Juniper

Exam Questions JN0-364

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

NEW QUESTION 1
You are troubleshooting a Level 1 IS-IS router that has an adjacency with a Level 1/2 router. Which two statements are correct in this scenario? (Choose two.)

A. The Level 1/2 router merges Level 1 and Level 2 into one complete topology.
B. The Level 1 router will learn the full topology of the Level 2 network.
C. The Level 1/2 router sees the Level 1 network and the Level 2 network as two separate topologies.
D. The Level 1 router will only learn the topology of the Level 1 network.

Answer: CD

Explanation:
In the context ofJuniper Networks Junos OSand theIS-IS (Intermediate System to Intermediate System) protocol, understanding the hierarchical relationship between router levels is critical for effective troubleshooting and design. IS-IS uses a two-level hierarchy to manage scalability:Level 1 (L1), which represents intra-area routing, andLevel 2 (L2), which represents inter-area backbone routing.
When a router is configured as aLevel 1/2 (L1/L2)device, it acts as a bridge between the two levels. According to Juniper technical documentation, an L1/L2 router maintains two completelyseparate Link- State Databases (LSDB)—one for Level 1 and one for Level 2. It does not merge these into a single topology. This separation ensures that local area topology changes (L1) do not necessarily flood into the backbone (L2) unless specific redistribution is configured, and vice versa. Therefore, statementCis correct because the L1/L2 router maintains distinct SPF (Shortest Path First) computations for each level.
Regarding the visibility of the Level 1 router, IS-IS is designed to keep L1 areas "stubby" by default. ALevel 1 routeronly possesses the topology information for its own area (the Level 1 LSDB). It does not receive specific L2 routes or the L2 topology. Instead, the L1/L2 router sets theAttached (ATT) bitin its L1 Link- State PDUs (LSPs) to signal to L1-only routers that it has a connection to the backbone. The L1 router then generates a default route pointing to the L1/L2 router to reach inter-area destinations. This confirms that statementDis correct: the L1 router's knowledge is limited to its local L1 topology.
Conversely, statements A and B are incorrect because merging topologies would violate the hierarchical scaling principles of IS-IS, and L1 routers never learn the full L2 topology without explicit, non-standard route leaking.

NEW QUESTION 2
Which two statements about graceful restart are correct? (Choose two.)

A. Graceful restart restarting router mode is not enabled by default.
B. Graceful restart helper mode is enabled by default.
C. Graceful restart requires that GRES be enabled.
D. Graceful restart uses nonstop bridging for forwarding operations.

Answer: AB

Explanation:
Graceful Restart (GR)is a high-availability mechanism designed to minimize the impact of a routing protocol process (rpd) restart or a Routing Engine (RE) switchover. It allows a router to continue forwarding traffic while the control plane is recovering, provided that the data plane (Packet Forwarding Engine) remains intact.
According to Juniper Networks documentation, Graceful Restart operates in two distinct roles:
Restarting Mode:This is the role of the router that is actually undergoing the restart. In Junos OS, this mode isnot enabled by default (Option A). An administrator must explicitly configure graceful-restart under the [edit routing-options] hierarchy to allow the router to signal its neighbors that it is attempting a graceful recovery.
Helper Mode:This is the role of the neighboring routers. When a neighbor sees a router restart, if it is in "helper mode," it will continue to forward traffic toward the restarting router and will not flush the associated routes from its forwarding table for a specified period. In Junos,helper mode is enabled by default (Option B)for most protocols (OSPF, BGP, IS-IS). This means that even if you haven't configured GR on your own router, it will automatically assist its neighbors if they perform a graceful restart.
Why other options are incorrect:
Option C:WhileGRES (Graceful Routing Engine Switchover)is often usedwithGraceful Restart to handle hardware-level RE failures, they are independent features. GR can function during a simple software process restart without dual REs or GRES.
Option D:Nonstop Bridging (NSB)is a separate high-availability feature for Layer 2 protocols (like STP). While it shares a similar goal, Graceful Restart is specifically a Layer 3 protocol mechanism (Layer 2 does not use "helper" routers in the same way).

NEW QUESTION 3
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 betweenFullandDown 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 theDead Interval. If a router does not receive a Hello packet within the Dead Interval (usually 40 seconds), it tears down the adjacency. Aphysical 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 theExchangeor ExStartstate. 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 theInitstate; 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.

NEW QUESTION 4
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 5
Exhibit:

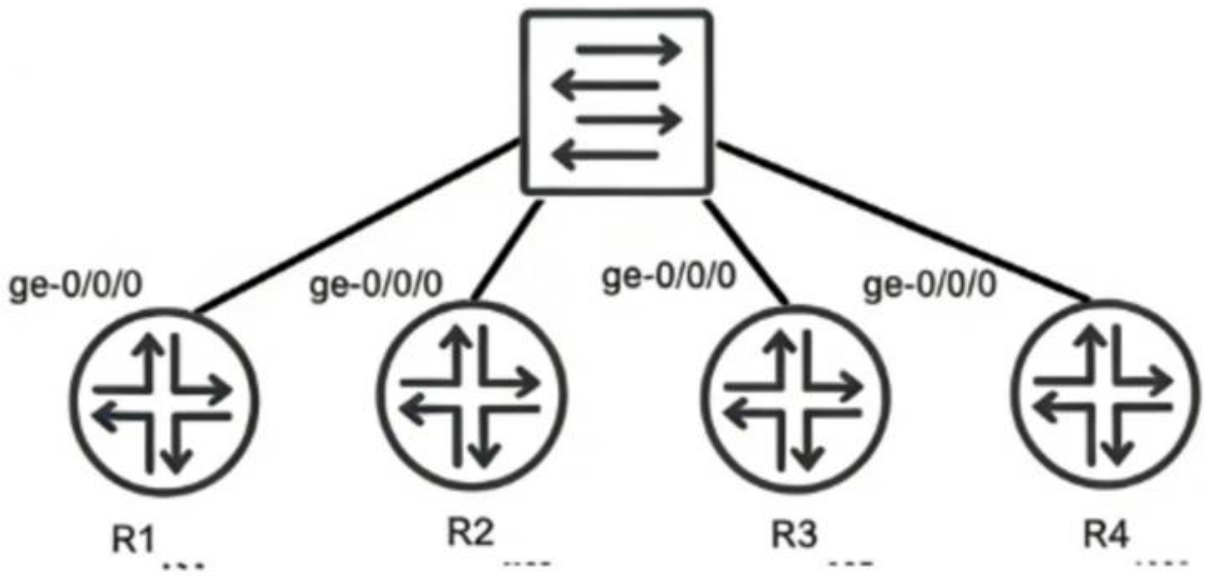


```
user@R1> show configuration routing-options
router-id 192.168.1.1;

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 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 7
Exhibit:

```
user@R1> show route 10.16.2.0/23 exact detail

inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)

10.16.2.0/23 (1 entry, 1 announced)

*Aggregate Preference: 130

Next hop type: Reject

Address: 0x8f3fd44

Next-hop reference count: 2

State:

Age: 1:39:21

Task: Aggregate

Announcement bits (1): 0-KRT

AS path: I (LocalAgg)

Flags: Depth: 0 Active

AS path list:

AS path: I Refcount: 2

Contributing Routes (2):

10.16.2.0/24 proto Direct

10.16.3.0/24 proto Direct
```

Which destination IP address will be matched by the aggregate route shown in the exhibit?

A. packets destined to 10.16.3.79
B. packets destined to 10.16.0.4
C. packets destined to 10.16.4.183
D. packets destined to 10.16.1.214

Answer: A

Explanation:
In the Juniper Networks Junos operating system,aggregate routesare used to represent a group of more specific routes with a single, shorter prefix. This technique is essential for reducing the size of routing tables and minimizing the volume of routing updates sent to neighbors. According to Juniper technical documentation, for a destination IP address to "match" a specific route, it must fall within the range defined by the network address and its associated CIDR mask.
The provided exhibit shows a detailed lookup for the aggregate route$10.16.2.0/23$. To determine the range of IP addresses covered by a $/23$ mask, we examine the binary representation of the third octet. A $/23$ mask means the first 23 bits are fixed. For the address $10.16.2.0$:
The first two octets (10.16) are fixed.
The third octet (2) is 00000010 in binary.
The 23rd bit is the second-to-last bit of this octet.
The $/23$ range allows the 24th bit (the last bit of the third octet) and all 8 bits of the fourth octet to vary.
This results in a range where the third octet can be either 2 (00000010) or 3 (00000011). Therefore, the aggregate route $10.16.2.0/23$ covers all IP addresses from$10.16.2.0$ to $10.16.3.255$. The exhibit further confirms this by listing the "Contributing Routes": $10.16.2.0/24$ and $10.16.3.0/24$.
Analyzing the provided options against this range:
* 10.16.3.79 (Option A):This address falls squarely within the $10.16.2.0$ to $10.16.3.255$ range.
* 10.16.0.4 (Option B):This address falls in the $10.16.0.0/23$ range (0.0 to 1.255).
* 10.16.4.183 (Option C):This address falls in the $10.16.4.0/23$ range (4.0 to 5.255).
* 10.16.1.214 (Option D):This address also falls in the $10.16.0.0/23$ range.
Consequently,10.16.3.79is the only destination listed that matches the aggregate route shown. It is also important to note theNext hop type: Rejectin the exhibit; this means that if a packet matches the aggregate but does not match any of the more specific contributing routes, the router will drop the packet and send an ICMP unreachable message to the source.

NEW QUESTION 8
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 9
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
By default, which routing table contains a list of all ingress LSPs?

A. inet.2
B. inet.3
C. inet.1
D. inet.0

Answer: B

Explanation:
In the Juniper Networks Junos operating system, the management of routing information is partitioned into several distinct routing tables (RIBs), each serving a specific architectural purpose. When dealing with Multiprotocol Label Switching (MPLS), understanding the distinction between inet.0 and inet.3 is fundamental for troubleshooting and traffic engineering.
Theinet.3routing table is specifically designed to store the egress IPv4 addresses ofLabel-Switched Paths (LSPs). When an ingress router successfully establishes an LSP (via RSVP or LDP), it places the host address of the egress router (the tail-end) into the inet.3 table. This table is not used for general packet forwarding; instead, it is primarily used by theBorder Gateway Protocol (BGP)for next-hop resolution. When BGP receives a route, it checks both inet.0 and inet.3 to resolve the next hop. If a matching entry exists in inet.3, the router knows it can reach that destination via an MPLS tunnel, allowing for the encapsulation of BGP traffic within MPLS.
In contrast,inet.0is the default unicast routing table used for standard IPv4 forwarding and contains routes learned via IGPs (OSPF, IS-IS) or static routing.inet.1is utilized for multicast forwarding (MBGP), andinet.2 is typically used for Multicast Source Discovery Protocol (MSDP) or RPF checks in multicast environments. By isolating LSP egress points in inet.3, Junos prevents MPLS-specific paths from interfering with standard IGP path selection unless the administrator explicitly chooses to merge them (e.g., using the traffic- engineering bgp-igp command). Therefore, by default, the ingress router maintains its list of reachable LSP endpoints in inet.3.

NEW QUESTION 10
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 12

Exhibit:

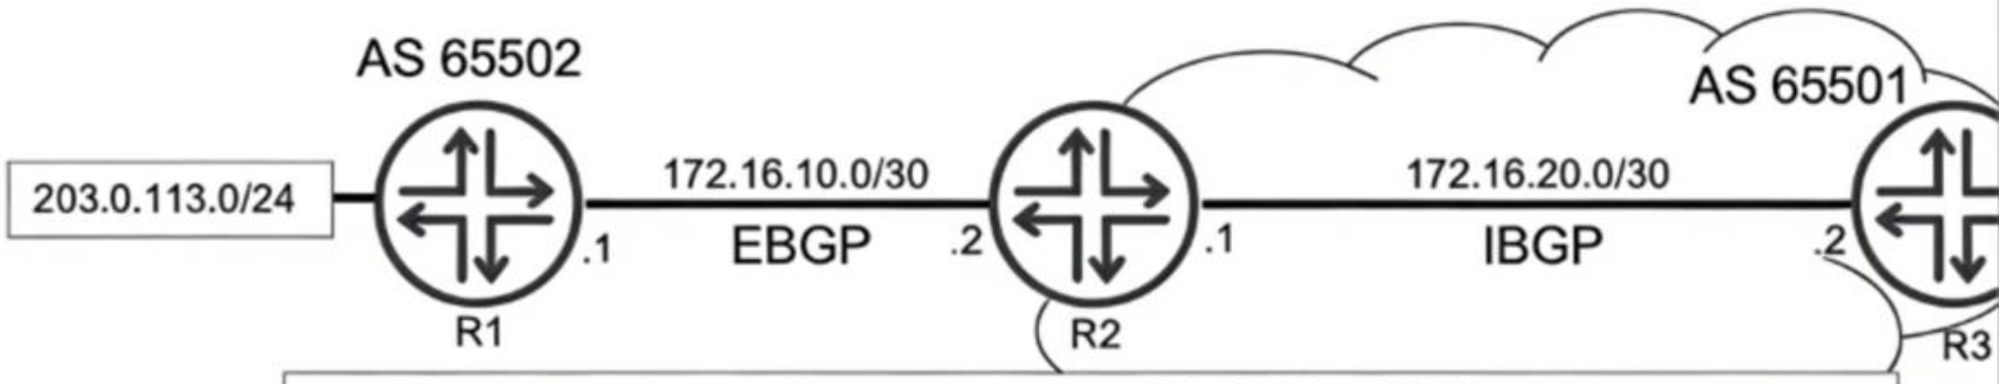


```
user@R3> show route receive-protocol bgp 172.16.20.1 hidden

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 16
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 19
What are two types of BGP messages exchanged while in the Established state? (Choose two.)

A. open
B. request
C. update
D. notification

Answer: CD

Explanation:
In theBorder Gateway Protocol (BGP)finite state machine (FSM), theEstablishedstate is the final and functional stage of a BGP peering session. According to Juniper Networks technical documentation, once a session reaches this state, the two peers have successfully exchanged Open messages and agreed upon session parameters (such as AS numbers, hold timers, and BGP identifiers). Only after the session is "Established" can the routers begin the actual exchange of network layer reachability information (NLRI).
The most frequent message type exchanged in the Established state is theUPDATEmessage. These messages are the heart of BGP operations; they are used to advertise new feasible routes to a peer or to withdraw routes that are no longer reachable. An UPDATE message contains path attributes (like AS-Path, Next-Hop, and Local Preference) and the associated prefixes. In a stable network, UPDATE messages are only sent when there is a change in the topology, adhering to BGP??s incremental update philosophy.
The second message type that can be exchanged in this state is theNOTIFICATIONmessage. While ideally, a session stays established, any detected error—such as a hold timer expiration, a malformed update, or a manual "clear" command—will trigger the transmission of a NOTIFICATION message. This message informs the peer of the specific error code and immediately causes the BGP session to transition back to the Idle state, tearing down the TCP connection.
It is important to note thatOPENmessages (Option A) are only used during the session initialization phase to transition from the OpenConfirm state to Established.REQUEST(Option B) is not a valid BGP message type defined in the standard (RFC 4271); the closest equivalent in functionality would be a Route-Refresh message, which is a separate extension. Therefore, in the context of standard BGP operations within the Established state, Updates and Notifications are the correct answers.

NEW QUESTION 20
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 24
Which term describes the router where traffic enters an MPLS label-switched path (LSP)?

A. egress router
B. transit router
C. penultimate router
D. ingress router

Answer: D

Explanation:
In the architecture of aLabel-Switched Path (LSP), routers are categorized based on their role in the handling of a specific packet's lifecycle through the MPLS network. Juniper Networks documentation defines these roles clearly:
TheIngress Router (Option D), also known as theIngress Label Edge Router (LER), is the entry point of the LSP. Its primary responsibility is to take an incoming "unlabeled" packet (usually a standard IPv4 or IPv6 packet), perform a route lookup, and determine which LSP the packet should follow. Once determined, the Ingress router performs aPushoperation, where it encapsulates the packet with an MPLS label header and forwards it toward the next hop. This is where the transition from IP-based forwarding to Label-based switching occurs.
To contrast this with the other options:
Transit Router (Option B):These are routers located between the ingress and egress. They performSwapoperations, replacing an incoming label with an outgoing label based on the Label Forwarding Information Base (LFIB).
Egress Router (Option A):This is the "tail-end" of the LSP where the packet exits the MPLS domain and the final label is removed (if it hasn't been removed already by the penultimate hop).
Penultimate Router (Option C):This is the second-to-last router in the path. As discussed in previous questions, it often performs thePopoperation (Penultimate Hop Popping) to remove the transport label before sending the packet to the Egress LER.
Therefore, the router where traffic first "enters" the LSP and receives its initial label is strictly defined as the Ingress router.

NEW QUESTION 27
Which IPv6 extension header is used to specify intermediate nodes for a packet's path?

A. hop-by-hop options
B. routing
C. fragment
D. destination options

Answer: B

Explanation:
In the IPv6 architecture, the base header is kept at a fixed size of 40 bytes to streamline processing. Any additional features or options are handled byExtension Headers, which are inserted between the IPv6 header and the upper-layer protocol. According to Juniper Networks technical documentation and RFC 8200, when a source node needs to list one or more intermediate nodes to be "visited" on the way to the final destination, it utilizes theRouting extension header (Option B).
The Routing header is functionally similar to the "Source Route" option in IPv4. When a packet contains a Routing header, it is addressed to the first intermediate node listed in the header. That node examines the header, swaps its own address with the next address in the list, and forwards the packet. This process continues until the packet reaches the final destination. This is a foundational component for technologies likeSegment Routing over IPv6 (SRv6), where the Routing header (specifically the Segment Routing Header or SRH) is used to steer traffic through a specific set of service instructions or nodes.
To distinguish this from the other options:
Hop-by-hop options (Option A):These carry information that must be examined byeverynode along the path (such as Router Alert), not just specific intermediate nodes.
Fragment (Option C):This is used only when the source node needs to fragment a packet that exceeds the path MTU.
Destination options (Option D):These carry optional information intended specifically for the destination node (or nodes listed in a Routing header), but they do not dictate the path themselves.

NEW QUESTION 29
You are the administrator for two Junos routers called R1 and R2. These two routers are directly connected to each other. These two routers run IS-IS and BFD. R1 is configured to send BFD packets every 300 milliseconds. R2 is configured to send BFD packets every 400 milliseconds. In this situation, what is the expected outcome?

A. Each router will send BFD packets at the rate that has been locally configured.
B. BFD will fail due to the mismatched timers.
C. Each router will negotiate to send BFD packets at the slowest of the two rates.
D. Each router will negotiate to send BFD packets at the fastest of the two rates.

Answer: C

Explanation:
In the context of Juniper Networks High Availability,Bidirectional Forwarding Detection (BFD)is a lightweight protocol designed to provide fast failure detection for the forwarding path. Unlike the slow "hello" mechanisms found in IGPs like OSPF or IS-IS, BFD can detect link or neighbor failures in sub-second intervals.
According to Juniper Networks technical documentation, BFD operates through a negotiation process. When two routers establish a BFD session, they exchange their locally configuredMinimum Transmit IntervalandMinimum Receive Intervalwithin the BFD control packets. The fundamental rule of BFD negotiation is that the routers must agree on a common timing value that accommodates the slower of the two devices to ensure stability and prevent "false positives" (detecting a failure when none exists simply because one router cannot keep up with the processing speed).
In this scenario, R1 expects to send at 300ms, while R2 is configured for 400ms. During the handshake, R1 informs R2 it is capable of 300ms, but R2 informs R1 it can only support a minimum of 400ms. Consequently, the routers will negotiate to use theslowest of the two rates (400ms). Specifically, the transmission interval of one router is matched to the receive interval of the other. By choosing the highest common denominator (the slowest rate), the BFD session ensures that both routers have sufficient time to process incoming control packets. This negotiation allows BFD to be highly flexible in heterogeneous environments where different hardware platforms may have varying CPU capabilities for handling rapid heartbeat packets.

NEW QUESTION 31
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 32
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

JN0-364 Practice Exam Features:

* JN0-364 Questions and Answers Updated Frequently

* JN0-364 Practice Questions Verified by Expert Senior Certified Staff

* JN0-364 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* JN0-364 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year