Juniper

Exam Questions JN0-364

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

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

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

Answer: CD

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

NEW QUESTION 2
Which two 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 3
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 4
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 5
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 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
What prevents routing loops in a single-area OSPF network?

A. The Dijkstra algorithm
B. Routing policies
C. The Bellman-Ford algorithm
D. Forwarding policies

Answer: A

Explanation:
In OSPF, loop prevention within a single area is achieved through the fundamental nature of its link-state architecture. Unlike distance-vector protocols that rely on "routing by rumor," OSPF ensures that every router within an area maintains an identicalLink-State Database (LSDB). This database acts as a complete map of the network topology.
Once the LSDB is synchronized, each router independently executes theShortest Path First (SPF) algorithm, which is formally known as theDijkstra algorithm. This mathematical process treats the local router as the "root" of a tree and calculates the shortest path to every other node (router) and prefix in the area based on the cumulative interface costs. Because every router uses the same synchronized map (the LSDB) and the same deterministic algorithm, they all arrive at a consistent, loop-free view of the best paths.
According to Juniper Networks technical documentation, the Dijkstra algorithm is superior to theBellman-Ford algorithm(used by distance-vector protocols like RIP) in this regard. Bellman-Ford is susceptible to "count-to-infinity" problems and loops because routers only know the distance and direction to a destination provided by their neighbors, rather than the full topology. In OSPF, even if a link fails, the updated Link-State Advertisement (LSA) is flooded rapidly, and the Dijkstra algorithm is re-run to find a new loop-free path.Routing policies(Option B) are used to manipulate path selection or filter routes but are not the primary mechanism for fundamental loop prevention in OSPF. Similarly,forwarding policies(Option D) govern how traffic is handled at the data plane level rather than determining the control plane's loop-free topology.

NEW QUESTION 8
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 9
Exhibit:

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

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

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

Answer: A

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

NEW QUESTION 10
Exhibit:

Exhibit

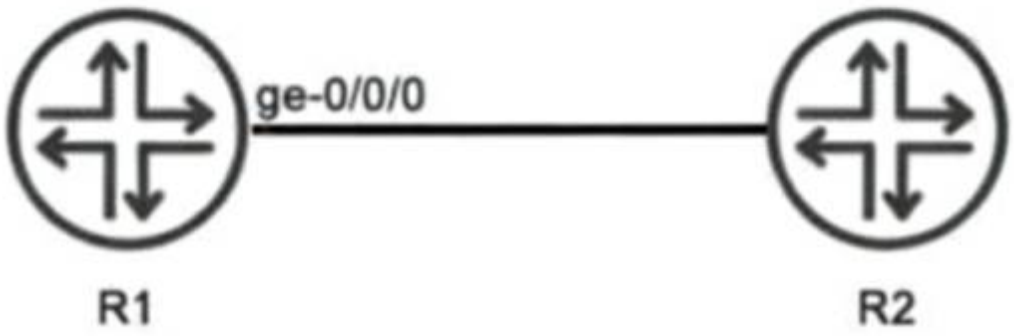


```
user@R1> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/0.0            2   0x1 Disabled          Point to Point        100/100
lo0.0                 2   0x1 Passive           Passive                 0/0
```

Referring to the exhibit, R1 and R2 are configured to run IS-IS. The IS-IS adjacency between R1 and R2 is up. What does the output of the show isis interface command tell you about R1?

A. R1 is not configured to use wide metrics.
B. R1 only forms a Level 2 adjacency with R2.
C. R1 advertises a Level 1 metric of 100 and a Level 2 metric of 100 toward R2 in its link-state PDU.
D. R1 sends Level 1 hello PDUs to R2.

Answer: B

Explanation:
In theIS-IS (Intermediate System to Intermediate System)protocol as implemented in Junos OS, routers can operate at two hierarchical levels:Level 1 (L1)for intra-area routing andLevel 2 (L2)for inter-area backbone routing. By default, a Juniper router and its interfaces are configured to act asLevel 1/2, meaning they will attempt to form adjacencies at both levels simultaneously.
According to Juniper Networks technical documentation, the show isis interface command provides a granular view of how the protocol is interacting with specific local links. In the provided exhibit, we must examine the L (Level)column and theDR (Designated Router)status columns to understand R1's operational state.
Level Configuration:Under the L column for both the physical interface ge-0/0/0.0 and the loopback lo0.0, the value is strictly2. This indicates that these interfaces have been explicitly configured to operate only at Level 2.
Adjacency Capabilities:For the interface ge-0/0/0.0, the Level 1 DR field is marked asDisabled. This confirms that R1 is not participating in Level 1 operations on this link; it will not transmit Level 1 Hello PDUs, nor will it listen for them. Consequently, R1 is incapable of forming a Level 1 adjacency with R2 on this segment.
Metric Implications:The exhibit shows an L1/L2 Metric of100/100. In Junos, "narrow" metrics (the default) are limited to a maximum value of 63 per interface. A metric of 100 indicates thatwide metrics (wide-metrics-only) have been enabled. Therefore, option A is incorrect because the routerisusing wide metrics.
Since the prompt states the adjacency is "up," and the interface is restricted to Level 2, we can conclude that R1 only forms a Level 2 adjacency with R2 (Option B). Even though an L1 metric of 100 is displayed in the table as a configured value, it is not actually "advertised" in a Link-State PDU because the Level 1 protocol is disabled on that interface.

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 15
Exhibit:

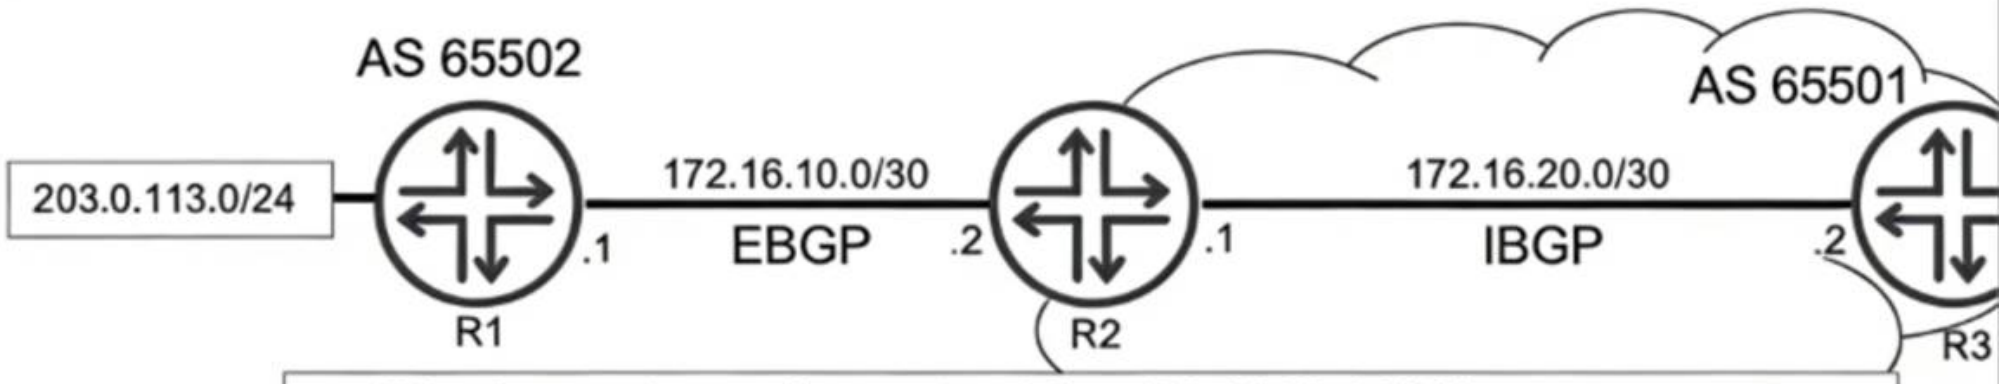


```
user@R3> show route receive-protocol bgp 172.16.20.1 hidden

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 20
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 24
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 27
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 30
You are designing a high availability solution for a Juniper router with dual Routing Engines (RE). You want to ensure that the routing protocol state is preserved during an RE switchover. You have already enabled graceful Routing Engine switchover (GRES) and you want to avoid relying on helper routers to maintain the routing protocol state. In this scenario, which feature would accomplish this behavior?

A. non-stop active bridging
B. bidirectional forwarding detection
C. graceful restart
D. non-stop active routing

Answer: D

Explanation:
When designing High Availability (HA) for Juniper Service Provider routers, understanding the interaction between the control plane and data plane is vital. The user has already enabledGraceful Routing Engine Switchover (GRES), which synchronizes the interface and kernel state between the primary and backup Routing Engines (REs). However, GRES by itself does not preserve the routing protocol state (like OSPF adjacencies or BGP sessions).
To achieve the preservation of the routing protocol state without relying on external "helper" routers, you must implementNon-Stop Active Routing (NSR).
According to Juniper Networks documentation, NSR uses the infrastructure provided by GRES to also synchronize the routing protocol process (rpd) information.

Under NSR, the backup RE maintains a "hot" standby state of all routing protocols. If the primary RE fails, the backup RE takes over immediately. Because it already possesses the full routing table and peer session states, the peering neighbors are unaware that a switchover occurred. No protocol adjacency resets occur, and traffic continues to flow uninterrupted.
It is crucial to differentiate NSR fromGraceful Restart (Option C). While Graceful Restart also aims to maintain traffic flow during a switchover, itdoesrequire help from neighboring routers (known as "helper mode"). If the neighbors do not support or are not configured for Graceful Restart, the sessions will drop. Since the user explicitly stated they want to "avoid relying on helper routers," Graceful Restart is not the correct solution.
Non-stop Active Bridging (Option A)provides a similar "hitless" failover but specifically for Layer 2 environments (STP/VLANs) rather than Layer 3 routing protocols.BFD (Option B)is a failure detection protocol used to speed up convergence but does not preserve state during an RE failover; in fact, without NSR, BFD would likely trigger a faster teardown of the session during a switchover. Therefore,NSRis the only feature that meets the requirement for independent control-plane preservation.

NEW QUESTION 32
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 34
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 36
Exhibit:

```
user@R2> show route 198.51.100.1

inet.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)

Restart Complete

+ = Active Route, - = Last Active, * = Both

198.51.100.1/32 *[Static/5] 5d 21:02:26

> to 203.0.113.65 via ge-0/0/3.0

user@R2> show route 172.20.110.0/24

inet.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)

Restart Complete

+ = Active Route, - = Last Active,

* = Both

172.20.110.0/24 *[Static/5] 10:43:01

> via gr-0/0/0.0
```

Referring to the exhibit, traffic destined to which network will be sent through the tunnel?

A. 172.20.110.0/24
B. 203.0.113.65
C. 0.0.0.0/0
D. 198.51.100.1/32

Answer: A

Explanation:
Explanation
To determine which traffic is being sent through a tunnel in a Junos OS environment, an administrator must analyze the routing table output for the exit interface associated with each destination prefix. The provided exhibit shows the results of the show route command on routerR2for two specific destination networks.
In the first output, the destination198.51.100.1/32is an active static route. The next-hop information specifies that traffic for this address is sent to the gateway 203.0.113.65 via the interfacege-0/0/3.0. According to Juniper Networks interface naming conventions, the prefixge-denotes aGigabit Ethernetinterface, which represents a standard physical connection. Therefore, this traffic does not traverse a tunnel.
In the second output, the destination172.20.110.0/24is also an active static route. However, the next-hop for this network is listed asvia gr-0/0/0.0. In the Junos operating system, thegr-prefix explicitly identifies a Generic Routing Encapsulation (GRE) tunnel interface. GRE is a widely used protocol in service provider networks to encapsulate various network layer protocols over an IP backbone, effectively creating a virtual point-to-point link. Because the routing table has installed the route for 172.20.110.0/24 specifically via the gr- interface, all traffic destined for this network will be encapsulated and sent through the tunnel.
The other choices are incorrect for the following reasons:
* 203.0.113.65 (Option B):This is the next-hop IP address for the physical Gigabit Ethernet path; it is not a destination network directed to a tunnel.
* 0.0.0.0/0 (Option C):There is no information in the exhibit regarding a default route.
* 198.51.100.1/32 (Option D):As identified by thege-interface prefix in the exhibit, traffic for this destination is sent via a physical Ethernet link.

NEW QUESTION 41
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 44
.....

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