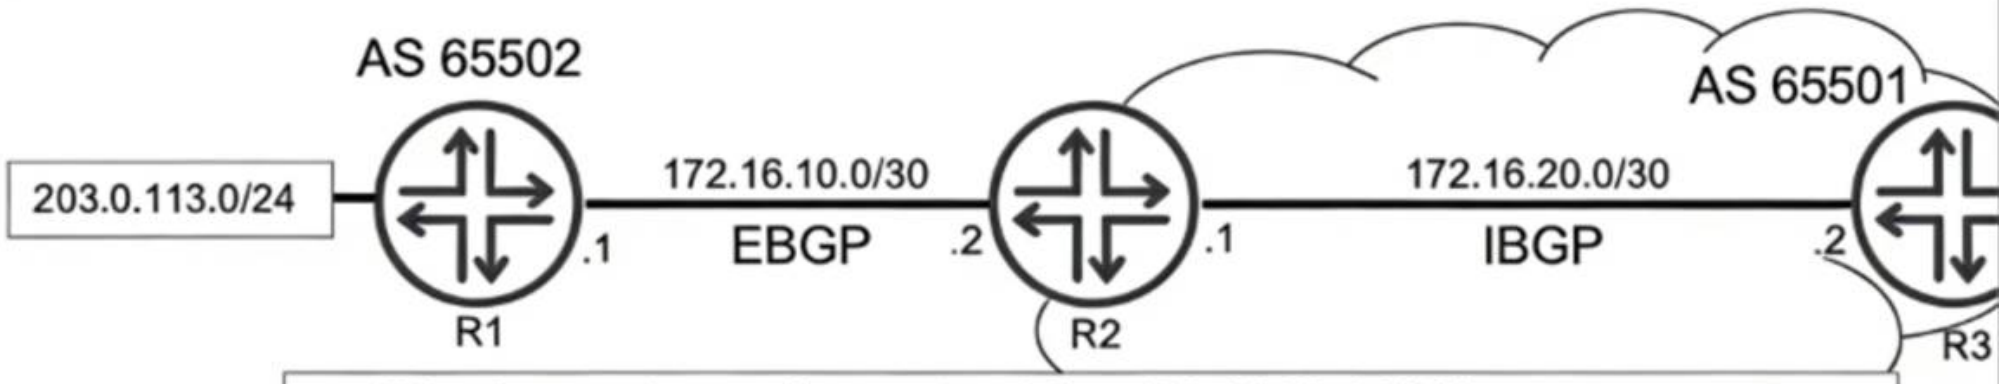


```
user@R3> show route receive-protocol bgp 172.16.20.1 hidden

inet.0: 9 dentinations, 9 routes (8 active, 0 hoiddown, 1 hidden)
  Prefix                  Nexthop              MED     Lolpref    AS path
  203.0.113.0/24          172.16.10.1                     100        65502 1
```

```
user@R2> show configuration protocols bgp
group EBGP {
    type external;
    neighbor 172.16.10.1 {
        peer-as 65502;
    }
}
group IBGP {
    type internal;
    export export-to-ibgp;
    neighbor 172.16.20.2 {
        peer-as 65501;
    }
}

user@R2> show configuration policy-options policy-statement export-to-ibgp
```

Referring to the exhibit, R1 is advertising prefix 203.0.113.0/24 to R2 over EBGP. R2 is configured to advertise this prefix into IBGP. R3 receives the 203.0.113.0/24 route, however the route is hidden. Which configuration statement do you need to add to R2 to solve this problem?

A. set policy-options policy-statement export-to-ibgp from route-filter 203.0.113.0/24 orlonger
B. set policy-options policy-statement export-to-ibgp then next-hop self
C. set protocols bgp group EBGP export export-to-ibgp
D. set policy-options policy-statement export-to-ibgp then local-preference 50

Answer: B

Explanation:
In Juniper Networks Junos OS, a "hidden" route in the BGP table typically signifies that the router has received the prefix but cannot install it into the active routing table because theBGP next hop is unreachable. This is a common occurrence in service provider environments when transitioning betweenExternal BGP (EBGP)andInternal BGP (IBGP).
According to Juniper technical documentation, when an EBGP speaker (R1) advertises a prefix to its peer (R2), it sets the next hop to its own interface IP address ($172.16.10.1$). By default, when R2 re-advertises that prefix to its IBGP peer (R3), itpreserves the original EBGP next-hop address. Unless R3 has a specific route in its Interior Gateway Protocol (IGP) or a static route to reach the $172.16.10.1$ subnet, it will mark the route as unusable (hidden).
In the exhibit, the show route output onR3explicitly shows the nexthop for $203.0.113.0/24$ as $172.16.10.1$. Since this route is marked "hidden," we can conclude R3 does not know how to reach R2's external peering link. To resolve this, the network administrator must modify the next-hop attribute before the route is sent to R3.
By adding the statementset policy-options policy-statement export-to-ibgp then next-hop self(Option B) on routerR2, R2 will replace the external next-hop ($172.16.10.1$) with its own internal peering address ($172.16.20.1$) before advertising the route to R3. Because R3 already has a direct or IGP connection to R2's internal address, it will successfully resolve the next hop, and the route will transition from "hidden" to "active."
Option A is unnecessary because the route is already being exported; Option C is redundant as the policy is already applied to the IBGP group; and Option D changes path preference but does not solve the underlying reachability problem.

NEW QUESTION 21
Which IS-IS adjacency state indicates that hello packets have been exchanged but the adjacency is not yet fully established?

A. loading
B. initializing
C. up
D. two-way

Answer: B

Explanation:
In the IS-IS (Intermediate System to Intermediate System)protocol, the process of forming an adjacency between two neighbors follows a specific sequence of states. While OSPF uses states like "Init," "Two-Way," and "Full," IS-IS uses a slightly different nomenclature within its state machine.
According to Juniper Networks technical documentation, when a router first sends anIS-IS Hello (IIH) PDU and receives one back from a neighbor, but has not yet confirmed that the neighbor "sees" it back, the adjacency enters the Initializing state. Specifically, on a point-to-point link, the state transitions from Downto Initializing as soon as the first PDU is received. On a broadcast network (like Ethernet), the Initializing state indicates that the local router has received a Hello PDU

from the neighbor, but the local router's own System ID is not yet listed in the neighbor's list of "seen" neighbors (the neighbor's Hello PDU does not yet contain the local router's MAC address).
The adjacency only moves to the Upstate (Option C) once bi-directional communication is confirmed— meaning both routers have seen each other's System IDs in the incoming Hello PDUs.
Why other options are incorrect:
Loading (Option A):This is an OSPF state, not an IS-IS state. In IS-IS, database synchronization happens after the adjacency is Up.
Two-Way (Option D):While functionally similar to the state IS-IS is achieving, "Two-Way" is the specific terminology for OSPF. In IS-IS, the intermediate step between knowing a neighbor exists and having a fully functional adjacency is strictly calledInitializing.

NEW QUESTION 23
How are routing loops prevented in external BGP networks?

A. By default, a router receiving a route with its own AS in the AS Path attribute will use the route.
B. Routing policies must be used to drop looped routes.
C. Routing policies must be used to accept valid routes.
D. By default, a router receiving a route with its own AS in the AS Path attribute will not use the route.

Answer: D

Explanation:
BGP is apath-vector protocol, and its primary mechanism for ensuring a loop-free topology across the global internet is theAS_PATHattribute. This attribute is a "well-known mandatory" attribute that records every Autonomous System (AS) a prefix has passed through.
According to Juniper Networks Service Provider documentation, the loop prevention rule forExternal BGP (EBGP)is straightforward: when a router receives a BGP Update from an EBGP peer, it examines the AS_PATH list. If the router's ownlocal AS numberis already present in the list, it indicates that the advertisement has already traversed the local AS and has returned. To prevent a routing loop, the routerwill not use the routeand will implicitly discard the update (Option D).
This behavior is a default, hard-coded function of the BGP protocol and does not require the administrator to write manualrouting policies(Options B and C) to achieve basic loop prevention. While there are advanced features like as-path-expand or allow-as-in that can modify this behavior for specific design requirements (such as in certain Hub-and-Spoke MPLS VPN topologies), the standard operational default is to reject any route where the local AS is detected in the path. This ensures that traffic does not circulate infinitely between Autonomous Systems.

NEW QUESTION 26
Exhibit:

Exhibit

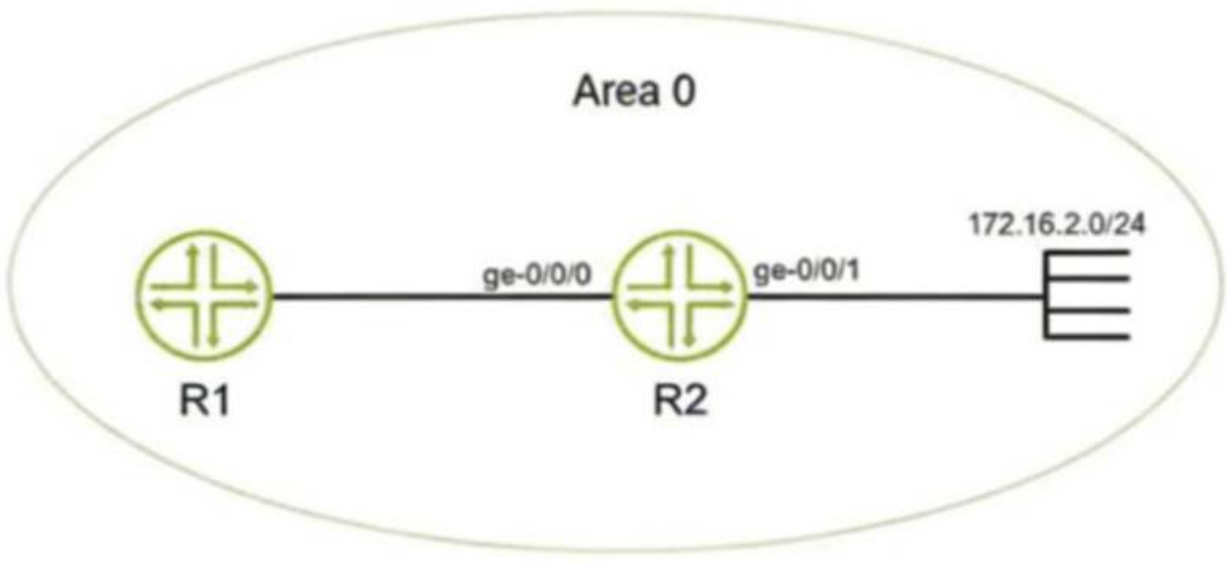


```
[edit]
user@R2# show protocols
ospf {
    area 0.0.0.0 {
        interface ge-0/0/0.0;
        interface lo0.0;
        interface ge-0/0/1.0;
    }
}
ospf3 {
    realm ipv4-unicast {
        area 0.0.0.0 {
            interface ge-0/0/0.0;
            interface ge-0/0/1.0;
            interface lo0.0;
        }
    }
    area 0.0.0.0 {
        interface ge-0/0/0.0;
        interface ge-0/0/1.0;
        interface lo0.0;
    }
}
```

You have configured IPv4 and IPv6 in your network and all OSPF neighbors are established. You apply the configuration shown in the exhibit. Which statement is true in this scenario?

A. There will only be an OSPFv2 entry in R1 for network 172.16.2.0/24.
B. There will be an OSPFv2 and OSPFv3 entry in R1 for network 172.16.2.0/24.
C. There will not be a route in R1 for network 172.16.2.0/24.
D. There will only be an OSPFv3 entry in R1 for network 172.16.2.0/24.

Answer: B

Explanation:
In a Juniper Networks environment running Junos OS, understanding the interaction between different versions of OSPF is essential for multi-protocol environments.OSPFv2(defined in RFC 2328) is the standard protocol used for routing IPv4 unicast traffic.OSPFv3(defined in RFC 5340) was originally developed to support IPv6 routing. However, OSPFv3 was later extended via RFC 5838 to support multiple address families (AF), allowing it to carry IPv4 unicast, IPv4 multicast, and other address types within a single OSPF instance.
According to Juniper technical documentation, Junos OS implements this multi-AF support in OSPFv3 through the use ofrealms. When the realm ipv4-unicast statement is configured under the [edit protocols ospf3] hierarchy, the OSPFv3 process becomes capable of calculating and advertising IPv4 routes.
In the provided exhibit, routerR2has a dual-protocol configuration. First, it is running standard OSPFv2, with the ge-0/0/1.0 interface (which is directly connected to the 172.16.2.0/24 network) participating in Area 0. This ensures that the prefix is advertised as a standard IPv4 LSA to its neighbor,R1. Second, R2 is running OSPFv3 with the realm ipv4-unicast specifically enabled on that same ge-0/0/1.0 interface. Because of this realm, OSPFv3 also treats the 172.16.2.0/24 prefix as a reachable IPv4 destination and advertises it to R1 as an OSPFv3 IPv4-unicast LSA.
As a result, whenR1(which is also running both protocols) receives these routing updates, it will see the same destination prefix advertised by two different protocols. Its routing table (inet.0) will contain one entry learned from the OSPFv2 process and a second, separate entry learned from the OSPFv3 process. While the Junos Routing Engine will ultimately select one as the "active" route based on route preference (both protocols have a default preference of 10), both entries will technically exist within the Routing Information Base (RIB). This confirms that statementBis the correct description of the operational state of the network.
=========

NEW QUESTION 31
You are evaluating BGP between two Juniper routers and the BGP session is stuck in the Idle state. What would cause this behavior?

A. The BGP hold time is too short.
B. The BGP group type is set to internal instead of external.
C. The local AS number is missing.
D. The peer IP address is incorrect.

Answer: D

Explanation:
In the BGP Finite State Machine (FSM), the Idle state is the first stage of any BGP connection. When a BGP session is "stuck" in Idle, it typically indicates that the router is unable to even begin the process of establishing a TCP connection with its neighbor. According to Juniper Networks documentation, before BGP can transition to the Connector Active states, it must have a valid route to the neighbor's IP address in the routing table and be able to initiate a three-way TCP handshake on port 179.
If the peer IP address is incorrect(Option D), the router may not have a route to that destination, or it may be attempting to connect to a non-existent or unreachable host. In many Junos configurations, if the underlying IGP (OSPF/IS-IS) or static routing cannot provide reach ability to the neighbor address defined in the BGP configuration, the BGP process will remain in the Idle state and periodically retry the connection.
Regarding the other options:
The local AS number is missing (Option C):In Junos, you cannot commit a BGP configuration if the local autonomous system is not defined at either the [edit routing-options] level or within the BGP group itself. The commit check would fail before the session could even attempt to start.
The BGP group type (Option B):Having a mismatch in group type (internal vs. external) usually results in the session reaching the Open Sentor Open Confirmstate before failing due to an "unacceptable AS" error in the OPEN message.
BGP hold time (Option A):Issues with hold timers or keep alives generally cause a session that is already in the Established state to drop; they do not prevent the session from leaving the Idle state.

NEW QUESTION 33
Exhibit:

```
user@switch1> show spanning-tree interface
Spanning tree interface parameters for instance 0
Interface    Port ID    Designated      Designated         Port   State  Role
                        port ID         bridge ID          Cost
ge-0/0/6.0   128:519      128:519   32768.0019e2552481     20000  FWD    DESG
ge-0/0/7.0   64:520       64:520    32768.0019e2552481     20000  FWD    DESG
ge-0/0/8.0   32:521       32:521    32768.0019e2552481     20000  FWD    DESG
ge-0/0/9.0   32:522       32:522    32768.0019e2552481     20000  FWD    DESG
ge-0/0/11.0  32:524       32:524    32768.0019e2552481     20000  FWD    DESG
ge-0/0/12.0  64:525       64:525    32768.0019e2552481     20000  FWD    DESG
ge-0/0/13.0  64:526       64:526    32768.0019e2552481     20000  FWD    DESG
```

Referring to the exhibit, which two statements are correct? (Choose two.)

A. The switch1 device is using VSTP.
B. The switch1 device is the root bridge.
C. The ge-0/0/8, ge-0/0/9, and ge-0/0/11 interfaces are using the default interface priority.
D. The bridge priority for switch1 is 32k.

Answer: BD

Explanation:
In the provided exhibit, the output of the command show spanning-tree interface for switch1 reveals critical details about the Spanning Tree Protocol (STP) operational state.
The first correct statement is thatthe switch1 device is the root bridge(Option B). This is determined by comparing the "Port ID" column with the "Designated port ID" column, as well as checking the "Designated bridge ID". In the exhibit, for every interface listed (from ge-0/0/6.0 to ge-0/0/13.0), the Port ID and the Designated port ID are identical. Furthermore, every port is in the "FWD" (Forwarding) state with the "DESG" (Designated) role. In a Spanning Tree topology, the root bridge is

the only device where all active participating interfaces serve as designated ports, as it has no need for a "Root" port role (which points toward a root bridge).
The second correct statement is thatthe bridge priority for switch1 is 32k(Option D). Looking at the "Designated bridge ID" column, we see the value 32768.0019e2552481. In Junos and general networking standards, the Bridge ID is composed of a bridge priority and the device's MAC address. The default priority for most Spanning Tree variants (STP, RSTP, MSTP) is 32,768, which is commonly referred to in shorthand as "32k".
Regarding the incorrect options:
Option A:There is no evidence of VSTP (VLAN Spanning Tree Protocol); the output shows "instance 0," which is typical for IEEE standard RSTP or STP.
Option C:The Port IDs for ge-0/0/8, ge-0/0/9, and ge-0/0/11 all start with "32" (e.g., 32:521), whereas the default port priority is typically 128 (as seen in ge-0/0/6.0 with 128:519). This indicates that the interface priorities for these specific ports have been manually tuned to a non-default value.

NEW QUESTION 37
What are three default BGP advertisement rules? (Choose three.)

A. EBGP peers advertise routes learned from IBGP or EBGP peers to other EBGP peers.
B. IBGP peers advertise routes received from EBGP peers to other IBGP peers.
C. IBGP peers advertise routes received from IBGP peers to other IBGP peers.
D. IBGP peers do not advertise routes received from IBGP peers to other IBGP peers.
E. IBGP peers do not advertise routes received from EBGP peers to other IBGP peers.

Answer: ABD

Explanation:
TheBorder Gateway Protocol (BGP)operates based on a strict set of advertisement rules designed to prevent routing loops while ensuring global reachability. These rules differ significantly depending on whether the relationship isExternal BGP (EBGP)orInternal BGP (IBGP).
* 1. EBGP Advertisement (Option A):In a standard EBGP scenario, a router acts as an exit/entry point for an Autonomous System. When an EBGP speaker receives a valid route from any peer (Internal or External), it will, by default, advertise that route to all of its other EBGP peers. This is the primary mechanism that allows prefixes to propagate across the global internet from one AS to another.
* 2. IBGP Split Horizon (Option D):
The most critical rule within an AS is theIBGP Split Horizonrule. To prevent loops within an AS, BGP dictates that a route learned from an IBGP peermust notbe advertised to any other IBGP peer. This is why BGP requires a "full mesh" of IBGP sessions or the use ofRoute Reflectorsto ensure all internal routers learn all routes. Without this rule, a route could circulate infinitely within the AS because IBGP does not update the AS_PATH attribute.
* 3. EBGP to IBGP Propagation (Option B):
When a router learns a route from an EBGP peer, it is permitted to advertise that route to all of its IBGP peers. This ensures that everyone inside the network knows how to reach external destinations. However, it is important to remember that in Junos OS, the BGP Next Hopis not modified by default when sending routes to IBGP peers, often requiring a "next-hop-self" policy to ensure internal reach ability.
Options C and E are incorrect because they directly contradict these fundamental BGP loop-prevention and propagation mechanisms.

NEW QUESTION 39
Exhibit:

```
user@Router-1> show route 172.24/16

inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

...

172.24.0.0/24 *[OSPF/150] 01:31:31, metric 0, tag 0

> to 172.20.0.2 via ge-0/0/2.0

to 172.20.1.2 via ge-0/0/3.0

user@Router-1> show route forwarding-table

Routing table: default.inet

Internet:

Destination Type RtRef Next hop Type Index NhRef Netif

...

172.24.0.0/24 user 0

172.20.0.2 ucst 551 2 ge-0/0/2.0

172.20.1.2 ucst 552 2 ge-0/0/3.0
```

Referring to the exhibit, which two statements are true? (Choose two.)

A. The router is performing default route load-balancing behavior.
B. The default route load-balancing behavior of this router has been modified.
C. This router will only choose the next hop with a > next to it in the routing table.
D. This router will choose both next hops in the routing table.

Answer: BD

Explanation:
In Junos OS, understanding the distinction between theRouting Information Base (RIB)and theForwarding Information Base (FIB)is fundamental to analyzing traffic patterns and load-balancing behavior. The RIB (show route) contains all prefixes learned via various protocols, while the FIB (show route forwarding-table) contains only the active next-hops that are actually programmed into the Packet Forwarding Engine (PFE).
According to Juniper Networks technical documentation, the default behavior for Junos OS when encounteringEqual-Cost Multipath (ECMP)routes is to select only a single next-hop from the available candidates in the RIB and install that single path into the FIB. In a default state, even if the show route output displays multiple next-hops for a destination like 172.24.0.0/24, only one would have the active route symbol (>) and only that one would appear in the forwarding table.
In the provided exhibit, the show route output shows two next-hops for 172.24.0.0/24, but only the first one (172.20.0.2) is marked with the>symbol as the active selection. However, the subsequent show route forwarding-table output reveals thatboth next-hops(172.20.0.2 and 172.20.1.2) are currently present in the forwarding table for that same destination. This discrepancy indicates that thedefault load-balancing behavior has been modified (Option B). This modification is typically achieved by creating a routing policy with the action then load-balance per-packet (which actually results in flow-based load balancing) and applying it to

the forwarding table via the export statement under [edit routing-options forwarding-table].
Because the forwarding table now contains both next-hops, the router is no longer restricted to a single path. Therefore, therouter will choose both next-hops in the routing table (Option D)for packet forwarding, distributing flows across the two available Gigabit Ethernet interfaces (ge-0/0/2.0 and ge-0/0/3.0). This ensures higher utilized bandwidth and provides redundancy at the data plane level.

NEW QUESTION 44
You are designing a high availability solution for a Juniper router with dual Routing Engines (RE). You want to ensure that the routing protocol state is preserved during an RE switchover. You have already enabled graceful Routing Engine switchover (GRES) and you want to avoid relying on helper routers to maintain the routing protocol state. In this scenario, which feature would accomplish this behavior?

A. non-stop active bridging
B. bidirectional forwarding detection
C. graceful restart
D. non-stop active routing

Answer: D

Explanation:
When designing High Availability (HA) for Juniper Service Provider routers, understanding the interaction between the control plane and data plane is vital. The user has already enabledGraceful Routing Engine Switchover (GRES), which synchronizes the interface and kernel state between the primary and backup Routing Engines (REs). However, GRES by itself does not preserve the routing protocol state (like OSPF adjacencies or BGP sessions).
To achieve the preservation of the routing protocol state without relying on external "helper" routers, you must implementNon-Stop Active Routing (NSR). According to Juniper Networks documentation, NSR uses the infrastructure provided by GRES to also synchronize the routing protocol process (rpd) information. Under NSR, the backup RE maintains a "hot" standby state of all routing protocols. If the primary RE fails, the backup RE takes over immediately. Because it already possesses the full routing table and peer session states, the peering neighbors are unaware that a switchover occurred. No protocol adjacency resets occur, and traffic continues to flow uninterrupted.
It is crucial to differentiate NSR fromGraceful Restart (Option C). While Graceful Restart also aims to maintain traffic flow during a switchover, itdoesrequire help from neighboring routers (known as "helper mode"). If the neighbors do not support or are not configured for Graceful Restart, the sessions will drop. Since the user explicitly stated they want to "avoid relying on helper routers," Graceful Restart is not the correct solution.
Non-stop Active Bridging (Option A)provides a similar "hitless" failover but specifically for Layer 2 environments (STP/VLANs) rather than Layer 3 routing protocols.BFD (Option B)is a failure detection protocol used to speed up convergence but does not preserve state during an RE failover; in fact, without NSR, BFD would likely trigger a faster teardown of the session during a switchover. Therefore,NSRis the only feature that meets the requirement for independent control-plane preservation.

NEW QUESTION 47
A service provider is onboarding a new enterprise customer that operates multiple branch offices, each with its own set of VLANs. The customer requires transparent Layer 2 connectivity between sites while maintaining separation of internal VLANs. The provider must also ensure that customer VLAN identifiers do not conflict with other customers on the shared infrastructure. Which solution would provide the desired results?

A. Extend customer VLANs using Q-in-Q tunneling.
B. Deliver Layer 3 VPN services using MPLS.
C. Aggregate customer traffic using GRE tunnels.
D. Provide Internet access with NAT and firewall services.

Answer: A

Explanation:
In a service provider environment,Q-in-Q tunneling(also known as 802.1ad or double-tagging) is the standard solution for transporting multiple customer VLANs over a shared provider backbone while maintaining total separation.
According to Juniper Networks documentation, Q-in-Q works by adding a second 802.1Q tag (theService Provider tagor S-tag) to the customer??s already tagged frames (theCustomer tagor C-tag). This creates a "tunnel" at Layer 2. This solution specifically addresses all the customer's requirements:
Transparent Layer 2 Connectivity:Because the provider simply encapsulates the customer's frames, the customer's internal BPDU traffic (like Spanning Tree) and VLAN tags are preserved and delivered transparently to the remote site.
Separation of Internal VLANs:The customer can run their own internal VLAN IDs (1-4094) without the provider needing to know or manage them.
Conflict Avoidance:Different customers on the same provider infrastructure are assigned unique S-tags. Even if two different customers both use "VLAN 10" internally, they remain isolated because their traffic is encapsulated in different provider S-tags.
Why other options are incorrect:
Layer 3 VPN (Option B):While MPLS L3VPNs are common, they provide Layer 3 (IP) connectivity, not the "transparent Layer 2" connectivity requested.
GRE Tunnels (Option C):GRE is a Layer 3 encapsulation and does not natively provide the transparent VLAN bridging required for a multi-site Layer 2 service.
NAT/Firewall (Option D):These are security and address-translation services for internet access and do not facilitate site-to-site Layer 2 bridging.

NEW QUESTION 51
Exhibit:

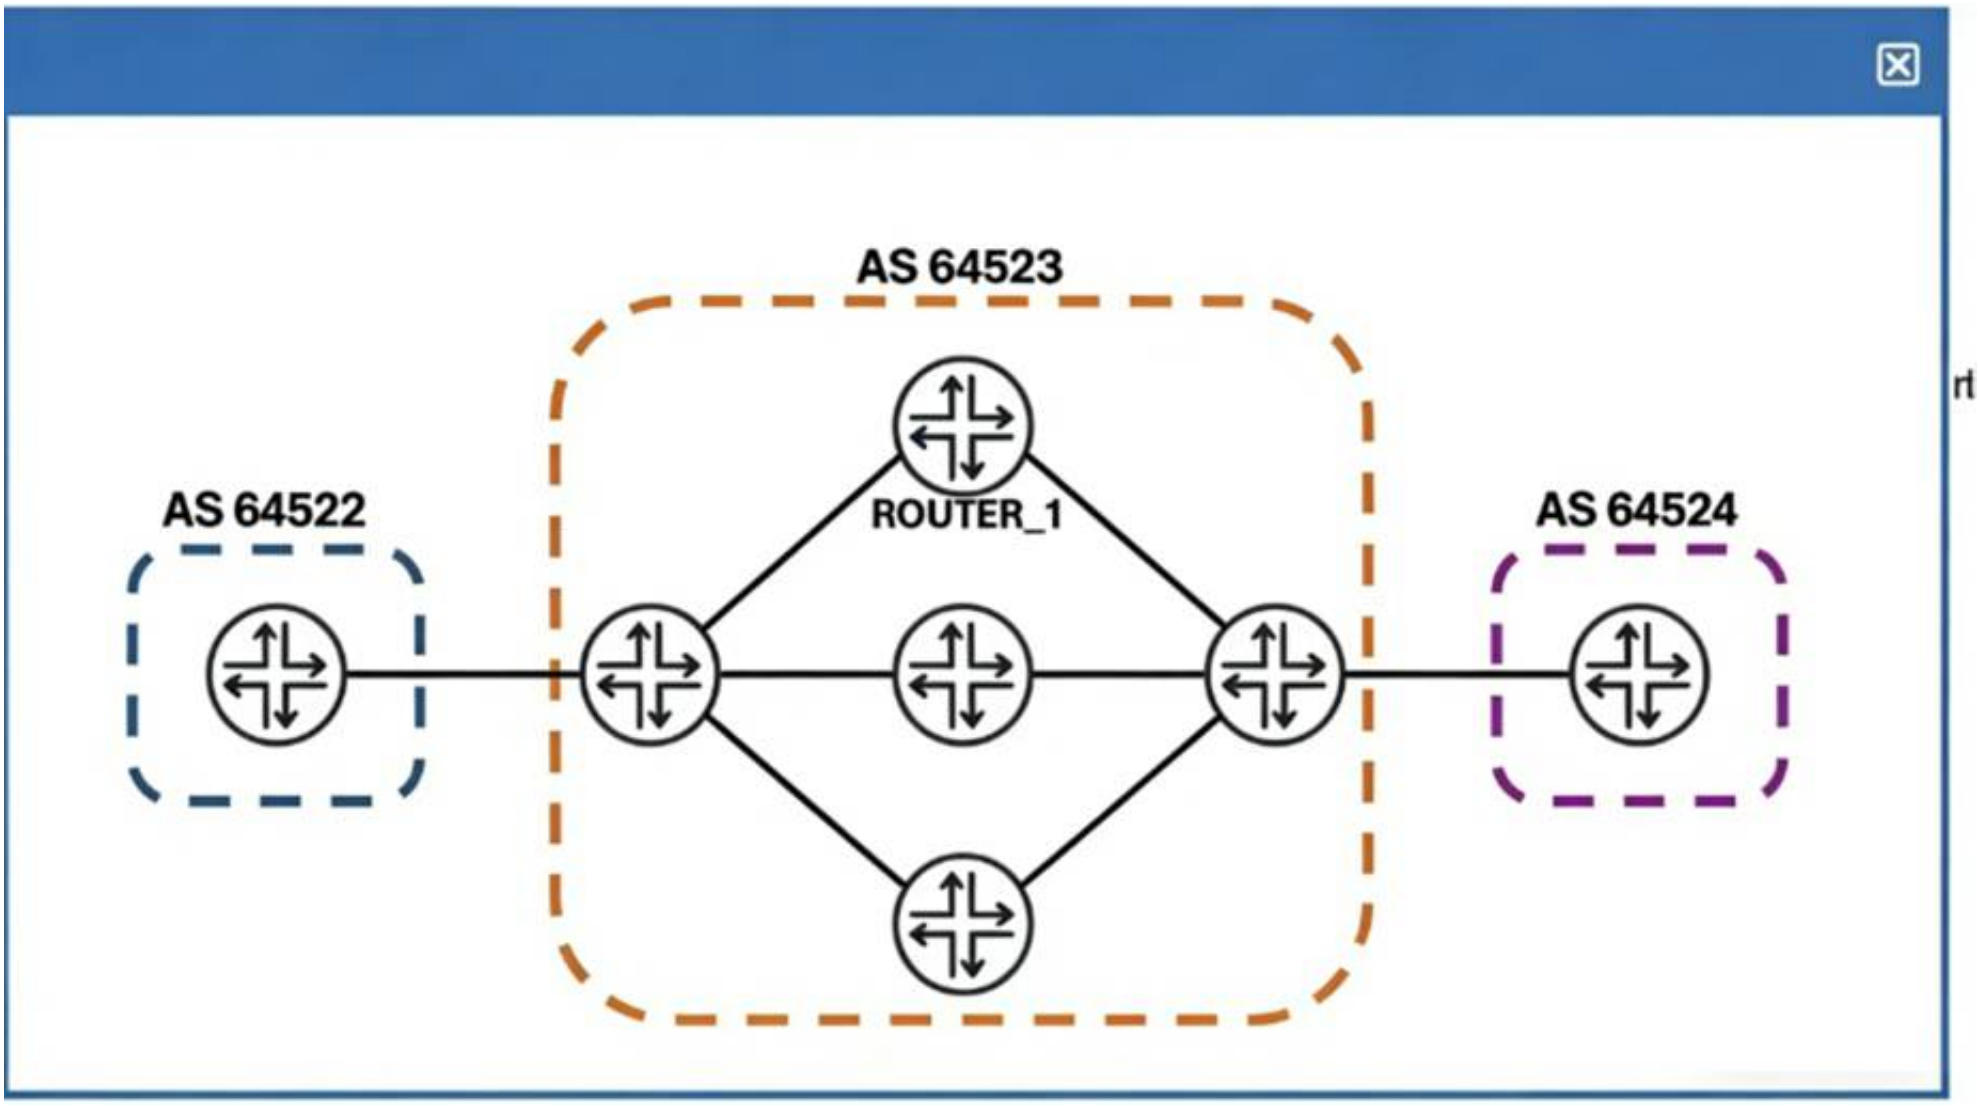

You must configure the router called ROUTER_1 to take all valid prefixes learned from internal BGP peers in AS 64523, and then re-advertise them to other internal BGP peers in the same autonomous system.
Referring to the exhibit, which configuration must you deploy on ROUTER_1 to accomplish this task?

A. Configure ROUTER_1's internal BGP group with a routing policy that exports prefixes learned from internal BGP.
B. Configure ROUTER_1's internal BGP group with the keyword cluster, followed by a unique 32-bit number.
C. Configure a routing policy on ROUTER_1 that removes the no-export BGP community from all received prefixes.
D. Configure ROUTER_1 to belong to a different autonomous system than the other BGP routers in your network.

Answer: B

Explanation:
In theBorder Gateway Protocol (BGP), theSplit Horizonrule is a fundamental loop-prevention mechanism for internal sessions. This rule dictates that a BGP speaker must not advertise a route learned from anInternal BGP (IBGP)peer to any other IBGP peer within the same Autonomous System (AS). This ensures that routes do not circulate infinitely inside a network, as IBGP does not modify the AS_PATH attribute. Consequently, to maintain full reachability, a network normally requires a "full mesh" of IBGP sessions, where every BGP-speaking router is directly peered with every other router.
In the provided exhibit,ROUTER_1is part of AS 64523. The requirement is for ROUTER_1 to take prefixes learned from its internal peers and re-advertise them to other internal peers in the same AS. This behavior is a direct violation of the standard Split Horizon rule. According to Juniper Networks technical documentation, the standard solution to scale IBGP without a full mesh is to configureRoute Reflection.
When a router is configured as aRoute Reflector (RR), it is permitted to "reflect" (re-advertise) routes learned from one IBGP peer to another. In Junos OS, the mechanism to enable Route Reflection is to configure acluster IDwithin the BGP group. By adding the cluster keyword followed by a unique 32-bit identifier (usually the router's loopback address) to the internal BGP group configuration, the router assumes the role of an RR. It then follows specific reflection rules:
Routes learned from anEBGP peerare reflected to all IBGP peers.
Routes learned from aRoute Reflector Clientare reflected to all other clients and non-clients.
Routes learned from anon-clientare reflected to all clients.
Option A is incorrect because BGP advertisement rules are hard-coded; a standard export policy cannot override the Split Horizon rule. Option C handles traffic engineering tags but does not enable route reflection. Option D would change the session to EBGP, which does not address the internal reachability requirement within AS 64523. Therefore, configuring the cluster ID is the only valid way to achieve the desired re-advertisement behavior.

NEW QUESTION 55
By default, which MPLS operation is performed by the penultimate router in an LSP on the transport label?

A. swap
B. push
C. rewrite
D. pop

Answer: D

Explanation:
In a Multiprotocol Label Switching (MPLS) environment, label operations are categorized into three primary actions:Push(adding a label),Swap(replacing a label), andPop(removing a label). The specific behavior described in the question refers to a mechanism calledPenultimate Hop Popping (PHP).
According to Juniper Networks technical documentation, the goal of PHP is to improve forwarding efficiency at the egress point of a Label-Switched Path (LSP). TheEgress Label Edge Router (LER), which is the final destination for the LSP, would normally have to perform two lookups if it received a labeled packet: first, it would look up the label in its MPLS table to see it is the destination, and second, it would look up the underlying IP payload in its IP routing table (inet.0) to forward the packet.
To alleviate this burden, the Egress LER signals a special label value calledImplicit Null (Label 3)to its upstream neighbor (the penultimate router) during the signaling process (RSVP or LDP). When thepenultimate routerreceives a packet destined for that egress LER, it sees the instruction to pop the transport label. Consequently, the penultimate router performs aPopoperation, stripping away the outer MPLS label and sending the raw IP packet (or the remaining inner service label) to the Egress LER.

This allows the Egress LER to perform only a single lookup. If the transport label was the only label, the Egress LER simply performs a standard IP lookup. If there is a VPN label remaining, it performs a single MPLS lookup for the VRF. This "default" behavior in Junos OS optimizes the performance of the egress router by offloading the final label removal to the penultimate hop. Note that ifUltimate Hop Popping (UHP)were configured (via the explicit-null command), the penultimate router would perform aSwapto Label 0 instead of a Pop.

NEW QUESTION 60
What is the default route preference for an aggregate route?

A. 180
B. 150
C. 130
D. 5

Answer: C

Explanation:
In the Junos OS architecture,route preference(often referred to as administrative distance in other vendor platforms) is the primary metric used by the Routing Engine to select the "best" path when multiple protocols provide a route to the same destination. Each routing protocol and route type is assigned a default numeric value; the lower the value, the more preferred the route.
According to Juniper Networks technical documentation, anaggregate routeis assigned a default preference of 130. gregate routes are a form of static-like route used to group specific routes into a single, broader prefix to reduce the size of routing tables and limit the scope of routing updates. They are "protocol-independent" because they are not learned from a dynamic neighbor but are manually defined by the administrator.
To understand where130fits in the hierarchy, it is helpful to compare it with other common Junos preferences:
Directly connected interfaces:0
Static routes:5
OSPF Internal:10
IS-IS Level 1/2:15/18
Aggregate routes: 130
OSPF AS External:150
BGP (Internal and External):170
Generated routes:150
By setting the aggregate route preference to 130, Junos ensures that specific routes learned via IGPs (like OSPF or IS-IS) are preferred over the aggregate. This is essential because an aggregate route is often used as a "catch-all" or a discard route when more specific path information is missing. If the aggregate had a lower preference (like 5), it might override dynamic routing information, leading to suboptimal routing or black- holed traffic.

NEW QUESTION 63
Exhibit:

```
user@R10> show configuration protocols isis

interface ge-0/0/1.0 {

point-to-point;

}

interface ge-0/0/2.0 {

point-to-point;

}

interface lo0.0;

source-packet-routing {

srgb start-label 300000 index-range 10000;

}

level 1 disable;

level 2 wide-metrics-only;

reference-bandwidth 100g;
```

You have a network of ten routers that have all been configured with an identical SRGB. The exhibit shows the IS-IS configuration from a router called R10. The other nine routers do not yet have an IPv4 shortest-path SR-MPLS LSP to this router. Which missing part of the configuration must you add on R10 to solve this problem?

A. R10 must be configured with an explicit binding SID.
B. R10 must be configured with explicit IPv4 adjacency SID.
C. R10 must tag its internal IPv4 BGP prefixes with a BGP prefix SID.
D. R10 must be configured with an explicit IPv4 node SID.

Answer: D

Explanation:
In aSegment Routing (SR-MPLS)architecture using IS-IS as the control plane, routers exchange labels (segments) to build Label-Switched Paths (LSPs) without the need for traditional signaling protocols like LDP or RSVP. According to Juniper Networks technical documentation, for a router to be reachable via a shortest-path LSP from other nodes in the network, it must advertise aPrefix Segment Identifier (Prefix SID).
A specific type of Prefix SID is theNode SID, which is assigned to a loopback address (typically lo0.0) to uniquely identify the router within the SR domain. In the provided exhibit, routerR10has been configured with aSegment Routing Global Block (SRGB)starting at label 300000. This configuration tells the router which label range to use for global segments, but it does not automatically assign a label to its own loopback interface.
Without aNode SIDconfiguration, R10 is not telling its neighbors which specific index or label within that SRGB corresponds to its own address. Consequently, the other nine routers in the IS-IS area can calculate the shortest path to R10 using standard SPF, but they cannot perform the "label-binding" necessary to push an SR- MPLS label onto the packets.
To solve this, a Node SID must be explicitly configured under the loopback interface within the IS-IS protocol hierarchy, such as:

set protocols isis interface lo0.0 level 2 ipv4-node-sid index <value>
Analysis of incorrect options:
Binding SID (Option A):This is used to encapsulate or steer traffic into a specific policy or tunnel (like a TE-LSP) and is not required for basic shortest-path reachability.
Adjacency SID (Option B):These are generated automatically by Junos for each link and represent a specific local hop; they are not used for global "shortest-path" forwarding to a distant node.
BGP Prefix SID (Option C):This is used for BGP Egress Peer Engineering (EPE) or prefix advertisement via BGP, which is not relevant for building the underlying IS-IS SR-MPLS transport.
Therefore, configuring anexplicit IPv4 node SIDis the mandatory step to enable the rest of the network to build a shortest-path SR-LSP toward R10.

NEW QUESTION 66
You are a network architect designing a brand new network. You want to deploy RSVP LSPs in this network. You are currently in the process of choosing whether to run OSPF or IS-IS as your interior gateway protocol. In this scenario, which two statements are correct about IGP traffic engineering extensions in an RSVP network? (Choose two.)

A. You must explicitly configure IS-IS to carry traffic engineering extensions.
B. In OSPF, traffic engineering extensions are enabled by default.
C. You must explicitly configure OSPF to carry traffic engineering extensions.
D. In IS-IS, traffic engineering extensions are enabled by default.

Answer: CD

Explanation:
In a Juniper Networks environment, deployingRSVP-signaled LSPsrequires a functionalTraffic Engineering Database (TED). This database is populated by the Interior Gateway Protocol (IGP) using specific extensions that carry link-state information beyond simple reachability, such as available bandwidth, administrative groups (link coloring), and Maximum Reservable Bandwidth.
The behavior of these extensions differs betweenOSPFandIS-ISin Junos OS:
OSPF (Option C):By default, OSPF is a "pure" routing protocol. To support RSVP-TE, it must carry Opaque LSAs (Type 10). According to Juniper documentation, you mustexplicitly configuretraffic engineering within the OSPF protocol hierarchy using the set protocols ospf traffic-engineering command. Without this command, OSPF will not flood the TE information required by the Constrained Shortest Path First (CSPF) algorithm, and LSPs will fail to establish.
IS-IS (Option D):IS-IS was designed to be extensible through the use of TLVs (Type, Length, Value). In Junos OS,IS-IS traffic engineering extensions are enabled by defaultonce the protocol is active. As soon as you enable IS-IS on an interface, it begins to advertise the wide metrics and TE TLVs (like TLV 22 and 135) necessary for building the TED.
This distinction is a common design consideration for network architects. While IS-IS simplifies the rollout of MPLS by having TE enabled "out of the box," OSPF requires that extra configuration step to transition from a standard IGP to a TE-aware protocol.

NEW QUESTION 70
Exhibit:

```
[edit interfaces]
user@switch# show
xe-0/0/4 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members 10;
            }
        }
    }
}
```

on a Juniper switch. It shows interfacexe-0/0/4withunit 0andfamily ethernet-switching. Under vlan, it lists members 10;`]
Referring to the exhibit, which two statements are true? (Choose two.)

A. The interface receives tagged traffic.
B. The interface is a part of a VLAN that uses VLAN ID 10.
C. The interface receives untagged traffic.
D. The interface is a member of the VLAN named 10.

Answer: CD

Explanation:

In Junos OS for switching platforms, an interface is configured for Layer 2 bridging under the family ethernet-switching hierarchy. The way an interface handles VLAN traffic depends on its port mode:accessortrunk.
According to Juniper Networks technical documentation, when an interface is configured simply with members, it defaults to anaccess port. In an access port configuration:
The port is a member of only a single VLAN.
The portreceives and sends untagged traffic (Option C). Any untagged frame arriving at this interface is implicitly associated with the configured VLAN member.
The interface does not expect or process 802.1Q tags in incoming frames.
In the exhibit, interface xe-0/0/4 has members 10;. In Junos, the members statement can reference either a VLAN name or aVLAN ID. However, when the configuration is shown as members 10; without further context of the specific ID mapping, the most precise interpretation of the CLI output provided is that the interface is a member of the VLAN named 10 (Option D). While "10" could be the numerical ID, Junos primarily maps members by their defined administrative name.
Why other options are incorrect:
Option A:Access ports do not receive tagged traffic; only trunk ports (which require the port-mode trunk and vlan members [...] statements) are designed to process tagged frames.
Option B:While the VLAN named 10likelyhas a VLAN ID of 10, the exhibit does not explicitly confirm the ID mapping. In Junos, a VLAN named "10" could technically have a different tag ID (e.g., VLAN "Office" with ID 10). Option D is the more accurate direct reading of the displayed member configuration.

NEW QUESTION 74
During OSPF neighbor establishment, which packet type is used to describe the contents of the link-state database?

A. Link-State Request (LSR)
B. Hello packet
C. Database Description (DBD)
D. Link-State PDU (LSP)

Answer: C

Explanation:
In the OSPF (Open Shortest Path First)protocol, ensuring that all routers within an area have a synchronized Link-State Database (LSDB)is fundamental to building a consistent loop-free topology. During the adjacency formation process—specifically when transitioning from theExStartstate to theExchangestate—routers must determine what information they are missing from their neighbors without sending the entire database at once, which would be highly inefficient.
The Database Description (DBD)packet, also known as a DDP, is the mechanism used for this summary exchange. According to Juniper Networks technical documentation, the DBD packet does not contain full Link-State Advertisements (LSAs). Instead, it contains only theLSA headers, which include the LSA type, the ID of the advertising router, and the sequence number.
By exchanging these headers, a Juniper router can compare the neighbor's database summary against its own local LSDB. If the router identifies a header in the DBD packet that represents a newer or missing entry, it records that LSA in its "Link-State Request List." This collaborative "handshake" ensures that only the necessary, updated information is requested in the subsequentLink-State Request (LSR)phase. It is important to distinguish this from theLink-State PDU (LSP)mentioned in Option D, which is actually the term used in the IS-IS protocol, not OSPF. In OSPF, the functional unit is the LSA, and the transport vehicle for the initial summary is the DBD packet. This methodical synchronization is what allows OSPF to scale effectively in large service provider environments.

NEW QUESTION 76
.....

Relate Links

100% Pass Your JN0-364 Exam with Exambible Prep Materials

https://www.exambible.com/JN0-364-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/