Juniper

Exam Questions JN0-364

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

NEW QUESTION 1
Which two events cause a static route to be removed from a routing table? (Choose two.)

A. The route is manually removed.
B. The outbound interface becomes unavailable.
C. The route has no traffic for 30 days.
D. Hosts two hops away become unreachable.

Answer: AB

Explanation:
In Junos OS, astatic routeis a manually configured entry in the routing table. Unlike dynamic routes, which have built-in timers and aging mechanisms, static routes are generally "permanent" as long as their conditions for validity are met.
* 1. Manual Removal (Option A):
Since static routes are explicitly defined by the administrator, the most direct way to remove one is through a configuration change. Using the delete routing-options static route <prefix> command followed by a commit will immediately remove the route from the Routing Information Base (RIB).
* 2. Next-Hop Reachability (Option B):
For a static route to be "active" and installed in the forwarding table, itsnext-hop must be reachable. If a static route points to a specific physical interface or an IP address on a local segment, and thatoutbound interface becomes unavailable(e.g., the link goes "Down"), the Junos kernel detects that the next-hop is no longer viable. Consequently, the route is marked as "hidden" or "inactive" and is removed from the active forwarding table to prevent traffic from being black-holed.
Why other options are incorrect:
Aging (Option C):Static routes do not have an expiration timer based on traffic. Even if no packet is sent for years, the route remains as long as the interface is up.
Remote Reachability (Option D):Standard static routes only track the status of the local interface or the immediate next-hop. They do not possess "end-to-end" visibility. If a host two hops away fails, the local router has no way of knowing this via the static route itself. To achieve this level of tracking, features likeRPM (Real-time Performance Monitoring)orBFD (Bidirectional Forwarding Detection)must be linked to the static route.

NEW QUESTION 2
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 3
You must ensure that your routing platform with redundant REs continues to forward packets, even if one RE fails. Which technology would you use to accomplish this task?

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

Answer: D

Explanation:
For Juniper platforms equipped with dualRouting Engines (REs), the fundamental technology required to provide high availability during a hardware or software failure of the primary RE isGraceful Routing Engine Switchover (GRES).
According to Juniper Networks technical documentation, GRES allows the backup RE to stay in a "hot" standby state. When GRES is enabled, the primary RE synchronizes critical state information with the backup RE, specifically thechassis stateand theinterface state. This synchronization includes the Packet Forwarding Engine (PFE) configuration.
When the primary RE fails, the backup RE takes over immediately. Because the PFE (which resides on the line cards) was already synchronized and is not restarted during the switchover, the routercontinues to forward packetsthat are already in flight or part of established flows. This prevents a complete network outage during an RE failover.
Comparison with other options:
NSB (Non-Stop Bridging - Option A):Focuses specifically on maintaining Layer 2 protocol states (like STP) during a switchover.
LAG (Link Aggregation - Option B):Provides redundancy for physical links, not the control plane or the RE.
BFD (Bidirectional Forwarding Detection - Option C):Is a protocol used for rapid detection of link or neighbor failures; it does not protect the RE or maintain forwarding during an internal switchover.

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

NEW QUESTION 4
What is the default export behavior of IS-IS in the Junos OS?

A. to export only IPv6 routes
B. to export only external routes
C. to export nothing
D. to export all learned prefixes

Answer: C

Explanation:
In the Junos OS, routing policy behavior is governed bydefault import and export rulesthat vary significantly between different protocols. ForIS-IS (Intermediate System to Intermediate System), the default export policy is "reject all." This means that, by default, an IS-IS process willexport nothingfrom the routing table into the IS-IS database.
According to Juniper Networks technical documentation, IS-IS automatically advertises its own direct interfaces that are configured under the [edit protocols isis] hierarchy. However, it does not automatically redistribute routes learned from other sources—such asStatic routes,OSPF, orBGP—into the IS-IS domain. This is a safety mechanism designed to prevent accidental routing loops or the flooding of unnecessary prefixes into the link-state database (LSDB), which could impact the stability of the SPF (Shortest Path First) algorithm.
To move routes from the routing table (inet.0) into IS-IS, an administrator must explicitly create arouting policyand apply it as anexport policywithin the IS-IS configuration. For example:
Code snippet set policy-options policy-statement REDIST-STATIC term 1 from protocol static set policy-options policy-statement REDIST-STATIC term 1 then accept
set protocols isis export REDIST-STATIC
Without such a policy, the IS-IS LSPs (Link-State PDUs) will only contain information about the IS-IS enabled interfaces and the reachability of other IS-IS neighbors. This behavior contrasts with protocols like BGP, which has different default rules for exporting active BGP routes to EBGP peers. In the context of IS-IS in a Juniper environment, "export nothing" is the standard operational baseline.

NEW QUESTION 5
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 6
Referring to the exhibit, which protocol would automatically create a full mesh of label-switched paths between MPLS-enabled routers?

Exhibit

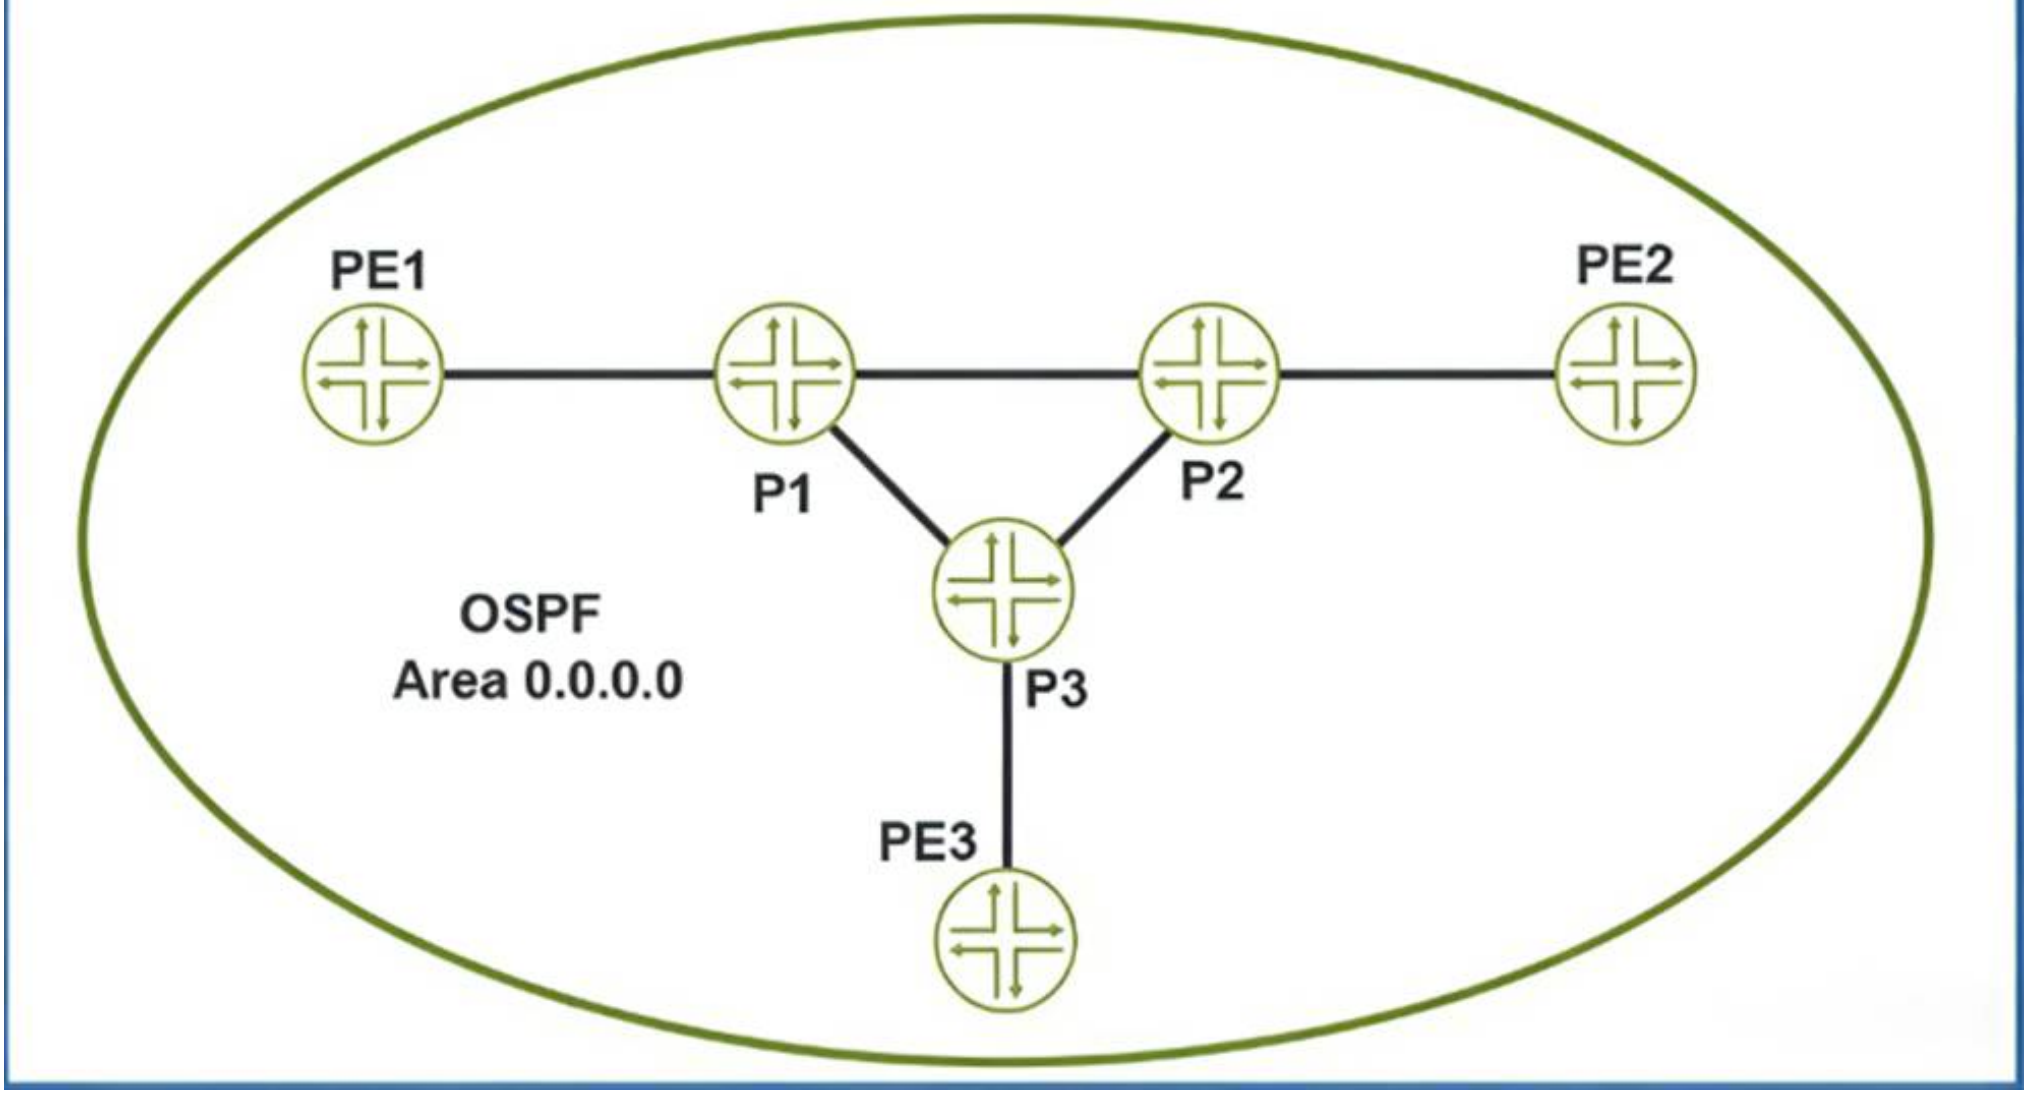

A. LDP
B. BFD
C. BGP
D. RSVP

Answer: A

Explanation:
In Juniper Networks Junos OS, theLabel Distribution Protocol (LDP)is specifically designed to automate the creation of Label Switched Paths (LSPs) based on the information provided by the underlying Interior Gateway Protocol (IGP), such as OSPF or IS-IS. When LDP is enabled on a set of interfaces within an OSPF area (as shown in the exhibit with Area 0.0.0.0), it automatically discovers neighbors and exchanges label mappings for all known unicast routes in the routing table.
The defining characteristic of LDP in this context is its "topology-driven" nature. Unlike RSVP (Resource Reservation Protocol), which typically requires the manual configuration of each LSP ingress point and destination, LDP follows the IGP's shortest path tree to automatically build afull meshof LSPs between all participating routers. This means that every Provider Edge (PE) and Provider (P) router in the exhibit—PE1, PE2, PE3, P1, P2, and P3—will establish label-switched connectivity to every other router without the administrator having to define individual tunnels.
LDP accomplishes this through a downstream-unsolicited label distribution mode by default in Junos. Each router assigns a local label for its loopback address and other prefixes and advertises these to its neighbors. Because every router is performing this action for every reachable prefix in the OSPF domain, a complete fabric of label-switched paths is formed. While RSVP is more robust for traffic engineering and bandwidth reservation, LDP is the preferred protocol for creating a simple, scalable full mesh of LSPs for applications like Layer 3 VPNs or internal BGP tunneling where complex path manipulation is not required. BFD is a failure detection protocol, and BGP is used for service signaling, making LDP the only correct choice for automatic mesh creation.

NEW QUESTION 7
In IS-IS, what would you use to control which external routes are installed in the routing table?

A. export policy
B. import policy
C. route preference
D. interface metric

Answer: B

Explanation:
In Junos OS, the flow of routing information is managed by policies that sit between the protocol's database (the RIB-In/LSDB) and the main routing table (inet.0). Understanding the direction of these policies is critical for correct configuration.
Animport policy (Option B)is used to control the movement of routes from a routing protocol into the routing table. According to Juniper Service Provider documentation, even though IS-IS is a link-state protocol that requires all routers in an area to have an identical Link-State Database (LSDB), animport policycan be used to filter which of those validated routes are actually placed into inet.0 for forwarding. For external routes (routes leaked into IS-IS from other areas or protocols), an import policy allows an administrator to selectively accept or reject prefixes based on specific criteria like prefix-lists or community tags.
It is important to distinguish this from anexport policy (Option A). In Junos, an export policy is used to take routesalready inthe routing table and push themoutto a protocol to be advertised to neighbors. For example, you would use an export policy to redistribute static routes into IS-IS.Route preference (Option C)is a global value used to select between different protocols for the same prefix, and theinterface metric (Option D)is used by the SPF algorithm to calculate the shortest path within the IS-IS database itself. Therefore, to specifically control which learned external routes are "installed" into the forwarding table, theimport policyis the correct tool.

NEW QUESTION 8
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 9
Exhibit:

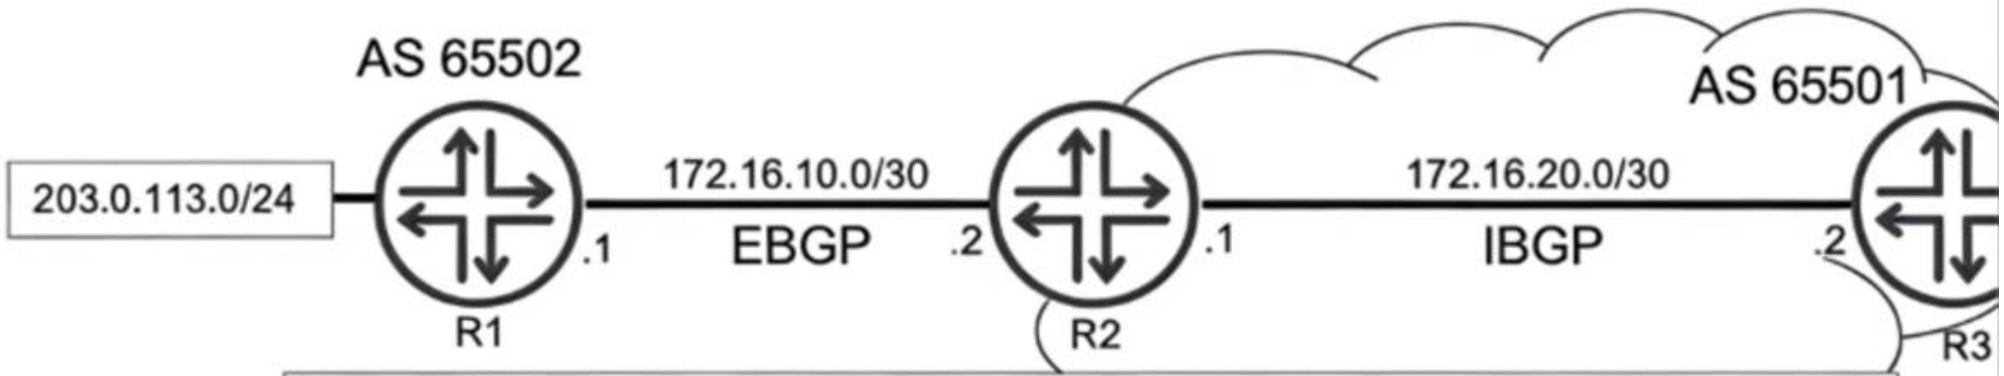


```
user@R3> show route receive-protocol bgp 172.16.20.1 hidden

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 10
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 10
Exhibit:

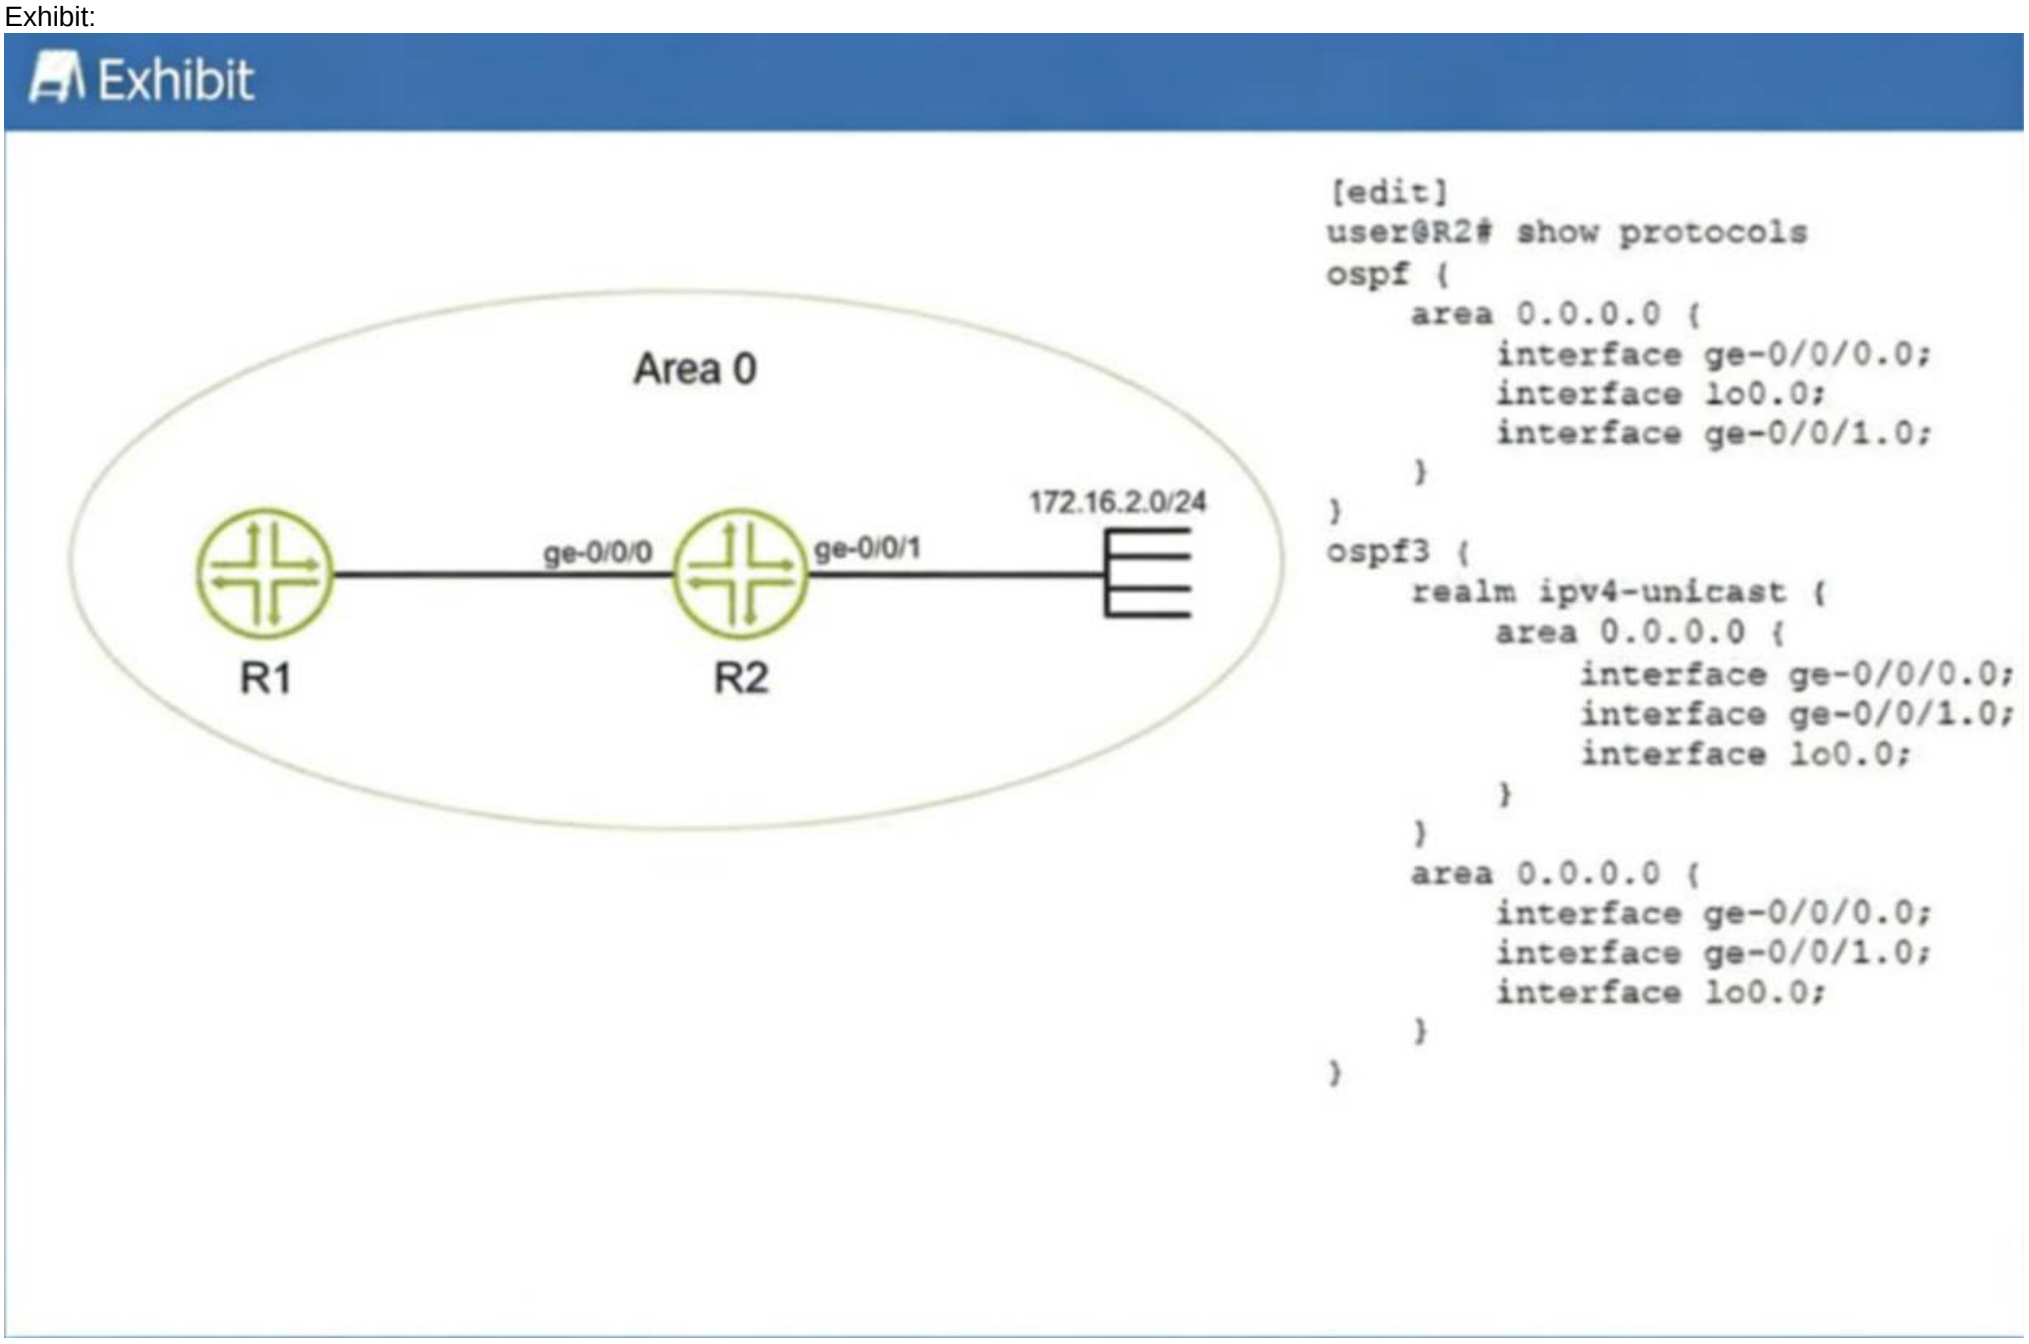

You have configured IPv4 and IPv6 in your network and all OSPF neighbors are established. You apply the configuration shown in the exhibit. Which statement is true in this scenario?

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 12
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 16
In an OSPF network, what is a purpose of a designated router?

A. to assign an OSPF router ID to all routers in the OSPF segment
B. to forward traffic within the configured subnet
C. to reduce OSPF traffic on the OSPF segment
D. to flood routes to all other OSPF devices in the entire domain

Answer: C

Explanation:
On multi-access network segments, such asEthernet, OSPF could potentially face a scalability issue. If every router on a segment formed a full adjacency with every other router, the number of adjacencies would follow the formula $n(n-1)/2$. In a segment with 10 routers, this would result in 45 adjacencies, each generating redundant flooding of Link-State Advertisements (LSAs) and excessive Hello traffic.
To solve this, OSPF elects a Designated Router (DR)and a Backup Designated Router (BDR). According to Juniper Networks documentation, the primary purpose of the DR is to act as a central point of contact for the segment, there by reducing OSPF traffic (Option C).
Instead of every router syncing with every other router, they all form aFull adjacencyonly with the DR and BDR. When a router (a DR-Other) has an update, it sends it to the multicast address224.0.0.6(All DR Routers). The DR then acknowledges the update and floods it to all other routers on the segment using the multicast address224.0.0.5(All OSPF Routers). This "hub-and-spoke" signaling model within the local segment significantly minimizes the bandwidth consumed by protocol overhead and reduces the CPU load on the participating routers.
It is important to note that the DR's scope is limited to the local segment; it does not "assign IDs" (Option A) nor does it flood routes to the "entire domain" (Option D), as that is the responsibility of individual routers within their respective areas.

NEW QUESTION 21
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 23
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 27
.....

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