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

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

Answer: C

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

NEW QUESTION 4
What are three extension headers supported by IPv6? (Choose three.)

A. destination options
B. hop-by-hop options
C. protocol
D. header checksum
E. fragment

Answer: ABE

Explanation:
One of the most significant architectural improvements in IPv6 is the move from a complex, variable-length header (as seen in IPv4) to a streamlined, fixed-length base header of 40 bytes. Additional functionality that was previously handled by "Options" in IPv4 is now moved toExtension Headers, which are inserted between the IPv6 base header and the upper-layer protocol (TCP/UDP).
According to Juniper Networks technical documentation and RFC 8200, the following are valid IPv6 Extension Headers:
Hop-by-Hop Options (Option B):This header carries optional information that must be examined by every node along the delivery path. It is used for features like the Router Alert and Jumbo Payload options.
Fragment (Option E):Unlike IPv4, where any router can fragment a packet, in IPv6, fragmentation is performed only by the source node. The Fragment header contains the information necessary for the destination to reassemble the packet (Offset, Identification, and More Fragments flag).
Destination Options (Option A):This header carries information intended only for the destination node. It can appear twice: once before a routing header and once after.
Why other options are incorrect:
Protocol (Option C):In IPv4, this was a field in the header. In IPv6, this is replaced by theNext Header field, which identifies the type of the following header (whether it's an extension header or the upper- layer protocol).
Header Checksum (Option D):This field was entirely removed in IPv6. IPv6 relies on the data link layer (Ethernet) and the transport layer (TCP/UDP) to perform error detection, significantly reducing the processing overhead for routers in the core of a service provider network.

NEW QUESTION 5
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 6
You are asked to configure a new network environment that will be based on IPv6 and use OSPF. In this scenario, which two statements correctly identify configuration task considerations? (Choose two.)

A. Participating interfaces must be configured with both IPv4 and IPv6 protocol families and addresses.
B. The router ID used must be based on a 128-bit identifier value.
C. The router ID used must be based on a 32-bit identifier value.
D. Participating interfaces are only required to be configured with the IPv6 protocol family and address.

Answer: CD

Explanation:
When transitioning to an IPv6 environment usingOSPFv3(the version of OSPF designed for IPv6), there are significant architectural differences compared to OSPFv2 (IPv4). According to Juniper Networks technical documentation, OSPFv3 was redesigned to be more protocol-agnostic.
Router ID (Option C):
Despite OSPFv3 routing IPv6 (which uses 128-bit addresses), the OSPFRouter IDremains a32-bit valueformatted like an IPv4 address (e.g., 1.1.1.1). This is a common point of confusion. In a pure IPv6 environment where no IPv4 addresses are configured on any interfaces, a Juniper router cannot automatically derive a Router ID. Therefore, the administrator must manually configure a 32-bit Router ID under [edit routing-options] for the OSPFv3 process to initialize.
Interface Configuration (Option D):
OSPFv3 runs directly over the IPv6 link-local scope. Unlike OSPFv2, it does not require an IPv4 address to function. Therefore, interfaces areonly required to be configured with family inet6(Option D). You do not need "dual-stack" (both IPv4 and IPv6) functionality just to run OSPFv3. The protocol uses the link-local address (fe80::/10) of the interface for neighbor adjacencies and as the next hop for routing updates. This separation allows OSPFv3 to carry multiple "address families" (both IPv4 and IPv6 unicast) if needed, but the base requirement for an IPv6-only network is simply the family inet6 configuration.

NEW QUESTION 7
Which two protocols would be used for dynamic routing in IPv6 environments? (Choose two.)

A. IGMP
B. IS-IS
C. OSPFv2
D. BGP

Answer: BD

Explanation:
The transition to IPv6 requires routing protocols that are capable of carrying 128-bit address information. Juniper Networks Junos OS supports several "IPv6-ready" protocols for dynamic routing.
* 1. IS-IS (Option B):
As discussed in previous questions,IS-ISis inherently extensible due to its use ofTLVs (Type, Length, Value). To support IPv6, the protocol did not need a major rewrite; instead, new TLVs (such as TLV 236 for IPv6 reachability and TLV 232 for IPv6 interface addresses) were added. A single IS-IS process in Junos can simultaneously carry both IPv4 and IPv6 routing information, making it a highly efficient choice for "dual-stack" service provider backbones.
* 2. BGP (Option D):
BGP was updated to support multiple protocols throughMultiprotocol Extensions (MP-BGP), defined in RFC 4760. By usingAddress Family Identifiers (AFI)andSubsequent Address Family Identifiers (SAFI), a single BGP session can exchange NLRI (Network Layer Reachability Information) for IPv4 unicast, IPv6 unicast, and even VPNv4/VPNv6 routes. In Junos, this is configured under the family inet6 unicast hierarchy within the BGP protocols configuration.
Why other options are incorrect:
IGMP (Option A):This is a management protocol for IPv4 multicast (Internet Group Management Protocol). Its IPv6 equivalent isMLD (Multicast Listener Discovery).
OSPFv2 (Option C):OSPF version 2 is strictly for IPv4. To run OSPF in an IPv6 environment,OSPFv3must be used, as it was specifically redesigned to handle the IPv6 address space and link-local communication.

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

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

Answer: B

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

NEW QUESTION 10
You are configuring BGP on a Juniper router to peer with an external provider. After committing the configuration, the BGP session remains in the Idle state. Which configuration issue would prevent the BGP session from progressing beyond the Idle state?

A. The peer IP address is unreachable.
B. The local AS number is higher than the peer??s AS number.
C. The peer is configured with a different router ID.
D. The BGP group type is set to internal instead of external.

Answer: A

Explanation:
In the BGP finite state machine, the Idle state is the "stop" or "start" point of the protocol. When a session is stuck in Idle, it means the BGP process is either administratively disabled or, more commonly, is unable to initiate the underlying TCP connection required for BGP.
According to Juniper Networks Service Provider documentation, the most common reason for a BGP session to remain in Idle is a lack of routing reach ability. For BGP to move to the Connect state, the Junos kernel must have a route to the IP address specified in the neighbor statement. If thepeer IP address is unreachable (Option A)—meaning there is no route in inet.0 (via OSPF, IS-IS, or static)—the router cannot initiate the TCP three-way handshake on port 179. Consequently, the state machine will never progress.
Analysis of incorrect options:
Option B:BGP does not care if the local AS is higher or lower than the peer's; it only cares if they match the configuration. AS numbers are identifiers, not priorities.
Option C:A mismatched Router ID does not prevent a session from leaving the Idle state. It would typically cause the session to reach the Open Confirm state, and then fail with a "Notification" message due to a collision or identification error.
Option D:While a mismatched group type(internal vs. external) will cause the session to fail, it usually fails during the Open message exchange (Open Sent state) because the AS numbers provided will not match the expected peer type (IBGP vs. EBGP).
Only the lack of a path to the neighbor (reach ability) keeps the session at the very beginning of the process: the Idle state.

NEW QUESTION 10
What happens if an IS-IS router receives a link-state PDU with a higher sequence number than the one in its database?

A. It ignores the link-state PDU.
B. It updates its database and floods the link-state PDU.
C. It sends a CSNP to request confirmation from the source of the link-state PDU.
D. It resets the adjacency with the source of the link-state PDU.

Answer: B

Explanation:
IS-IS is a link-state protocol that relies on the rapid and consistent flooding ofLink-State PDUs (LSPs)to ensure that every router in an area has an identical view of the topology. To manage the "freshness" of information, IS-IS uses aSequence Number—a 32-bit unsigned integer that increments every time the originating router makes a change to its LSP.
According to Juniper Networks technical documentation, when a router receives an LSP, it performs a comparison between the received LSP and the version it currently holds in itsLink-State Database (LSDB). If the received LSP has ahigher sequence number, the router concludes that this is "newer" and more accurate information. The router will then perform two immediate actions:
Update:It replaces the older LSP in its LSDB with the newly received version.
Flood:It propagates the new LSP to all other neighbors (except the one that sent it) to ensure the entire area converges on the new data.
If the sequence numbers were equal, the router would ignore the incoming PDU as it already has the information. If the received sequence number were lower, the router would conclude its own database is more recent and would actually send its own "newer" version back to the neighbor to bring them up to date (a process called "poisoning" or refreshing the neighbor).Complete Sequence Number PDUs (CSNPs)(Option C) are used during initial database synchronization or periodic checks on broadcast links, but the primary response to a "newer" LSP is immediate database update and flooding.

NEW QUESTION 13
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 16
By default, which routing table contains a list of all ingress LSPs?

A. inet.2
B. inet.3
C. inet.1
D. inet.0

Answer: B

Explanation:
In the Juniper Networks Junos operating system, the management of routing information is partitioned into several distinct routing tables (RIBs), each serving a specific architectural purpose. When dealing with Multiprotocol Label Switching (MPLS), understanding the distinction between inet.0 and inet.3 is fundamental for troubleshooting and traffic engineering.
Theinet.3routing table is specifically designed to store the egress IPv4 addresses ofLabel-Switched Paths (LSPs). When an ingress router successfully establishes an LSP (via RSVP or LDP), it places the host address of the egress router (the tail-end) into the inet.3 table. This table is not used for general packet forwarding; instead, it is primarily used by theBorder Gateway Protocol (BGP)for next-hop resolution. When BGP receives a route, it checks both inet.0 and inet.3 to resolve the next hop. If a matching entry exists in inet.3, the router knows it can reach that destination via an MPLS tunnel, allowing for the encapsulation of BGP traffic within MPLS.
In contrast,inet.0is the default unicast routing table used for standard IPv4 forwarding and contains routes learned via IGPs (OSPF, IS-IS) or static routing.inet.1is utilized for multicast forwarding (MBGP), andinet.2 is typically used for Multicast Source Discovery Protocol (MSDP) or RPF checks in multicast environments. By isolating LSP egress points in inet.3, Junos prevents MPLS-specific paths from interfering with standard IGP path selection unless the administrator explicitly chooses to merge them (e.g., using the traffic- engineering bgp-igp command). Therefore, by default, the ingress router maintains its list of reachable LSP endpoints in inet.3.

NEW QUESTION 17
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 22
Exhibit:

```
user@R1> show configuration protocols mpls
label-switched-path to-r3 {
    to 192.168.100.3;
}
interface ge-0/0/0.0;
user@R1> show configuration protocols ospf
area 0.0.0.0 {
    interface ge-0/0/0.0;
    interface lo0.0;
}
user@R1> show route 192.168.100.3
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
192.168.100.3/32   *[OSPF/10] 00:05:39, metric 2
                    > to 172.16.1.2 via ge-0/0/0.0
user@R1> show mpls lsp detail
Ingress LSP: 1 sessions
192.168.100.3
From: 192.168.100.1, State: Dn, ActiveRoute: 0, LSPname: to-r3
ActivePath: (none)
LSPtype: Static Configured, Penultimate hop popping
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
Primary                          State: Dn
  Priorities: 7 0
  SmartOptimizeTimer: 180
  Will be enqueued for recomputation in 27 second(s).
 17 Sep 14 20:29:00.840 CSPF: could not determine self
Total 1 displayed, Up 0, Down 1
Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
Transit LSP: 0 sessions
```

```
    Total 0 displayed, Up 0, Down 0
    user@R1> show configuration interfaces
    ge-0/0/0 {
     unit 0 {
         family inet {
             address 172.16.1.1/24;
         }
         family mpls;
     }
}
fxp0 {
    unit 0 {
        family inet {
            address 10.0.1.11/24;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
```

```
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 10.0.1.11/24;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.100.1/32;
            }
        }
    }
```

You have configured an MPLS LSP to 192.168.100.3. However, the LSP is in the down state. Referring to the exhibit, which two actions would solve this problem? (Choose two.)

A. Issue the set routing-options rib inet.3 static route 192.168.100.1 command and commit.
B. Issue the set protocols mpls label-switched-path to-r3 no-cspf command and commit.
C. Issue the set interfaces lo0 family mpls command on router R1 and commit.
D. Issue the set protocols ospf traffic-engineering command and commit.

Answer: BD

Explanation:
In a Juniper Networks environment, establishing a functionalMultiprotocol Label Switching (MPLS)Label- Switched Path (LSP) requires synchronized control plane operations. According to Juniper technical documentation, the most common reason for an LSP to remain in the "Down" state at the ingress router is a failure of theConstrained Shortest Path First (CSPF)algorithm during the path computation phase.
The provided exhibit for routerR1reveals a critical error in the show mpls lsp detail output: "CSPF: could not determine self". This specific error indicates that the CSPF process is unable to find its own local router ID within theTraffic Engineering Database (TED). For CSPF to build a valid TED, the underlying Interior Gateway Protocol (IGP), such as OSPF, must be configured to flood opaque link-state advertisements (Type 10 LSAs) that carry traffic engineering attributes. As seen in the OSPF configuration, traffic engineering is not enabled. Therefore, issuing theset protocols ospf traffic-engineeringcommand (Option D) will allow R1 to populate the TED with its own local information and that of its neighbors, enabling CSPF to calculate a valid path.
Alternatively, an administrator can choose to bypass the requirement for a TED entirely by disabling CSPF on the specific LSP. By issuing theset protocols mpls label-switched-path to-r3 no-cspfcommand (Option B), the router will stop attempting to perform a constrained path calculation. Instead, the signaling protocol (RSVP) will rely on the standardinet.0routing table to determine the hop-by-hop path to the egress destination (192.168.100.3), allowing the LSP to establish without traffic engineering constraints.
Regarding the other options, whilefamily mplsis required on all transit interfaces, the ingress loopback interface (lo0) generally does not require it for standard LSP signaling unless it's used as a transit hop. Furthermore, adding a static route toinet.3(Option A) is used for next-hop resolution of BGP routes over LSPs but does not assist in the signaling or establishment of the LSP itself.

NEW QUESTION 25

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 26
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 27
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 29
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 30
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 35
Exhibit:

```
user@Router-1> show route 172.24/16

inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)

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

...

172.24.0.0/24 *[OSPF/150] 01:31:31, metric 0, tag 0

> to 172.20.0.2 via ge-0/0/2.0

to 172.20.1.2 via ge-0/0/3.0

user@Router-1> show route forwarding-table

Routing table: default.inet

Internet:

Destination Type RtRef Next hop Type Index NhRef Netif

...

172.24.0.0/24 user 0

172.20.0.2 ucst 551 2 ge-0/0/2.0

172.20.1.2 ucst 552 2 ge-0/0/3.0
```

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

A. The router is performing default route load-balancing behavior.
B. The default route load-balancing behavior of this router has been modified.
C. This router will only choose the next hop with a > next to it in the routing table.
D. This router will choose both next hops in the routing table.

Answer: BD

Explanation:
In Junos OS, understanding the distinction between theRouting Information Base (RIB)and theForwarding Information Base (FIB)is fundamental to analyzing traffic patterns and load-balancing behavior. The RIB (show route) contains all prefixes learned via various protocols, while the FIB (show route forwarding-table) contains only the active next-hops that are actually programmed into the Packet Forwarding Engine (PFE).
According to Juniper Networks technical documentation, the default behavior for Junos OS when encounteringEqual-Cost Multipath (ECMP)routes is to select only a single next-hop from the available candidates in the RIB and install that single path into the FIB. In a default state, even if the show route output displays multiple next-hops for a destination like 172.24.0.0/24, only one would have the active route symbol (>) and only that one would appear in the forwarding table.
In the provided exhibit, the show route output shows two next-hops for 172.24.0.0/24, but only the first one (172.20.0.2) is marked with the>symbol as the active selection. However, the subsequent show route forwarding-table output reveals thatboth next-hops(172.20.0.2 and 172.20.1.2) are currently present in the forwarding table for that same destination. This discrepancy indicates that thedefault load-balancing behavior has been modified (Option B). This modification is typically achieved by creating a routing policy with the action then load-balance per-packet (which actually results in flow-based load balancing) and applying it to

the forwarding table via the export statement under [edit routing-options forwarding-table].
Because the forwarding table now contains both next-hops, the router is no longer restricted to a single path. Therefore, therouter will choose both next-hops in the routing table (Option D)for packet forwarding, distributing flows across the two available Gigabit Ethernet interfaces (ge-0/0/2.0 and ge-0/0/3.0). This ensures higher utilized bandwidth and provides redundancy at the data plane level.

NEW QUESTION 39
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 40
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 45
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 47
In the exhibit, Site A is sending traffic to Site B. R1 adds MPLS label 7166 to direct the traffic to R5.

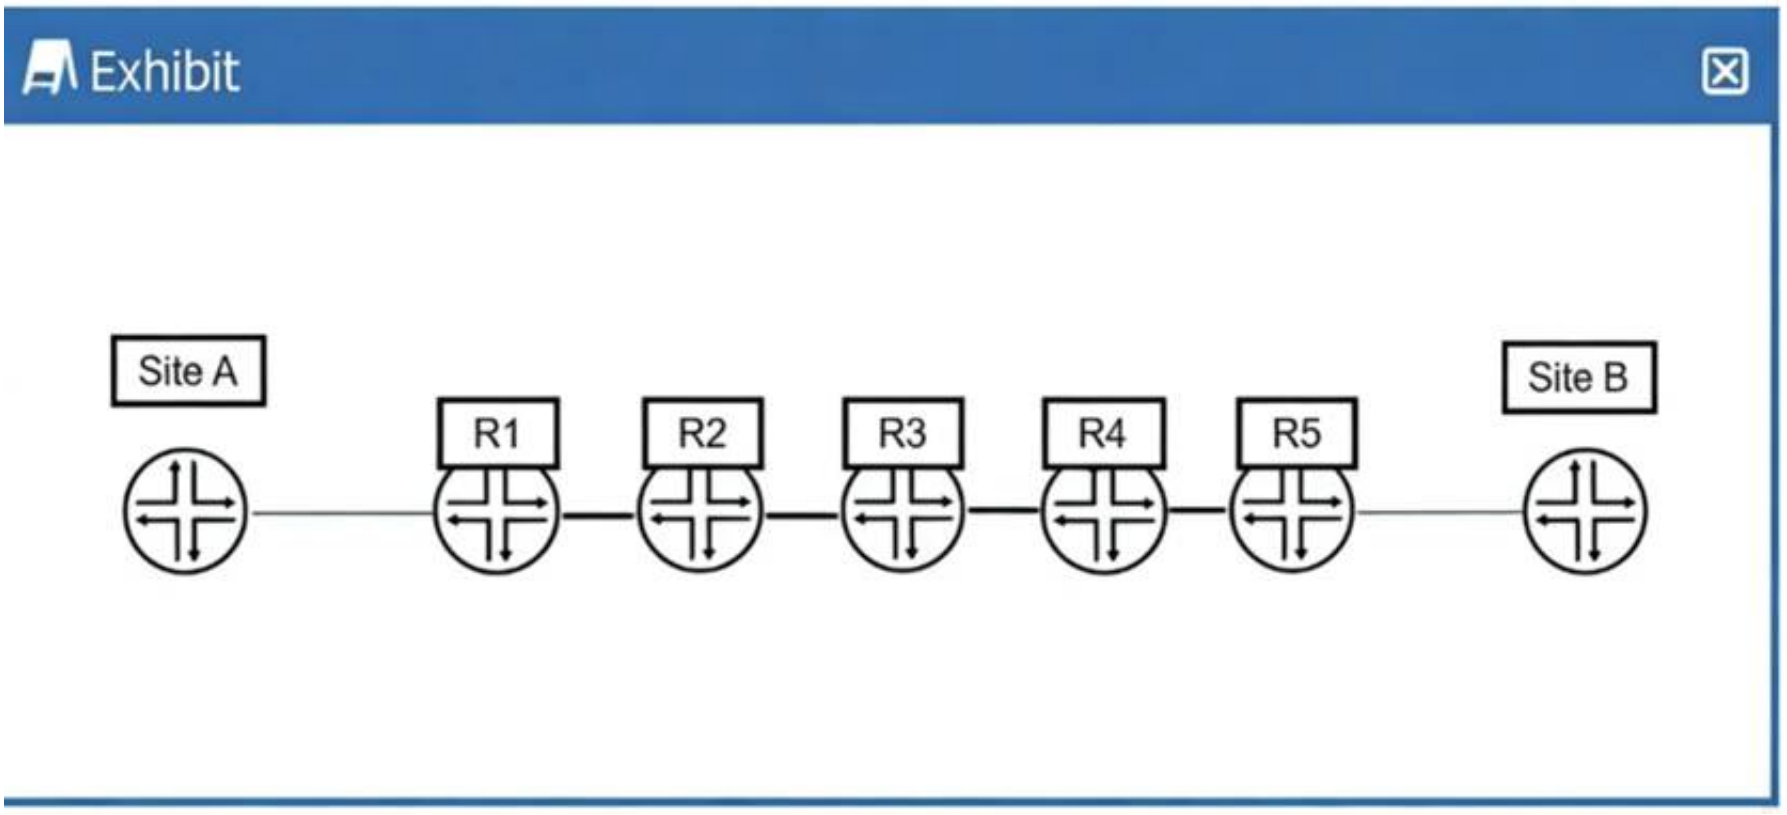

Which two criteria did R1 use to determine which label number to add to the traffic? (Choose two.)

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 50
Which IPv6 extension header is used to specify intermediate nodes for a packet's path?

A. hop-by-hop options
B. routing
C. fragment
D. destination options

Answer: B

Explanation:
In the IPv6 architecture, the base header is kept at a fixed size of 40 bytes to streamline processing. Any additional features or options are handled byExtension Headers, which are inserted between the IPv6 header and the upper-layer protocol. According to Juniper Networks technical documentation and RFC 8200, when a source node needs to list one or more intermediate nodes to be "visited" on the way to the final destination, it utilizes theRouting extension header (Option B).
The Routing header is functionally similar to the "Source Route" option in IPv4. When a packet contains a Routing header, it is addressed to the first intermediate node listed in the header. That node examines the header, swaps its own address with the next address in the list, and forwards the packet. This process continues until the packet reaches the final destination. This is a foundational component for technologies likeSegment Routing over IPv6 (SRv6), where the Routing header (specifically the Segment Routing Header or SRH) is used to steer traffic through a specific set of service instructions or nodes.
To distinguish this from the other options:
Hop-by-hop options (Option A):These carry information that must be examined byeverynode along the path (such as Router Alert), not just specific intermediate nodes.
Fragment (Option C):This is used only when the source node needs to fragment a packet that exceeds the path MTU.
Destination options (Option D):These carry optional information intended specifically for the destination node (or nodes listed in a Routing header), but they do not dictate the path themselves.

NEW QUESTION 54
By default, which MPLS operation is performed by the penultimate router in an LSP on the transport label?

A. swap
B. push

C. rewrite
D. pop

Answer: D

Explanation:
In a Multiprotocol Label Switching (MPLS) environment, label operations are categorized into three primary actions:Push(adding a label),Swap(replacing a label), andPop(removing a label). The specific behavior described in the question refers to a mechanism calledPenultimate Hop Popping (PHP).
According to Juniper Networks technical documentation, the goal of PHP is to improve forwarding efficiency at the egress point of a Label-Switched Path (LSP). TheEgress Label Edge Router (LER), which is the final destination for the LSP, would normally have to perform two lookups if it received a labeled packet: first, it would look up the label in its MPLS table to see it is the destination, and second, it would look up the underlying IP payload in its IP routing table (inet.0) to forward the packet.
To alleviate this burden, the Egress LER signals a special label value calledImplicit Null (Label 3)to its upstream neighbor (the penultimate router) during the signaling process (RSVP or LDP). When thepenultimate routerreceives a packet destined for that egress LER, it sees the instruction to pop the transport label. Consequently, the penultimate router performs aPopoperation, stripping away the outer MPLS label and sending the raw IP packet (or the remaining inner service label) to the Egress LER.
This allows the Egress LER to perform only a single lookup. If the transport label was the only label, the Egress LER simply performs a standard IP lookup. If there is a VPN label remaining, it performs a single MPLS lookup for the VRF. This "default" behavior in Junos OS optimizes the performance of the egress router by offloading the final label removal to the penultimate hop. Note that ifUltimate Hop Popping (UHP)were configured (via the explicit-null command), the penultimate router would perform aSwapto Label 0 instead of a Pop.

NEW QUESTION 58
You are the administrator for two Junos routers called R1 and R2. These two routers are directly connected to each other. These two routers run IS-IS and BFD. R1 is configured to send BFD packets every 300 milliseconds. R2 is configured to send BFD packets every 400 milliseconds. In this situation, what is the expected outcome?

A. Each router will send BFD packets at the rate that has been locally configured.
B. BFD will fail due to the mismatched timers.
C. Each router will negotiate to send BFD packets at the slowest of the two rates.
D. Each router will negotiate to send BFD packets at the fastest of the two rates.

Answer: C

Explanation:
In the context of Juniper Networks High Availability,Bidirectional Forwarding Detection (BFD)is a lightweight protocol designed to provide fast failure detection for the forwarding path. Unlike the slow "hello" mechanisms found in IGPs like OSPF or IS-IS, BFD can detect link or neighbor failures in sub-second intervals.
According to Juniper Networks technical documentation, BFD operates through a negotiation process. When two routers establish a BFD session, they exchange their locally configuredMinimum Transmit IntervalandMinimum Receive Intervalwithin the BFD control packets. The fundamental rule of BFD negotiation is that the routers must agree on a common timing value that accommodates the slower of the two devices to ensure stability and prevent "false positives" (detecting a failure when none exists simply because one router cannot keep up with the processing speed).
In this scenario, R1 expects to send at 300ms, while R2 is configured for 400ms. During the handshake, R1 informs R2 it is capable of 300ms, but R2 informs R1 it can only support a minimum of 400ms. Consequently, the routers will negotiate to use theslowest of the two rates (400ms). Specifically, the transmission interval of one router is matched to the receive interval of the other. By choosing the highest common denominator (the slowest rate), the BFD session ensures that both routers have sufficient time to process incoming control packets. This negotiation allows BFD to be highly flexible in heterogeneous environments where different hardware platforms may have varying CPU capabilities for handling rapid heartbeat packets.

NEW QUESTION 61
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 64
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 66
.....

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