JN0-364 Dumps

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

https://www.certleader.com/JN0-364-dumps.html

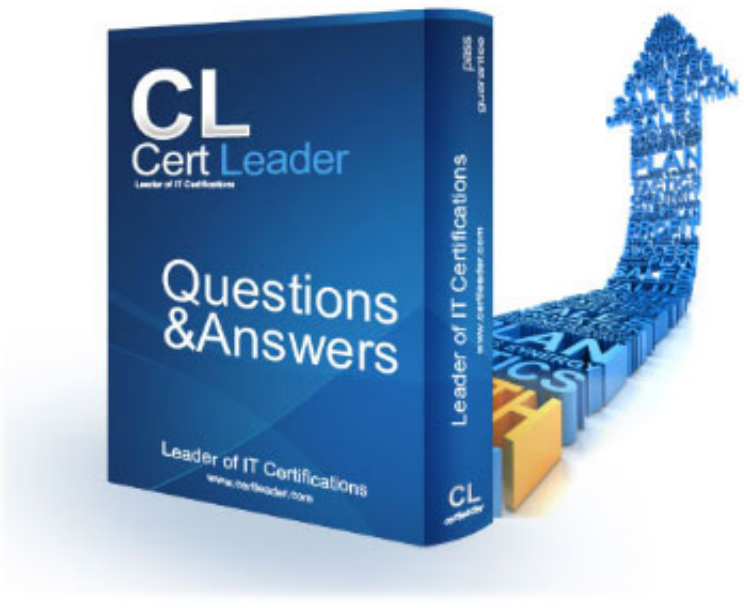

NEW QUESTION 1
You must ensure that your routing platform with redundant REs continues to forward packets, even if one RE fails. Which technology would you use to accomplish this task?

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 2
Exhibit:

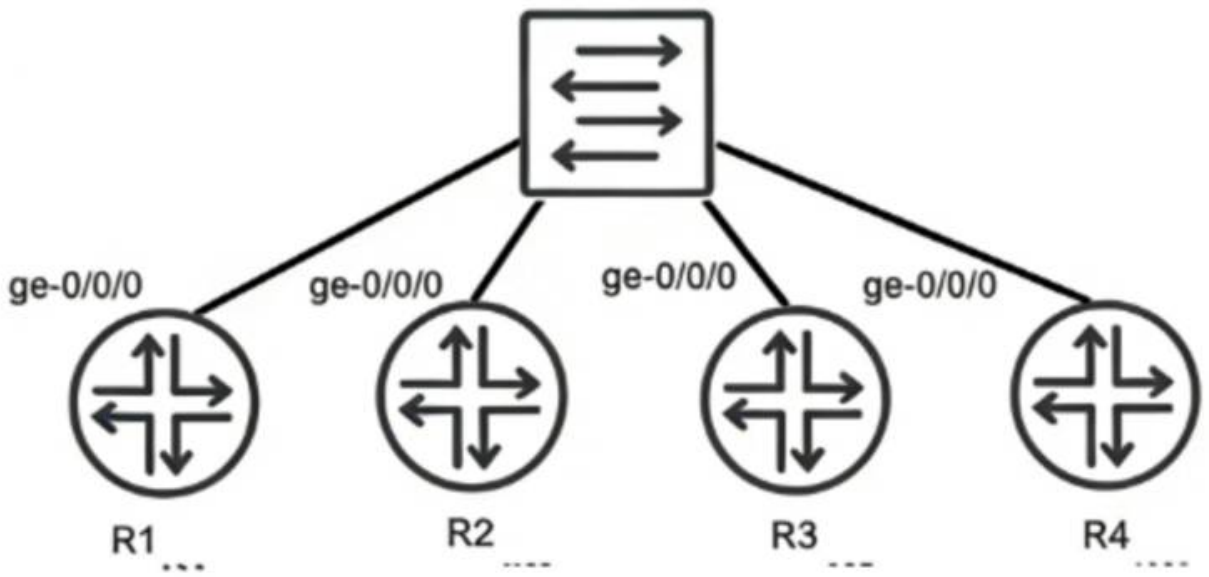


```
user@R1> show configuration routing-options
router-id 192.168.1.1;

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 3
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 4
The MPLS Label Information Base (LIB) is stored in which table?

A. inet6.0
B. mpls.0
C. inet.3
D. inet.0

Answer: B

Explanation:
In Junos OS, the Routing Engine maintains several different tables to manage various types of reachability and forwarding information. When a router is running MPLS, it must track both IP routes and label-to-label mappings.
Thempls.0table is the primary repository for theLabel Information Base (LIB)and theLabel Forwarding Information Base (LFIB). According to Juniper Networks documentation, mpls.0 is used by transit and egress routers to perform label lookups. When a labeled packet arrives at an interface, the router looks at the top label and references the mpls.0 table to determine the next action. This table stores the mapping of incoming labels to their corresponding operations:Pop(remove the label),Swap(replace the label), orPush(add an additional label).
It is crucial to understand the roles of the other tables to avoid confusion:
inet.0 (Option D):This is the default unicast routing table for IPv4, used for standard IP-to-IP forwarding.
inet.3 (Option C):This is theMPLS Path Table. It stores the egress loopback addresses of LSPs and is used by BGP for next-hop resolution to determine if a destination can be reached via an MPLS tunnel. While inet.3 knowsaboutLSPs, the actual label-switching instructions reside in mpls.0.
inet6.0 (Option A):This is the default unicast routing table for IPv6.
Therefore, for the specific purpose of storing the label base used for transit switching operations,mpls.0is the correct and only table used in the Junos architecture.

NEW QUESTION 5
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 6
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 7
You are evaluating BGP between two Juniper routers and the BGP session is stuck in the Idle state. What would cause this behavior?

A. The BGP hold time is too short.
B. The BGP group type is set to internal instead of external.
C. The local AS number is missing.
D. The peer IP address is incorrect.

Answer: D

Explanation:
In the BGP Finite State Machine (FSM), the Idle state is the first stage of any BGP connection. When a BGP session is "stuck" in Idle, it typically indicates that the router is unable to even begin the process of establishing a TCP connection with its neighbor. According to Juniper Networks documentation, before BGP can transition to the Connector Active states, it must have a valid route to the neighbor's IP address in the routing table and be able to initiate a three-way TCP handshake on port 179.
If the peer IP address is incorrect(Option D), the router may not have a route to that destination, or it may be attempting to connect to a non-existent or unreachable host. In many Junos configurations, if the underlying IGP (OSPF/IS-IS) or static routing cannot provide reach ability to the neighbor address defined in the BGP configuration, the BGP process will remain in the Idle state and periodically retry the connection.
Regarding the other options:
The local AS number is missing (Option C):In Junos, you cannot commit a BGP configuration if the local autonomous system is not defined at either the [edit routing-options] level or within the BGP group itself. The commit check would fail before the session could even attempt to start.
The BGP group type (Option B):Having a mismatch in group type (internal vs. external) usually results in the session reaching the Open Sentor Open Confirmstate before failing due to an "unacceptable AS" error in the OPEN message.
BGP hold time (Option A):Issues with hold timers or keep alives generally cause a session that is already in the Established state to drop; they do not prevent the session from leaving the Idle state.

NEW QUESTION 8
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 9
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 10
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 10
In the exhibit, Site A is sending traffic to Site B. R1 adds MPLS label 7166 to direct the traffic to R5.

Exhibit

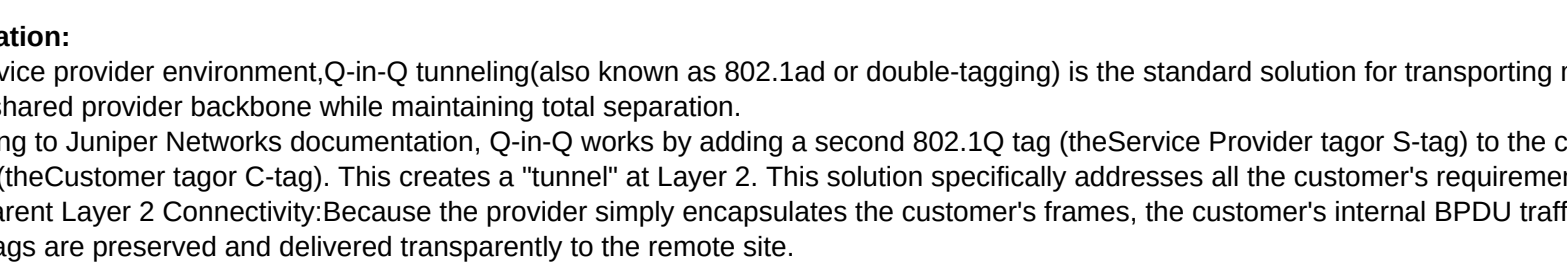

Which two criteria did R1 use to determine which label number to add to the traffic? (Choose two.)

A. the source address of the traffic
B. a label number received from R5
C. the destination address of the traffic
D. a label number advertisement received from R2

Answer: CD

Explanation:
In a Juniper Networks MPLS environment, the process by which a router determines how to forward traffic involves both the control plane and the data plane. When R1 (acting as an Ingress Label Edge Router, or LER) receives an IP packet from Site A destined for Site B, it must perform a lookup to decide whether to forward the packet via standard IP routing or via an MPLS Label Switched Path (LSP).
The first criterion R1 uses is thedestination address of the traffic(Option C). Upon receiving the native IP packet, R1 looks up the destination IP in its routing table (typically inet.0). If the destination matches a prefix that is associated with an LSP—such as the loopback address of R5 or a prefix reachable via R5—the router identifies the appropriate Forwarding Equivalence Class (FEC). The FEC essentially groups packets that should be forwarded in the same manner over the same path. Without identifying the destination, the router cannot map the traffic to the correct MPLS tunnel.
The second criterion is thelabel number advertisement received from R2(Option D). MPLS relies on downstream label allocation. In this topology, R2 is the immediate downstream "next hop" for R1 on the path to Site B. For the LSP to be established, R2 must signal a label to R1 using a protocol like LDP (Label Distribution Protocol) or RSVP (Resource Reservation Protocol). This label (in this case, 7166) tells R1: "If you want to send traffic to the destination associated with this LSP, wrap it in this specific label so I know how to process it."
R1 does not use the source address (Option A) for standard label mapping, nor does it receive the label directly from R5 (Option B) in a hop-by-hop signaling model; it must use the label provided by its direct neighbor, R2. Therefore, by combining the destination IP (to find the path) and the label provided by the next hop (to encapsulate the packet), R1 successfully directs the traffic through the MPLS core.

NEW QUESTION 14
Which term describes the router where traffic enters an MPLS label-switched path (LSP)?

A. egress router
B. transit router
C. penultimate router
D. ingress router

Answer: D

Explanation:
In the architecture of aLabel-Switched Path (LSP), routers are categorized based on their role in the handling of a specific packet's lifecycle through the MPLS network. Juniper Networks documentation defines these roles clearly:
TheIngress Router (Option D), also known as theIngress Label Edge Router (LER), is the entry point of the LSP. Its primary responsibility is to take an incoming "unlabeled" packet (usually a standard IPv4 or IPv6 packet), perform a route lookup, and determine which LSP the packet should follow. Once determined, the Ingress router performs aPushoperation, where it encapsulates the packet with an MPLS label header and forwards it toward the next hop. This is where the transition from IP-based forwarding to Label-based switching occurs.
To contrast this with the other options:
Transit Router (Option B):These are routers located between the ingress and egress. They performSwapoperations, replacing an incoming label with an outgoing label based on the Label Forwarding Information Base (LFIB).
Egress Router (Option A):This is the "tail-end" of the LSP where the packet exits the MPLS domain and the final label is removed (if it hasn't been removed already by the penultimate hop).
Penultimate Router (Option C):This is the second-to-last router in the path. As discussed in previous questions, it often performs thePopoperation (Penultimate Hop Popping) to remove the transport label before sending the packet to the Egress LER.
Therefore, the router where traffic first "enters" the LSP and receives its initial label is strictly defined as the Ingress router.

NEW QUESTION 16
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 18
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your JN0-364 Exam with Our Prep Materials Via below:

https://www.certleader.com/JN0-364-dumps.html