Juniper

Exam Questions JN0-364

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

NEW QUESTION 1
How are routing loops prevented in internal BGP networks?

A. Internal BGP routes are never readvertised to other internal BGP neighbors.
B. External BGP routes are never readvertised to other external BGP neighbors.
C. External BGP routes are never readvertised to other internal BGP neighbors.
D. Internal BGP routes are never readvertised to other external BGP neighbors.

Answer: A

Explanation:
The prevention of routing loops within an Autonomous System (AS) is handled differently than loop prevention between ASes. While External BGP (EBGP) uses the AS_PATH attribute to detect loops,Internal BGP (IBGP)does not modify the AS_PATH. Therefore, a different mechanism is required to ensure that a route does not circulate infinitely inside the network.
This mechanism is known as theIBGP Split Horizon rule. According to Juniper Networks documentation and the BGP standard (RFC 4271), a BGP speakermust not advertise a route learned via an IBGP peer to any other IBGP peer. In simpler terms, "what is learned internally, stays local." This rule ensures that a route only travels one "hop" inside the AS—from the router that learned it from an external source to all other internal routers.
Because of this rule, IBGP routers do not naturally propagate routes through each other. This creates a requirement for afull meshof IBGP sessions, where every BGP-speaking router in the AS must have a direct peering session with every other BGP-speaking router. To mitigate the scaling issues of a full mesh in large service provider networks, architects useRoute ReflectorsorConfederations, which are authorized exceptions to the Split Horizon rule.
Option B is incorrect because EBGP peersdoadvertise EBGP routes to other EBGP peers (this is how the internet works). Option C is incorrect because EBGP-learned routesmustbe sent to IBGP peers so the internal network knows how to reach the outside world. Option D is incorrect because internal routesmustbe sent to external peers to advertise your network to the internet.

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

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

Answer: AB

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

NEW QUESTION 3
You are monitoring OSPF on a router and notice frequent state changes between Full and Down. Which condition would cause this behavior?

A. physical interface flapping
B. route preference mismatch
C. area ID mismatch
D. MTU mismatch

Answer: A

Explanation:
When troubleshooting OSPF in a service provider environment, distinguishing between "stuck" adjacencies and "flapping" adjacencies is the first step. A session that transitions frequently betweenFullandDown indicates that the relationship can be established successfully (meaning parameters match), but it cannot be maintained.
According to Juniper Networks documentation, the most common cause for a session to drop from Full to Down is the expiration of theDead Interval. If a router does not receive a Hello packet within the Dead Interval (usually 40 seconds), it tears down the adjacency. Aphysical interface flapping (Option A)is the primary trigger for this. If the physical link or the underlying transport (like a leased line or a microwave link) goes down even momentarily, the OSPF process immediately detects the interface failure, flushes the
neighbors, and moves the state to Down. As soon as the interface comes back up, the routers perform the Hello exchange and reach the Full state again, creating the flapping cycle.
Analysis of other options:
MTU Mismatch (Option D):This typically causes the adjacency to get "stuck" in theExchangeor ExStartstate. The routers can exchange small Hello packets, but when they try to send larger Database Description (DBD) packets that exceed the MTU, the packets are dropped, preventing the session from ever reaching "Full."
Area ID Mismatch (Option C):This prevents the adjacency from even reaching theInitstate; the routers will never form a neighbor relationship.
Route Preference (Option B):This affects which route is chosen for the forwarding table but has no impact on the OSPF neighbor state machine itself.

NEW QUESTION 4

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

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

Answer: AD

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

NEW QUESTION 5
Exhibit:

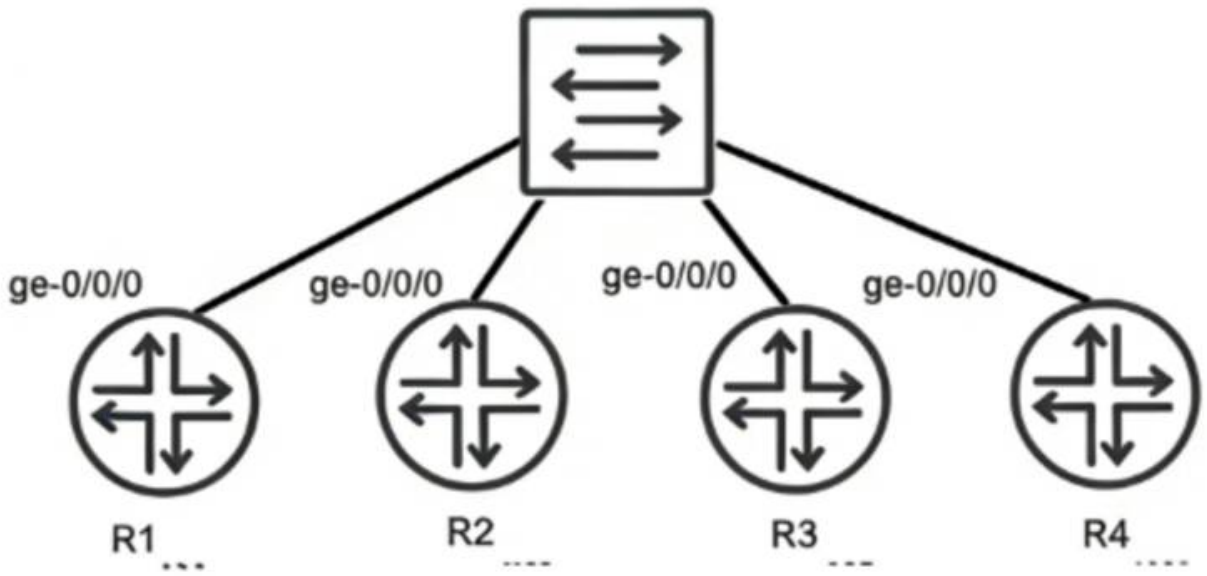


```
user@R1> show configuration routing-options
router-id 192.168.1.1;

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

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

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

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

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

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

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

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

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

Answer: D

Explanation:
In OSPF networks, when multiple routers are connected to a shared multi-access broadcast segment (like an Ethernet switch), they undergo an election process to select aDesignated Router (DR)and aBackup Designated Router (BDR). This mechanism is essential for reducing the number of adjacencies and limiting the

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

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

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

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

NEW QUESTION 7
Exhibit:

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

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

10.16.2.0/23 (1 entry, 1 announced)

*Aggregate Preference: 130

Next hop type: Reject

Address: 0x8f3fd44

Next-hop reference count: 2

State:

Age: 1:39:21

Task: Aggregate

Announcement bits (1): 0-KRT

AS path: I (LocalAgg)

Flags: Depth: 0 Active

AS path list:

AS path: I Refcount: 2

Contributing Routes (2):

10.16.2.0/24 proto Direct

10.16.3.0/24 proto Direct
```

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

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

Answer: A

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

NEW QUESTION 8
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 9
You are configuring LDP in a service provider network. After enabling LDP on core interfaces, you notice that labels are being advertised for every loopback IPv4 address that is in your IGP. Which label distribution mode is being used in this scenario?
A. conservative retention
B. ordered control
C. downstream unsolicited
D. downstream on demand

Answer: C

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

NEW QUESTION 10
Exhibit:

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

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

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

Answer: A

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

NEW QUESTION 10
Which IS-IS packet type will establish and maintain neighbor relationships?

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

Answer: B

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

NEW QUESTION 11
Exhibit:

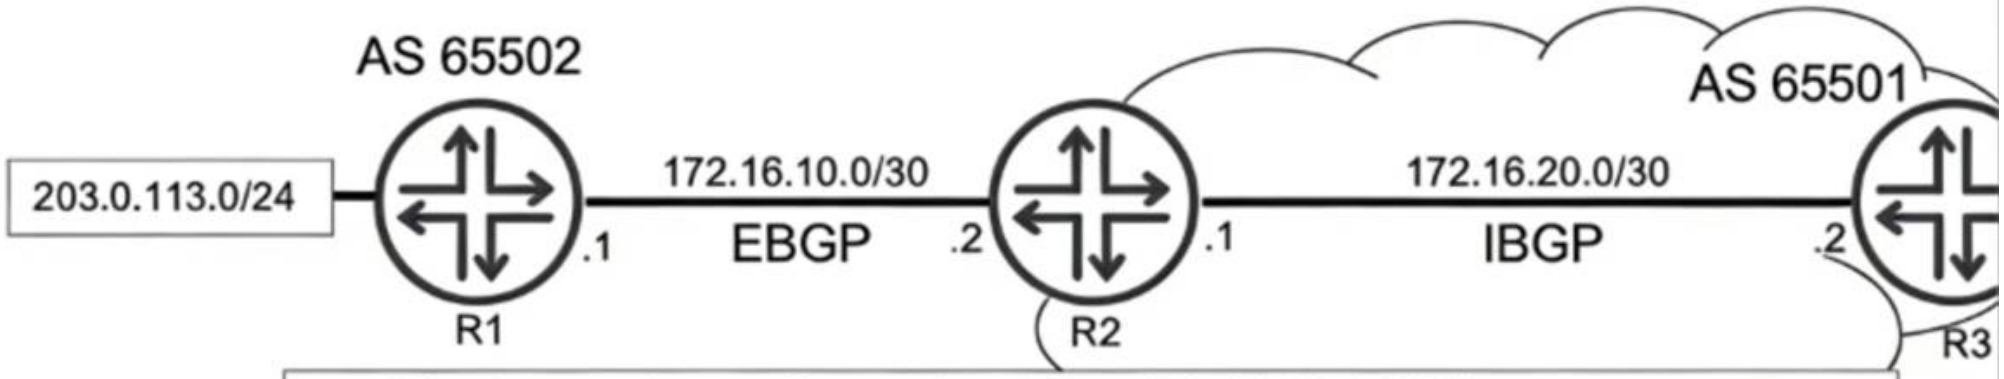


```
user@R3> show route receive-protocol bgp 172.16.20.1 hidden

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 12
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 15
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 16
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 18
Exhibit:

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

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

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

Answer: BD

Explanation:
In the provided exhibit, the output of the command show spanning-tree interface for switch1 reveals critical details about the Spanning Tree Protocol (STP) operational state.
The first correct statement is thatthe switch1 device is the root bridge(Option B). This is determined by comparing the "Port ID" column with the "Designated port ID" column, as well as checking the "Designated bridge ID". In the exhibit, for every interface listed (from ge-0/0/6.0 to ge-0/0/13.0), the Port ID and the Designated port ID are identical. Furthermore, every port is in the "FWD" (Forwarding) state with the "DESG" (Designated) role. In a Spanning Tree topology, the root bridge is the only device where all active participating interfaces serve as designated ports, as it has no need for a "Root" port role (which points toward a root bridge).
The second correct statement is thatthe bridge priority for switch1 is 32k(Option D). Looking at the "Designated bridge ID" column, we see the value 32768.0019e2552481. In Junos and general networking standards, the Bridge ID is composed of a bridge priority and the device's MAC address. The default priority for most Spanning Tree variants (STP, RSTP, MSTP) is 32,768, which is commonly referred to in shorthand as "32k".
Regarding the incorrect options:
Option A:There is no evidence of VSTP (VLAN Spanning Tree Protocol); the output shows "instance 0," which is typical for IEEE standard RSTP or STP.
Option C:The Port IDs for ge-0/0/8, ge-0/0/9, and ge-0/0/11 all start with "32" (e.g., 32:521), whereas the default port priority is typically 128 (as seen in ge-0/0/6.0 with 128:519). This indicates that the interface priorities for these specific ports have been manually tuned to a non-default value.

NEW QUESTION 19
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 20
Which two statements are correct about TLVs in IS-IS? (Choose two.)

A. LSPs can only contain one TLV.
B. TLVs only support encoding IPv4 routing information.
C. TLVs allow flexible encoding of routing information.
D. LSPs can contain multiple TLVs.

Answer: CD

Explanation:
In the IS-IS protocol,TLVs (Type, Length, Value)are the fundamental building blocks used to carry information withinLink-State PDUs (LSPs). Unlike some other protocols that have a fixed, rigid packet format, IS-IS was designed from the ground up to be modular and extensible. This extensibility is achieved through the use of TLVs, which allow the protocol to carry different types of data without requiring changes to the core protocol state machine.
According to Juniper Networks technical documentation,TLVs allow flexible encoding of routing information (Option C). Each TLV specifies the "Type" of information it carries (such as neighbor information or IP reachability), the "Length" of that information, and the "Value" (the actual data). This architecture is what allowed IS-IS to easily support IPv6 by simply adding new TLVs (like TLV 236 for IPv6 reachability) without redesigning the protocol. It also supports Traffic Engineering (TE) extensions used in MPLS environments by adding TLVs that describe link bandwidth and administrative groups.
Furthermore, a singleLSP can contain multiple TLVs (Option D). When a Juniper router generates an LSP, it packs all the necessary information—such as the router's area addresses, its neighbors, and its local interface prefixes—into various TLVs and places them into a single PDU. If the amount of information exceeds the Maximum Transmission Unit (MTU) of the interface, the router will generate additional LSPs (fragmented LSPs) to carry the remaining TLVs.
Options A and B are incorrect because restricting an LSP to a single TLV would make the protocol incredibly inefficient, and the very nature of IS-IS is its ability to support multiple network layer protocols (not just IPv4) through its agnostic TLV-based transport.

NEW QUESTION 23
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 27
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 29
A service provider is onboarding a new enterprise customer that operates multiple branch offices, each with its own set of VLANs. The customer requires transparent Layer 2 connectivity between sites while maintaining separation of internal VLANs. The provider must also ensure that customer VLAN identifiers do not conflict with other customers on the shared infrastructure. Which solution would provide the desired results?

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

Answer: A

Explanation:
In a service provider environment,Q-in-Q tunneling(also known as 802.1ad or double-tagging) is the standard solution for transporting multiple customer VLANs over a shared provider backbone while maintaining total separation.

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

NEW QUESTION 31
You are designing an MPLS network and want to ensure that traffic traverses an LSP between PE routers that follow an explicit path through the core. Which protocol would accomplish this task?

A. BGP
B. RSVP
C. IS-IS
D. LDP

Answer: B

Explanation:
In a Juniper Networks MPLS environment, the selection of a signaling protocol depends heavily on the requirement for traffic engineering and path control. To satisfy the requirement for anexplicit path—where the network architect defines specific hop-by-hop routers that the traffic must traverse—theResource Reservation Protocol (RSVP)is the necessary choice.
According to Juniper documentation, RSVP (specifically RSVP-TE) supports the use ofExplicit Route Objects (EROs). When you configure an LSP in Junos OS, you can define a path consisting of a series of IP addresses (strict or loose hops). RSVP then signals the LSP along that exact sequence of routers, reserving resources and establishing labels as it goes. This allows for precise control over the network's traffic patterns, enabling administrators to steer traffic away from congested links or toward specific high-bandwidth paths.
In contrast,LDP (Label Distribution Protocol)(Option D) is a "best-effort" signaling protocol. LDP strictly follows the Interior Gateway Protocol (IGP) shortest path. It does not support explicit paths or traffic engineering constraints; it simply builds a "mesh" of labels based on the existing routing table.IS-IS(Option C) is an IGP used to populate the routing table and TED but does not signal labels.BGP(Option A) is used for service delivery (like L3VPNs) but relies on an underlying transport LSP (built by RSVP or LDP) to reach its next hop. Therefore, only RSVP provides the mechanism for explicit path manipulation.

NEW QUESTION 35
You are configuring BGP for IPv6 operations. In this scenario, which two statements are correct? (Choose two.)

A. The Autonomous System Number (ASN) must be a 64-bit value.
B. The router ID uses a 128-bit identifier value.
C. The router ID uses a 32-bit identifier value.
D. The Autonomous System Number (ASN) can be either a 32-bit or 64-bit value.

Answer: CD

Explanation:
When implementingMultiprotocol BGP (MP-BGP)for IPv6, several architectural constants remain consistent with the original BGP design, while others have evolved to accommodate larger network scales.
Router ID (Option C):
A critical point in Juniper's Service Provider documentation is that theBGP Router IDremains a32-bit value, even when the protocol is carrying 128-bit IPv6 prefixes. The Router ID is typically represented in dotted-quad notation (e.g., 192.168.1.1). In an IPv6-only environment, a Juniper router cannot automatically derive this ID from an interface address, so it must be manually defined under [edit routing-options]. This 32-bit ID is essential for BGP tie-breaking and loop prevention within the AS.
Autonomous System Number (Option D):
TheAutonomous System Number (ASN)was originally a 16-bit value (0 to 65535). However, to address the exhaustion of available ASNs, the standard was extended to32-bit ASNs(documented in RFC 6793). In Junos OS, you can configure BGP using either the older 16-bit format or the newer 32-bit format (often represented in "asplain" or "asdot" notation). While the question mentions a 64-bit value, there is currently no standard for a 64-bit ASN in BGP; the transition from 16-bit to 32-bit satisfies current global scalability needs. Therefore, Option D is the most accurate within the context of current networking standards, as it acknowledges the coexistence of different ASN lengths.

NEW QUESTION 39
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 43
Exhibit:

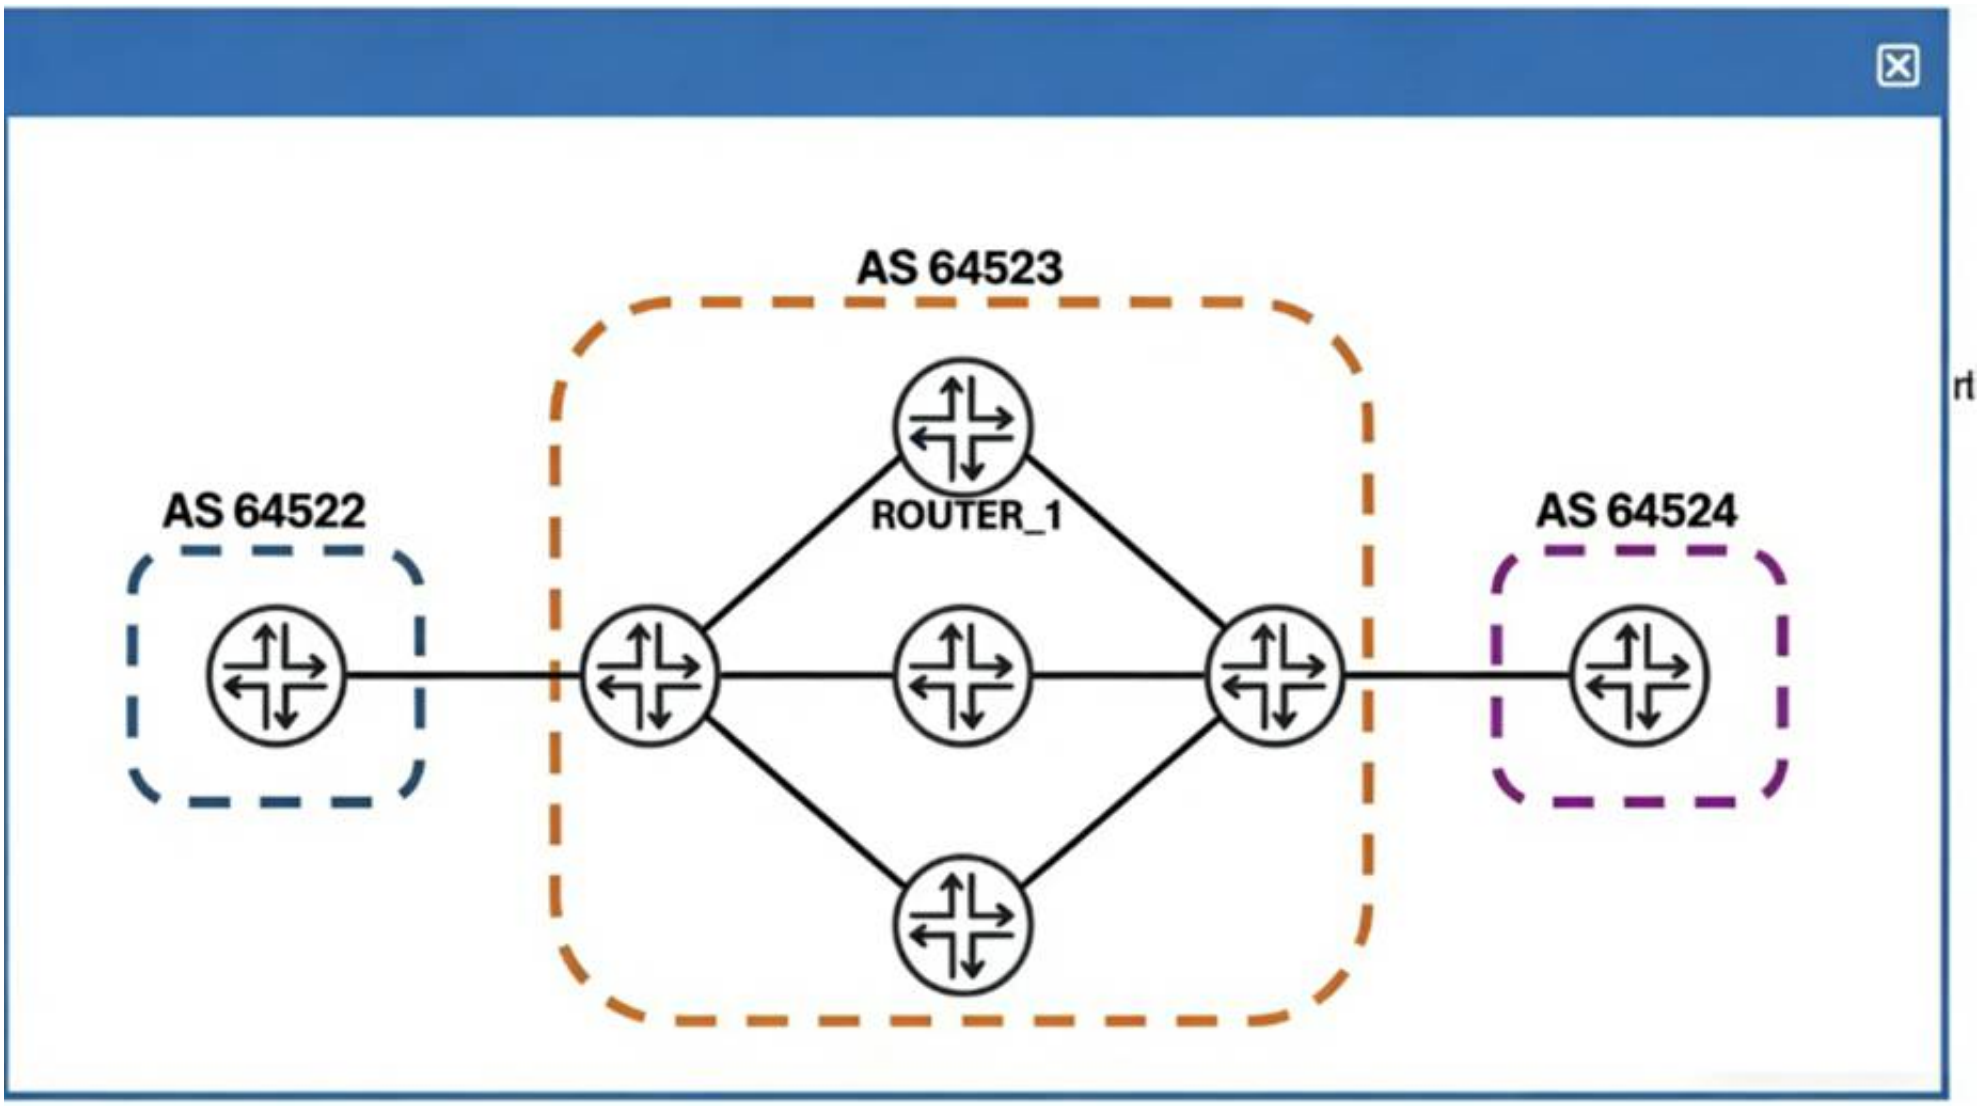

You must configure the router called ROUTER_1 to take all valid prefixes learned from internal BGP peers in AS 64523, and then re-advertise them to other internal BGP peers in the same autonomous system.
Referring to the exhibit, which configuration must you deploy on ROUTER_1 to accomplish this task?

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

Answer: B

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

NEW QUESTION 47
.....

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