



Juniper JN0-364 JNCIS-SP Practice Questions

Shared by Dotson on 19-10-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: MultipleChoice

What information is determined by using the AS path attribute included in the BGP update message? (Choose two.)

Options:

- A- the origin of a route from IGP or EGP
- B- the presence of a routing loop
- C- the shortest AS path to reach a prefix
- D- the total number of next-hop devices to reach a prefix

Answer:

B, C

Explanation:

The AS_PATH attribute is a 'well-known mandatory' attribute in BGP, meaning it must be present in every BGP Update message exchanged between External BGP (eBGP) peers. It records the sequence of Autonomous System numbers that a route has traversed. Per Juniper Networks Service Provider documentation, this attribute serves two fundamental purposes:

1. Loop Prevention (Option B):

This is the most critical function of the AS_PATH. When a BGP router receives an update from an eBGP peer, it scans the AS_PATH attribute for its own AS number. If the router finds its local AS number already listed in the path, it concludes that the route has already passed through its network and has 'looped' back. To prevent an infinite routing loop, the router will immediately discard the update. This mechanism is the cornerstone of BGP's stability as a path-vector protocol.

2. Path Selection / Shortest Path Determination (Option C):

BGP uses a complex 'tie-breaking' algorithm to select the best path among multiple candidates. One of the highest-ranking criteria in this algorithm (after Weight, Local Preference, and AS_PATH length) is the length of the AS_PATH. A shorter AS_PATH (fewer AS numbers listed) is generally preferred over a longer one, as it typically represents a more direct path through the internet hierarchy.

Why other options are incorrect:

Option A: The 'origin' of a route (IGP, EGP, or Incomplete) is determined by the ORIGIN attribute, which is a separate well-known mandatory attribute.

Option D: BGP does not count individual 'next-hop devices' (which would be an IGP metric like hop count in RIP); it only tracks Autonomous Systems. A single AS in the path might contain hundreds of internal routers (next-hops), but BGP only sees it as one 'hop' in the AS_PATH.

Question 2

Question Type: MultipleChoice

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? (Select two.)

Options:

- 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:

C, D

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 the destination 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 the label 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.

Question 3

Question Type: MultipleChoice

What prevents routing loops in a single-area OSPF network?

Options:

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

Answer:

A

Explanation:

In OSPF, loop prevention within a single area is achieved through the fundamental nature of its link-state architecture. Unlike distance-vector protocols that rely on 'routing by rumor,' OSPF ensures that every router within an area maintains an identical Link-State Database (LSDB). This database acts as a complete map of the network topology.

Once the LSDB is synchronized, each router independently executes the Shortest Path First (SPF) algorithm, which is formally known as the Dijkstra algorithm. This mathematical process treats the local router as the 'root' of a tree and calculates the shortest path to every other node (router) and prefix in the area based on the cumulative interface costs. Because every router uses the same synchronized map (the LSDB) and the same deterministic algorithm, they all arrive at a consistent, loop-free view of the best paths.

According to Juniper Networks technical documentation, the Dijkstra algorithm is superior to the Bellman-Ford algorithm (used by distance-vector protocols like RIP) in this regard. Bellman-Ford is susceptible to 'count-to-infinity' problems and loops because routers only know the

distance and direction to a destination provided by their neighbors, rather than the full topology. In OSPF, even if a link fails, the updated Link-State Advertisement (LSA) is flooded rapidly, and the Dijkstra algorithm is re-run to find a new loop-free path. Routing policies (Option B) are used to manipulate path selection or filter routes but are not the primary mechanism for fundamental loop prevention in OSPF. Similarly, forwarding policies (Option D) govern how traffic is handled at the data plane level rather than determining the control plane's loop-free topology.

Question 4

Question Type: MultipleChoice

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? (Select two.)

Options:

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

Answer:

A, D

Explanation:

To support dual 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 <interface-name>]).

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, and gigether-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.

Question 5

Question Type: MultipleChoice

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

Options:

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

Answer:

A

Explanation:

In Juniper Networks Junos OS, the Label Distribution Protocol (LDP) is specifically designed to automate the creation of Label Switched Paths (LSPs) based on the information provided by the underlying Interior Gateway Protocol (IGP), such as OSPF or IS-IS. When LDP is enabled on a set of interfaces within an OSPF area (as shown in the exhibit with Area 0.0.0.0), it automatically discovers neighbors and exchanges label mappings for all known unicast routes in the routing table.

The defining characteristic of LDP in this context is its 'topology-driven' nature. Unlike RSVP (Resource Reservation Protocol), which typically requires the manual configuration of each LSP ingress point and destination, LDP follows the IGP's shortest path tree to automatically build a full mesh of LSPs between all participating routers. This means that every Provider Edge (PE) and Provider (P) router in the exhibit---PE1, PE2, PE3, P1, P2, and P3---will establish label-switched connectivity to every other router without the administrator having to define individual tunnels.

LDP accomplishes this through a downstream-unsolicited label distribution mode by default in Junos. Each router assigns a local label for its loopback address and other prefixes and advertises these to its neighbors. Because every router is performing this action for every reachable prefix in the OSPF domain, a complete fabric of label-switched paths is formed. While RSVP is more robust for traffic engineering and bandwidth reservation, LDP is the preferred protocol for creating a simple, scalable full mesh of LSPs for applications like Layer 3 VPNs or internal BGP tunneling where complex path manipulation is not required. BFD is a failure detection protocol, and BGP is used for service signaling, making LDP the only correct choice for automatic mesh creation.

Question 6

Question Type: MultipleChoice

Exhibit:

P2P
exams

P2P
exams

```

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 {
            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.)

Options:

- 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:

B, D

Explanation:

In a Juniper Networks environment, establishing a functional Multiprotocol 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 the Constrained Shortest Path First (CSPF) algorithm during the path computation phase.

The provided exhibit for router R1 reveals 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 the Traffic 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 the set protocols ospf traffic-engineering command (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 the set protocols mpls label-switched-path to-r3 no-cspf command (Option B), the router will stop attempting to perform a constrained path calculation. Instead, the signaling protocol (RSVP) will rely on the standard inet.0 routing 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, while family mpls is 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 to inet.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.

Question 7

Question Type: MultipleChoice

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

Options:

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

Answer:

C, D

Explanation:

In the context of Juniper Networks Junos OS and the IS-IS (Intermediate System to Intermediate System) protocol, understanding the hierarchical relationship between router levels is critical for effective troubleshooting and design. IS-IS uses a two-level hierarchy to manage scalability: Level 1 (L1), which represents intra-area routing, and Level 2 (L2), which represents inter-area backbone routing.

When a router is configured as a Level 1/2 (L1/L2) device, it acts as a bridge between the two levels. According to Juniper technical documentation, an L1/L2 router maintains two completely separate Link-State Databases (LSDB)---one for Level 1 and one for Level 2. It does not merge these into a single topology. This separation ensures that local area topology changes (L1) do not necessarily flood into the backbone (L2) unless specific redistribution is configured, and vice versa. Therefore, statement C is correct because the L1/L2 router maintains distinct SPF (Shortest Path First) computations for each level.

Regarding the visibility of the Level 1 router, IS-IS is designed to keep L1 areas 'stubby' by default. A Level 1 router only possesses the topology information for its own area (the Level 1 LSDB). It does not receive specific L2 routes or the L2 topology. Instead, the L1/L2 router sets the Attached (ATT) bit in its L1 Link-State PDUs (LSPs) to signal to L1-only routers that it has a connection to the backbone. The L1 router then generates a default route pointing to the L1/L2 router to reach inter-area destinations. This confirms that statement D is correct: the L1 router's knowledge is limited to its local L1 topology.

Conversely, statements A and B are incorrect because merging topologies would violate the hierarchical scaling principles of IS-IS, and L1 routers never learn the full L2 topology without explicit, non-standard route leaking.

Question 8

Question Type: MultipleChoice

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?

Options:

- 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 an explicit path---where the network architect defines specific hop-by-hop routers that the traffic must traverse---the Resource Reservation Protocol (RSVP) is the necessary choice.

According to Juniper documentation, RSVP (specifically RSVP-TE) supports the use of Explicit 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.

Question 9

Question Type: MultipleChoice

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?

Options:

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

Answer:

D

Explanation:

For Juniper platforms equipped with dual Routing Engines (REs), the fundamental technology required to provide high availability during a hardware or software failure of the primary RE is Graceful 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 the chassis state and the interface 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 router continues to forward packets that 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 the forwarding state, it does not by itself

maintain the routing protocol state (adjacencies). To keep OSPF or BGP sessions from dropping during the switchover, GRES must be paired with Non-Stop Active Routing (NSR). However, as the question focuses on the core requirement of continuing to forward packets, GRES is the foundational technology.



To Get Premium Files for JN0-364 Visit

<https://www.p2pexams.com/products/jn0-364>

For More Free Questions Visit

<https://www.p2pexams.com/juniper/pdf/jn0-364>

20%
DISCOUNT

P2P
exams