Juniper

Exam Questions JN0-364

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

NEW QUESTION 1
You have configured an MPLS LSP that begins on R1 and terminates on R5 using the Junos default settings. Referring to the exhibit, which router will perform only label swap operations?

A. R4
B. R3
C. R5
D. R1

Answer: B

Explanation:
In an MPLS network, routers are categorized by their role along a Label Switched Path (LSP). In this scenario, the LSP originates onR1 (Ingress LER)and terminates onR5 (Egress LER). Between these two endpoints are the Provider (P) routers, also known as Transit Label Switching Routers (LSRs), which include R2, R3, and R4.
To identify which router performsonlylabel swap operations, we must look at the standard Junos data plane behavior:
R1 (Ingress LER):Performs aPushoperation. It receives native IP traffic from Networks 1 or 2, looks up the destination, and imposes (pushes) an MPLS label onto the packet before sending it to R2.
R2 and R3 (Transit LSRs):These routers perform aSwapoperation. They receive a labeled packet, look up the incoming label in their Label Forwarding Information Base (LFIB), replace it with an outgoing label provided by the downstream neighbor, and forward it.
R4 (Penultimate Hop):By default, Junos usesPenultimate Hop Popping (PHP). Because R4 is the second-to-last router before the egress (R5), the egress router R5 advertises an "implicit-null" label (Label 3) to R4. This instructs R4 to perform aPopoperation—removing the MPLS label entirely—and sending the native IP packet to R5.
R5 (Egress LER):Receives the packet (which is already unlabeled due to PHP) and performs a standard IP route lookup to reach the final destination in Network 3 or 4.
Among the options provided,R3is the only router that is a transit LSR butnotthe penultimate hop. While R2 also performs a swap, it is not an option. R4 performs a Pop (due to PHP), R1 performs a Push, and R5 performs an IP lookup. Therefore, R3 is the correct answer as it solely performs the label swap operation.

NEW QUESTION 2
What are three extension headers supported by IPv6? (Choose three.)

A. destination options
B. hop-by-hop options
C. protocol
D. header checksum
E. fragment

Answer: ABE

Explanation:
One of the most significant architectural improvements in IPv6 is the move from a complex, variable-length header (as seen in IPv4) to a streamlined, fixed-length base header of 40 bytes. Additional functionality that was previously handled by "Options" in IPv4 is now moved toExtension Headers, which are inserted between the IPv6 base header and the upper-layer protocol (TCP/UDP).
According to Juniper Networks technical documentation and RFC 8200, the following are valid IPv6 Extension Headers:
Hop-by-Hop Options (Option B):This header carries optional information that must be examined by every node along the delivery path. It is used for features like the Router Alert and Jumbo Payload options.
Fragment (Option E):Unlike IPv4, where any router can fragment a packet, in IPv6, fragmentation is performed only by the source node. The Fragment header contains the information necessary for the destination to reassemble the packet (Offset, Identification, and More Fragments flag).
Destination Options (Option A):This header carries information intended only for the destination node. It can appear twice: once before a routing header and once after.
Why other options are incorrect:
Protocol (Option C):In IPv4, this was a field in the header. In IPv6, this is replaced by theNext Header field, which identifies the type of the following header (whether it's an extension header or the upper- layer protocol).
Header Checksum (Option D):This field was entirely removed in IPv6. IPv6 relies on the data link layer (Ethernet) and the transport layer (TCP/UDP) to perform error detection, significantly reducing the processing overhead for routers in the core of a service provider network.

NEW QUESTION 3
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 4
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 5
Which two statements regarding GRE and IP-IP tunnels are correct? (Choose two.)

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

Answer: AD

Explanation:
In Juniper Networks Junos OS,Generic Routing Encapsulation (GRE)andIP-in-IP (IP-IP)are common tunneling mechanisms used to transport packets across a network by encapsulating them within another protocol. Understanding the header structure and the limitations of these protocols is essential for proper MTU (Maximum Transmission Unit) management and security design.
Overhead (Option A):
Both GRE and IP-IP tunnels operate by adding an additional IP header to the original packet. An IP-IP tunnel (Protocol 4) adds a20-byteIPv4 header. A GRE tunnel (Protocol 47) adds the same20-bytedelivery IP header plus a minimum4-byteGRE header (totaling 24 bytes, which can increase if keys or sequencing are used). Because these headers are added to the payload, the total size of the packet increases. This "overhead" means that if the original packet was already at the MTU limit (e.g., 1500 bytes), the encapsulated packet will exceed it, potentially leading to fragmentation or the need to adjust theTCP MSS (Maximum Segment Size).
Encryption (Option D):
Crucially, according to Juniper Service Provider documentation, neither GRE nor IP-IP provides nativeencryptionor data confidentiality. They are encapsulation protocols, not security protocols. The payload remains in cleartext and is visible to any device along the path. If security and encryption are required for data traversing these tunnels, they must be combined withIPsec (IP Security). While GRE is often used as the "carrier" for IPsec (to allow multicast or dynamic routing protocols which IPsec alone does not support), the GRE protocol itself remains an unencrypted delivery mechanism. Therefore, statements A and D accurately describe the architectural behavior of these tunnel types.

NEW QUESTION 6
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 7
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 8
You are using EBGP to connect to two upstream peers in the same AS. You want to make one of the links less preferred for traffic entering your network from the peer's AS. Which feature should you use to achieve this goal?

A. a route reflector
B. origin code
C. AS-path prepending
D. local preference

Answer: C

Explanation:
In the world of BGP, controllinginbound traffic(traffic entering your network) is significantly more challenging than controlling outbound traffic because it requires influencing a decision made by an external Autonomous System (AS). According to Juniper Networks documentation, when you have multiple links to the same AS or even different ASes, the BGP path selection process is used by the upstream neighbor to decide which path to take to reach your prefixes.
AS-Path Prependingis the standard technique used to make a path appear less attractive to external peers. By artificially lengthening the AS_PATH attribute on the BGP advertisements sent over a specific link, you exploit the BGP best-path algorithm rule that prefers a shorter AS path. When you prepend your own AS number multiple times to the update sent to the "less preferred" peer, that peer??s BGP routers will see a longer path compared to the alternative link and will naturally prefer the shorter, unprepended route.
It is important to distinguish why other options are incorrect for this specific goal:
Local Preference (Option D):This is a well-known discretionary attribute used to influenceoutbound traffic. It is not advertised to EBGP peers; therefore, your upstream neighbor cannot see your local preference settings.
Origin Code (Option B):While the origin code (IGP, EGP, or Incomplete) is a tie-breaker in the selection process, it is rarely used for traffic engineering and lacks the granular control provided by prepending.
Route Reflector (Option A):This is an Internal BGP (IBGP) scaling mechanism used to reduce the need for a full mesh of peers within an AS; it does not directly influence external path selection by an upstream provider.
Junos OS allows you to easily implement prepending viarouting policiesapplied as an "export" policy to the EBGP neighbor. By using the as-path-prepend action within a policy term, you can selectively degrade a path's attractiveness to manage your inbound bandwidth.

NEW QUESTION 9
For two or more switches to participate in the same MSTP region, which parameter must match?

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

Answer: A

Explanation:
Multiple Spanning Tree Protocol (MSTP), as defined in IEEE 802.1s and implemented in Juniper Networks Junos OS, allows for the grouping of VLANs into specific spanning tree instances. This provides significant scalability and load-balancing advantages over traditional STP or RSTP. To achieve this, switches must be grouped into logical "Regions."
According to Juniper documentation, for two or more switches to be considered part of the sameMSTP Region, they must possess an identicalMSTP Configuration Identifier. This identifier consists of three specific attributes that must match exactly across all participating switches:
MSTI Name (Region Name):A descriptive string (up to 32 characters) that identifies the region.
MSTI Revision Level:A numerical value (0–65535) used to track configuration changes.
VLAN-to-Instance Mapping:The specific table that defines which VLAN IDs are associated with which Multiple Spanning Tree Instances (MSTIs).
If even one of these parameters—such as theRegion name(Option A)—differs, the switches will treat each other as being in separate regions. When switches are in different regions, they interact using theCommon Spanning Tree (CST), effectively seeing the other region as a single "virtual bridge," which limits the granularity of traffic engineering.
The Extended system ID(Option B) is a component of the Bridge ID used to carry VLAN information in PVST+ but is not a region-matching requirement.Root bridge priority(Option C) andRoot bridge ID(Option D) are variables used during the STP election process to determine the topology's root, but they do not define the boundaries of an MSTP region itself.

NEW QUESTION 10
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 10
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 11
Which OSPF packet type is used to initiate and maintain neighbor relationships?

A. Hello
B. Database Description
C. Link-State Update
D. Link-State Acknowledgment

Answer: A

Explanation:
TheHello packetis the most basic, yet most vital, component of the OSPF protocol. It serves as the primary mechanism for neighbor discovery, parameter negotiation, and "keepalive" functionality. Per Juniper Networks' routing documentation, OSPF routers use the Hello protocol to dynamically discover other OSPF-enabled routers on their directly connected segments.
When OSPF is enabled on a Junos interface, the router begins multicasting Hello packets (typically to the224.0.0.5"All OSPF Routers" address). This initiates the neighbor relationship. For two routers to move beyond theInitstate and become neighbors, they must agree on several critical parameters contained within the Hello packet:
Area ID:Routers must be in the same OSPF area.
Authentication:Passwords or keys must match.
Timers:The Hello and Dead intervals must be identical.
Options:Such as Stub area flags.
Beyond the initial "initiation," the Hello packet is used tomaintainthe relationship. By continuously sending these packets at a fixed interval (the Hello interval), a router signals to its peers that it is still functional. If a router stops receiving Hello packets from a neighbor for a duration exceeding theDead Interval, it declares the neighbor "down," flushes the associated LSAs from the database, and triggers a new SPF calculation.
Furthermore, on multi-access networks like Ethernet, the Hello packet is the vehicle for the election of theDesignated Router (DR)andBackup Designated Router (BDR). By exchanging priority values and Router IDs within the Hello packets, the segment can elect a central point of contact to minimize the number of adjacencies required on the wire.

NEW QUESTION 14
You are asked to add next-hop redundancy using VRRP for an IPv6 enabled service. The configured primary router must always be active when available, and the servers connected to the network must be able to ping their gateway. Which VRRP element is required to accomplish this requirement?

A. The backup router requires the track parameter to track the primary router's interface.
B. The preempt parameter must be added to the VRRP configuration.
C. Both routers running VRRP will require a static ARP entry to be configured for the VRRP VIP.
D. The accept-data parameter must be added to the VRRP configuration.

Answer: D

Explanation:
InVirtual Router Redundancy Protocol (VRRP), the primary goal is to provide a highly available default gateway for end hosts. However, there is a specific operational behavior in the VRRP standard (RFC 3768/RFC 5798) regarding how the "Virtual Router" responds to traffic destined for its own Virtual IP (VIP).
According to Juniper Networks documentation, by default, a VRRP router that is in the Master state will only respond to packets destined for the VIP if that router is theIP Address Owner(meaning its physical interface IP matches the VIP). If the router is a "non-owner" (a common configuration in many networks), it will forward traffic on behalf of the VIP but will not respond to management traffic, such asICMP Echo Requests (Pings), directed at the VIP itself.
To satisfy the requirement that "servers connected to the network must be able to ping their gateway," theaccept-data (Option D)parameter must be configured. In Junos OS, the accept-data statement allows the VRRP Master to respond to traffic destined for the virtual IP address even if it is not the address owner. This includes responding to Pings and allowing other management connections like SSH or Telnet to the VIP.
Regarding the other options:
Preempt (Option B):While preempt is often used to ensure the primary router regains control, in Junos, a router with the highest priority (255) defaults to preemptive behavior, and accept-data is specifically what solves the "pinging the gateway" requirement.
Track (Option A):Tracking is used for failover logic but doesn't affect the ability to ping the VIP.
Static ARP (Option C):This is unnecessary as VRRP uses a virtual MAC address to ensure hosts can resolve the VIP via standard NDP (for IPv6) or ARP (for IPv4).

NEW QUESTION 17
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 18
Exhibit:

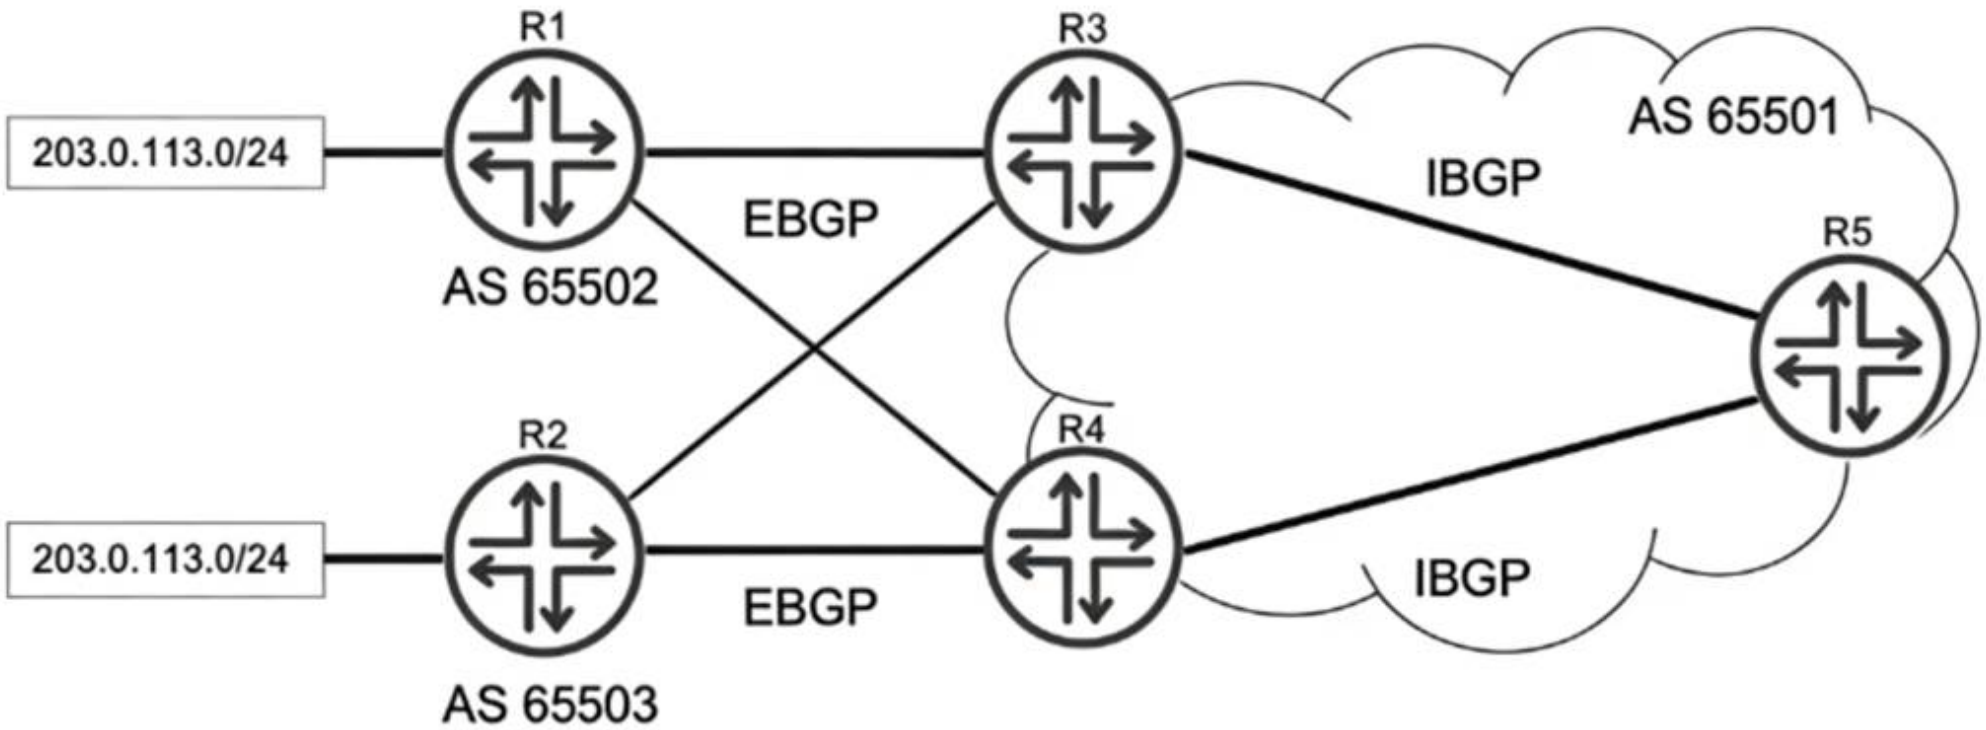


```
user@R3> show configuration policy-options
policy-statement export-to-ibgp {
    from {
        route-filter 203.0.113.0/24 exact;
    }
    then {
        local-preference 150;
        next-hop self;
        accept;
    }
}
```

```
user@R4> show configuration policy-options
policy-statement export-to-ibgp {
    from {
        route-filter 203.0.113.0/24 exact;
    }
    then {
        local-preference 200;
        next-hop self;
        accept;
    }
}
```

Referring to the exhibit, R1 and R2 are advertising the same prefix 203.0.113.0/24 to R3 and R4 over EBGP. R3 and R4 both advertise this prefix to R5. Which advertisement does R5 choose to install in its routing table?

A. The advertisement from R4 is chosen.
B. The advertisements from both R3 and R4, but R3 is chosen for forwarding.
C. The advertisement from R3 is chosen.
D. The advertisements from both R3 and R4, but R4 is chosen for forwarding.

Answer: A

Explanation:
In a Juniper Networks environment, when a router receives multiple BGP paths for the same destination prefix, it utilizes theBGP Path Selection Algorithmto determine the single "best" path to install in the routing table and advertise to other peers. This selection process follows a strict hierarchy of attributes.
According to Juniper Networks technical documentation, the very first attribute evaluated by the BGP process (after ensuring the next hop is reachable) is theLocal Preference. Local preference is a well-known discretionary attribute used to communicate a preference for a specific exit point from the local Autonomous System (AS). A higher local preference value is always preferred over a lower one.
Analyzing the exhibit:
R3receives the prefix from R1 and applies an export policy to its IBGP session that sets thelocal preference to 150.
R4receives the same prefix from R2 and applies an export policy to its IBGP session that sets thelocal preference to 200.
R5receives both of these IBGP updates from R3 and R4.
When R5 runs the best-path algorithm for the 203.0.113.0/24 prefix, it compares the local preference values. Since the path from R4 has a local preference of 200 and the path from R3 has a local preference of 150, R5 immediately selects the path fromR4as the best route. Because BGP is designed to prevent loops and maintain a consistent view, only this single best path is installed as the active route in R5's routing table (inet.0). Options B and D are incorrect because they imply multiple paths are installed for forwarding, which only occurs if specific multipath load-balancing is configured, which is not indicated here.

NEW QUESTION 19
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 24
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 25
.....

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