Juniper

Exam Questions JN0-364

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

NEW QUESTION 1
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 2
Which BGP attribute is optional, transitive, and is passed unchanged to other BGP peers if not recognized?

A. Origin
B. AS Path
C. Community
D. MED

Answer: C

Explanation:
BGP attributes are categorized into four distinct types based on how they are handled by a BGP speaker:Well- known mandatory,Well-known discretionary,Optional transitive, andOptional non-transitive. Understanding these categories is essential for traffic engineering and ensuring consistent policy across an Autonomous System.
According to Juniper Networks technical documentation, theCommunityattribute is classified as anoptional transitiveattribute. The term "optional" implies that a BGP implementation is not required to support or recognize the attribute. However, because it is "transitive," if a Juniper router receives an update containing a community tag that it does not recognize or has no specific policy for, it must accept the attribute and pass it along to other BGP peers unchanged. This ensures that community-based policies can be signaled across intermediate ASes that may not be configured to act upon those specific tags.
In contrast:
Origin (Option A)andAS Path (Option B)arewell-known mandatoryattributes. Every BGP update must include these, and every BGP-compliant router must recognize them.
MED (Option D)(Multi-Exit Discriminator) is anoptional non-transitiveattribute. If a router receives a MED and advertises that route to an EBGP peer, the MED is typically stripped away (unless specific configurations like path-selection cisco-non-deterministic are used), as it is intended only to influence the immediate neighboring AS.
The Community attribute (defined in RFC 1997) is a powerful tool in Junos OS, often used for tagging routes to trigger specific routing policies, such as setting local preference or identifying the geographic origin of a prefix. By being transitive, it allows for sophisticated administrative control across complex multi-provider environments.

NEW QUESTION 3
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 4
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 5
What information is determined by using the AS path attribute included in the BGP update message? (Choose two.)

A. the origin of a route from IGP or EGP
B. the presence of a routing loop
C. the shortest AS path to reach a prefix
D. the total number of next-hop devices to reach a prefix

Answer: BC

Explanation:
TheAS_PATHattribute is a "well-known mandatory" attribute in BGP, meaning it must be present in every BGP Update message exchanged between External BGP (eBGP) peers. It records the sequence of Autonomous System numbers that a route has traversed. Per Juniper Networks Service Provider documentation, this attribute serves two fundamental purposes:
* 1. Loop Prevention (Option B):
This is the most critical function of the AS_PATH. When a BGP router receives an update from an eBGP peer, it scans the AS_PATH attribute for its own AS number. If the router finds its local AS number already listed in the path, it concludes that the route has already passed through its network and has "looped" back. To prevent an infinite routing loop, the router will immediately discard the update. This mechanism is the cornerstone of BGP's stability as a path-vector protocol.
* 2. Path Selection / Shortest Path Determination (Option C):
BGP uses a complex "tie-breaking" algorithm to select the best path among multiple candidates. One of the highest-ranking criteria in this algorithm (after Weight, Local Preference, and AS_PATH length) is the length of the AS_PATH. A shorter AS_PATH (fewer AS numbers listed) is generally preferred over a longer one, as it typically represents a more direct path through the internet hierarchy.
Why other options are incorrect:
Option A:The "origin" of a route (IGP, EGP, or Incomplete) is determined by theORIGINattribute, which is a separate well-known mandatory attribute.
Option D:BGP does not count individual "next-hop devices" (which would be an IGP metric like hop count in RIP); it only tracks Autonomous Systems. A single AS in the path might contain hundreds of internal routers (next-hops), but BGP only sees it as one "hop" in the AS_PATH.

NEW QUESTION 6
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 7

Exhibit:

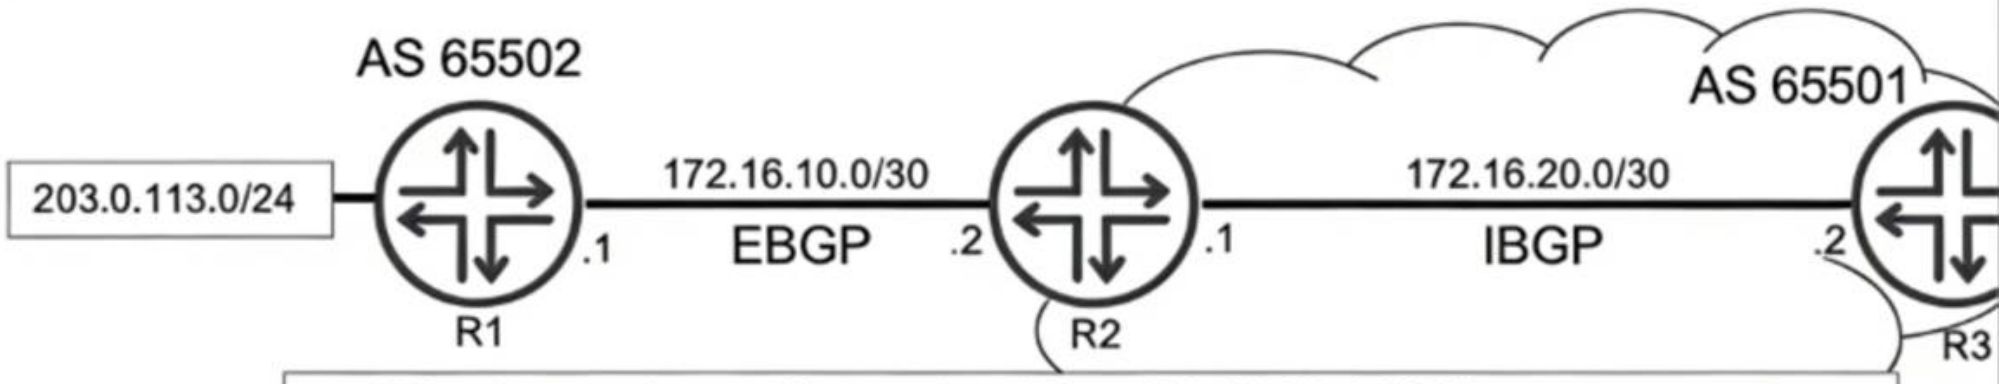


```
user@R3> show route receive-protocol bgp 172.16.20.1 hidden

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 8
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 9
Exhibit:

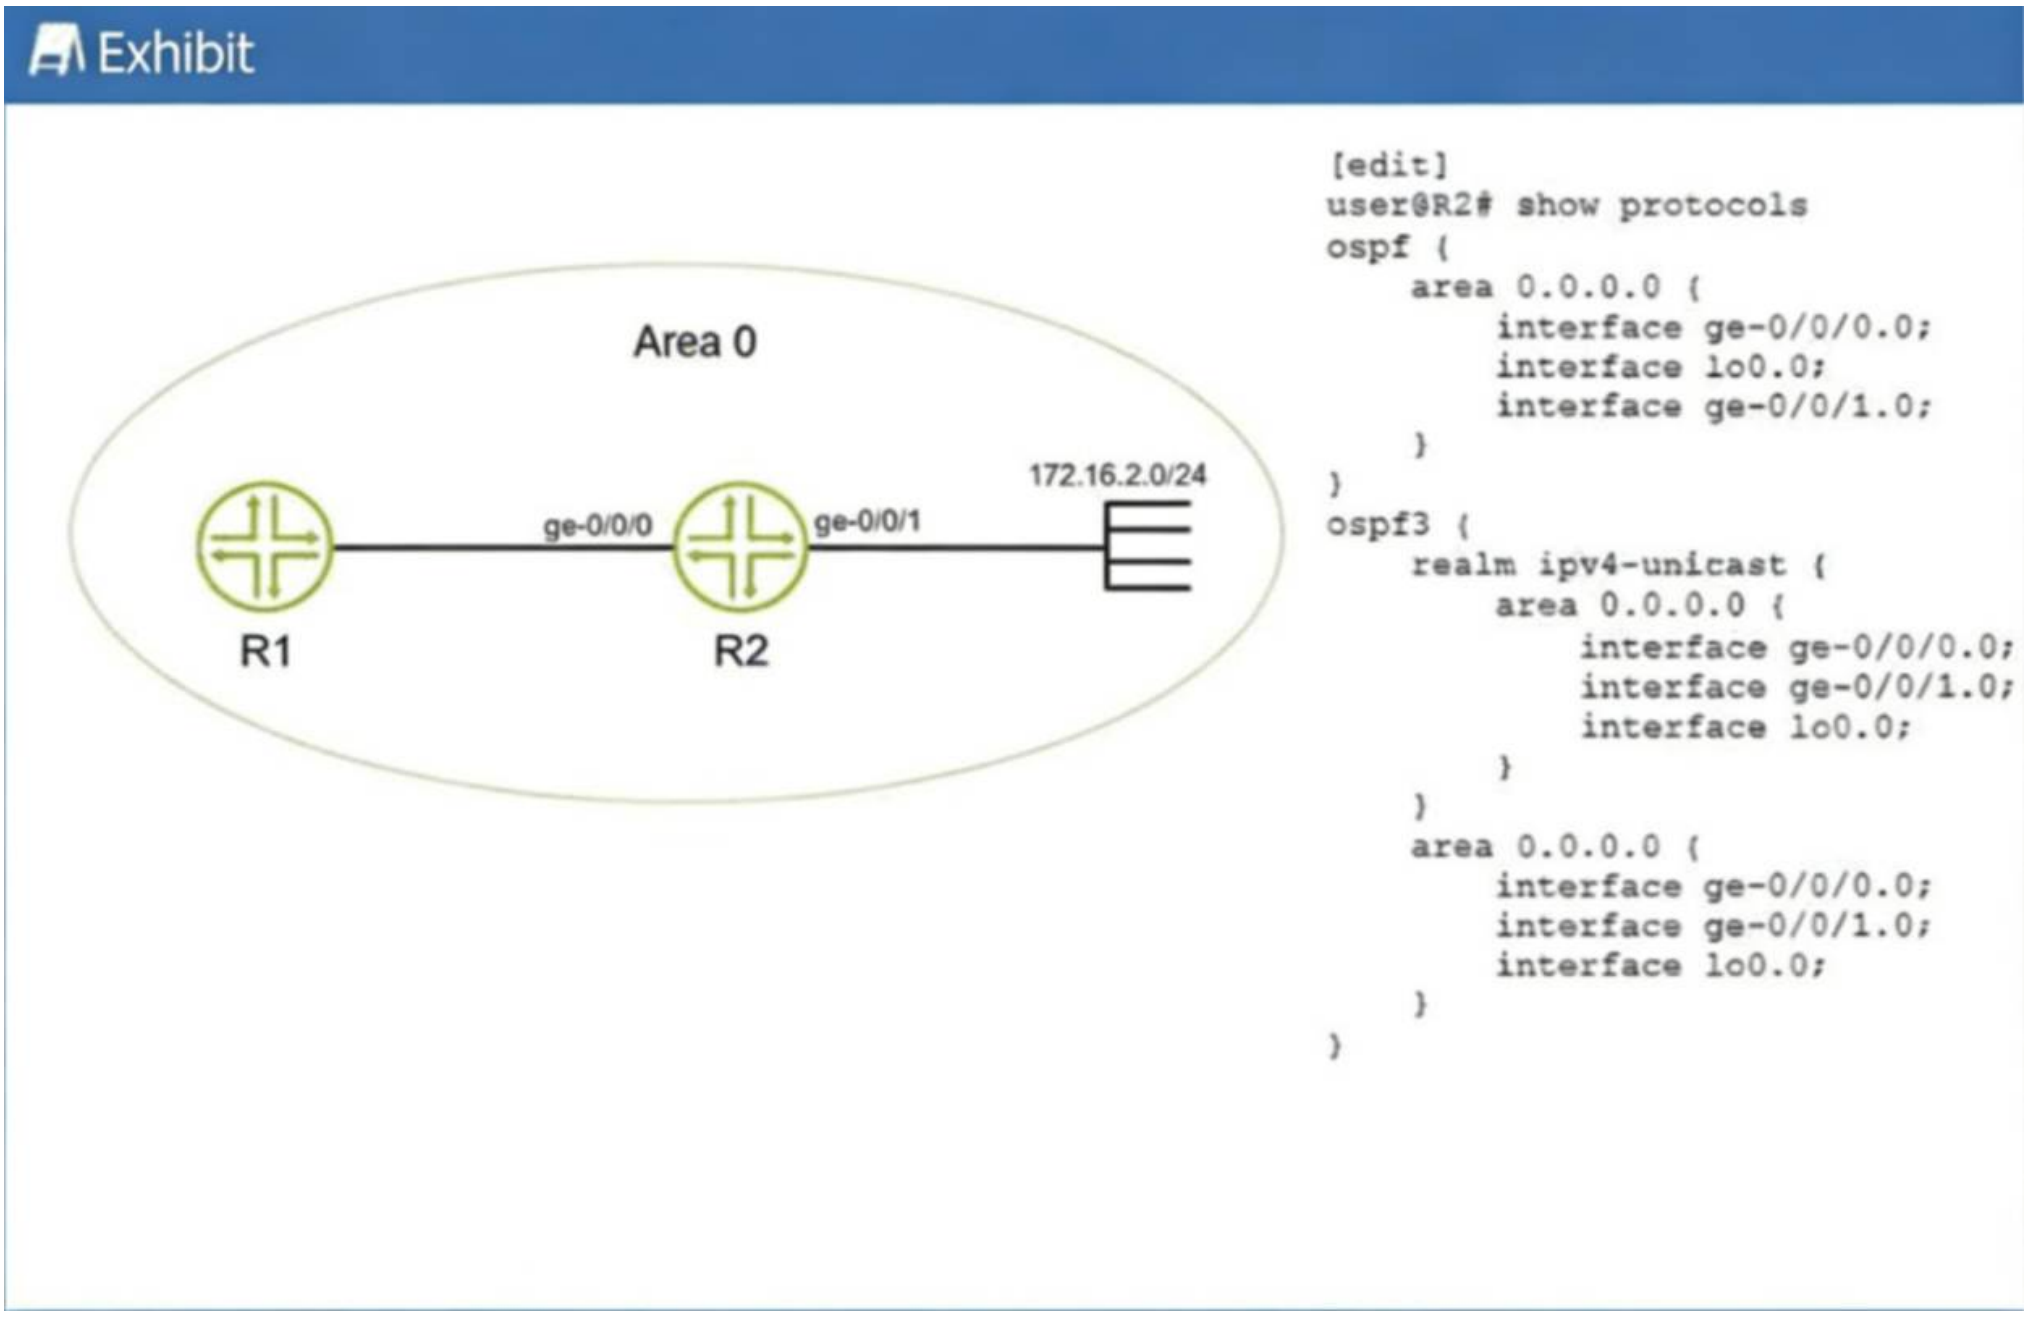

You have configured IPv4 and IPv6 in your network and all OSPF neighbors are established. You apply the configuration shown in the exhibit. Which statement is true in this scenario?

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 10
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 12
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 14
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 18
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 21
Exhibit:

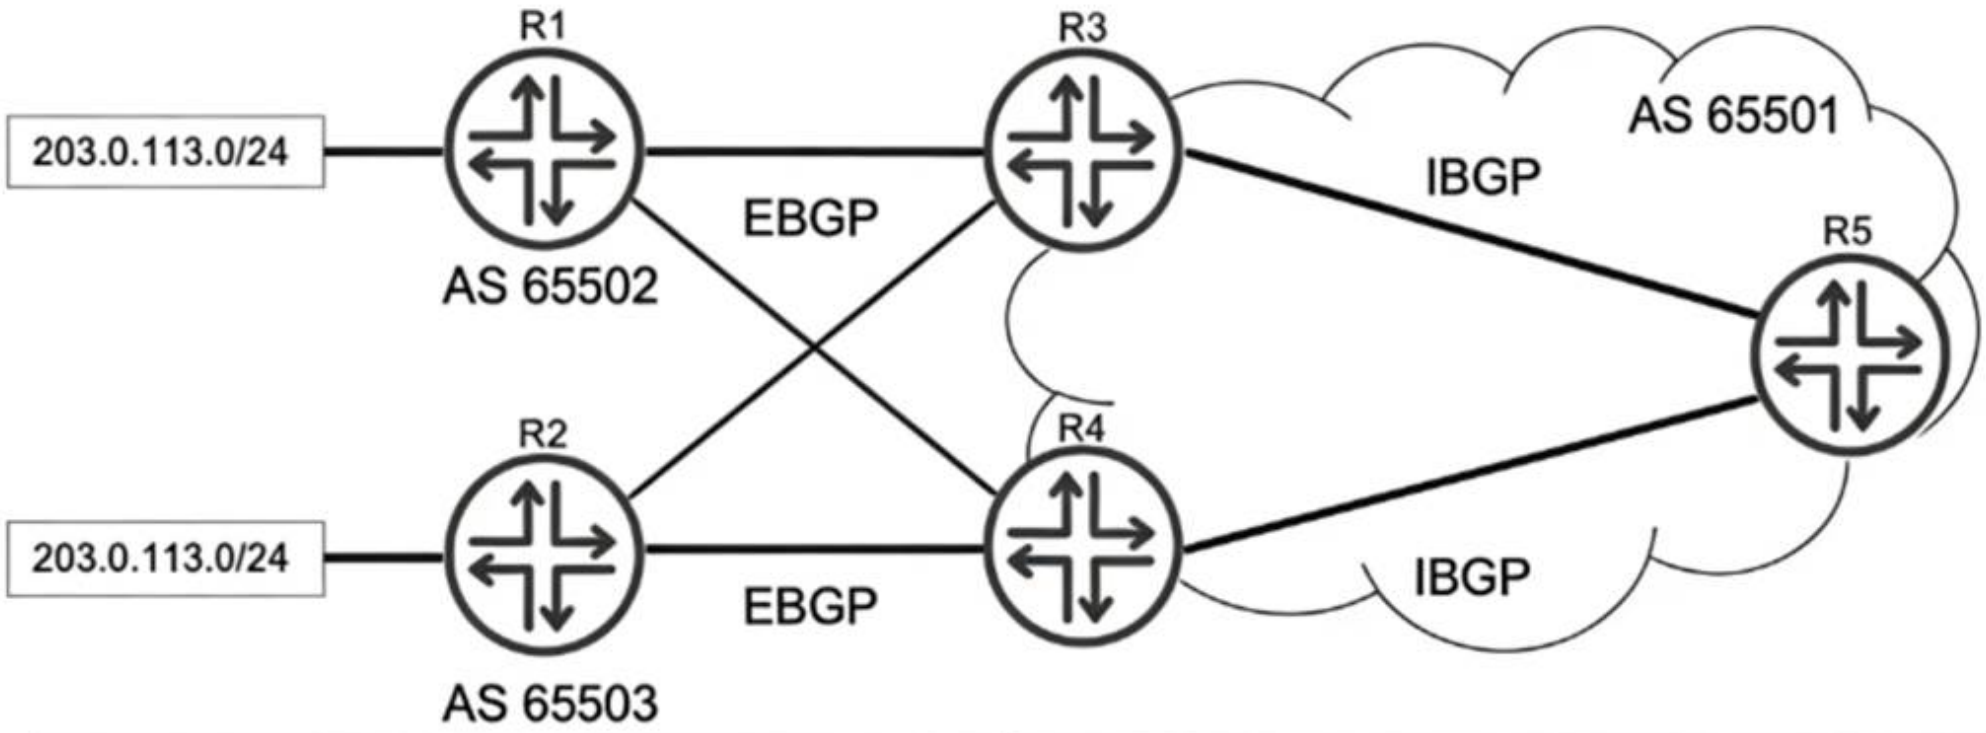


```
user@R3> show configuration policy-options
policy-statement export-to-ibgp {
    from {
        route-filter 203.0.113.0/24 exact;
    }
    then {
        local-preference 150;
        next-hop self;
        accept;
    }
}
```

```
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 23
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 27
In OSPF, which three fields must match between neighbors before forming an adjacency? (Choose three.)

A. router priority
B. hello interval
C. network mask
D. dead interval
E. designated router

Answer: BCD

Explanation:
For OSPF routers to transition from the "Init" state to a full adjacency, they must agree on several parameters exchanged within theirHello packets. If these parameters do not match, the routers will refuse to form a neighbor relationship, a common point of failure in service provider networks.
According to Juniper Networks documentation, the following fields are mandatory matches:
Hello Interval (Option B):The frequency at which Hello packets are sent. Default is 10 seconds on broadcast networks.
Dead Interval (Option D):The time a router waits without receiving a Hello before declaring a neighbor down. Default is 4 times the Hello interval.
Network Mask (Option C):On broadcast and NBMA (Non-Broadcast Multi-Access) segments, the subnet masks must match because OSPF uses the mask to determine the network boundaries for the link-state advertisements.
Area ID:Routers must belong to the same logical OSPF area.
Authentication:If configured, the type and password/key must be identical.
Why other options are incorrect:
Router Priority (Option A):This is used to influence the election of the Designated Router (DR). It doesnotneed to match; in fact, different priorities are often used to ensure a specific router becomes the DR.
Designated Router (Option E):The DR is theresultof an election that happens after the initial Hello exchange. It is not a field that must match beforehand to start the process.
By ensuring the Hello/Dead timers and the Subnet Mask are synchronized, OSPF guarantees a stable and predictable environment for the subsequent exchange of Link-State Advertisements (LSAs).

NEW QUESTION 31
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 35
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 36
.....

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