Juniper JN0-281 Practice Questions

Shared by Harrell on 19-10-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You are creating a static route with a next hop that is not directly connected. Which feature should be used to accomplish this task?

Options:

A- install
B- resolve
C- qualified next hop
D- retain

Answer:

B

Explanation:

In Junos, a static route normally expects its next hop to be directly reachable on a connected interface. When the configured next hop is not directly connected, the router must determine how to reach that next hop using an existing route in the routing table. The resolve feature provides this behavior by allowing Junos to recursively resolve the configured static next hop through another route, such as an IGP learned route, a directly connected route to an intermediate device, or even another static route. Once the system can resolve the next hop to a usable outgoing interface and a final forwarding next hop, the static route becomes active and can be installed in the forwarding table.

This is common in data center environments where you want to point a static route at a loopback address, a service node address, or a next hop that is reachable through the fabric underlay rather than a directly connected subnet. With resolve enabled, the static route's validity follows the reachability of the recursive path. If the supporting route used for resolution disappears, the static route is withdrawn, helping avoid blackholing traffic toward an unreachable next hop.

The qualified next hop feature is used to define primary and backup next hops with different preferences for the same static route, not to solve indirect reachability by itself. The install and retain options influence route installation and retention behaviors but do not provide recursive resolution of a non-direct next hop.

Question 2

Question Type: MultipleChoice

In the Junos OS, which feature is used to create an alternate next hop with a unique preference for a static route?

Options:

A- Preference
B- Resolve
C- Next-hop
D- Qualified-next-hop

Answer:

D

Explanation:

In Junos OS, the qualified-next-hop feature is used to specify an alternate next hop for a static route, along with a unique preference value.

Step-by-Step Breakdown:

Qualified-Next-Hop:

A qualified-next-hop allows you to define multiple next hops for a static route, each with its own preference. This provides flexibility by allowing the router to choose the best available next hop based on reachability and preference.

Use Case:

If the primary next hop becomes unreachable, the router can automatically switch to the alternate next hop defined by the qualified-next-hop with a higher preference value.

Command Example:

set routing-options static route 10.10.10.0/24 qualified-next-hop 192.168.1.1 preference 5

set routing-options static route 10.10.10.0/24 qualified-next-hop 192.168.1.2 preference 10

Preference:

The next hop with the lowest preference is chosen first. If it becomes unavailable, the router will use the higher preference next hop.

Juniper Reference:

Qualified-Next-Hop: This feature is used to configure backup or alternate next hops for static routes in Juniper devices.

Question 3

Question Type: MultipleChoice

Which BGP attribute is used to prevent routing loops in a network?

Options:

A- MED
B- LOCAL_PREF
C- NEXT_HOP
D- AS_PATH

Answer:

D

Explanation:

BGP prevents inter-domain routing loops primarily through the AS_PATH attribute. Each time a route advertisement crosses an autonomous system boundary, the advertising router prepends its own AS number to the AS_PATH before sending the route onward. When a router receives a BGP update, it can inspect the AS_PATH and reject routes that already contain its own AS number. This simple rule blocks the route from circulating back into the same AS, which is the fundamental loop prevention behavior in BGP.

In data center IP fabrics that use EBGP for the underlay, AS_PATH still plays an important role even though private AS numbers are used. The fabric can be designed so that leaf and spine devices are in distinct private ASs or use structured AS assignment, and AS_PATH tracking ensures clean propagation and helps operators understand where a route has traveled. It also supports policy decisions, such as preferring shorter paths or applying policy based on specific AS sequences.

The other attributes serve different purposes. MED influences inbound path selection into an AS, LOCAL_PREF is used within an AS to influence outbound selection, and NEXT_HOP indicates the forwarding next hop and is validated for reachability. None of those attributes provide the explicit ''contains my AS'' loop detection that AS_PATH does.

Question 4

Question Type: MultipleChoice

What are two available modes when using LACP with an aggregated Ethernet bundle? Select two.

Options:

A- aggressive
B- mixed
C- passive
D- active

Answer:

C, D

Explanation:

On Junos devices used in data centers, an aggregated Ethernet bundle can run either static bundling or dynamic bundling using LACP. When LACP is used, Junos supports two negotiation modes: active and passive. These modes control whether the device initiates LACP negotiation by transmitting LACP Data Units or whether it waits to respond to LACP Data Units sent by its peer. In active mode, the system periodically sends LACP control frames to begin and maintain the negotiation. In passive mode, the system does not initiate negotiation but will respond if it receives LACP control frames from the neighbor.

From an operational perspective, active is typically recommended on at least one side of the link so the bundle reliably forms even if the peer is configured to be passive. If both ends are passive, each side waits for the other to start, and the aggregated link might not come up as expected. This is a common cause of down or partially formed bundles in leaf spine uplinks and server dual-homing scenarios.

Options aggressive and mixed are not Junos LACP modes for aggregated Ethernet. In Junos, you configure LACP under the aggregated-ether-options hierarchy and select active or passive behavior, along with optional timing such as periodic fast for quicker detection.

Question 5

Question Type: MultipleChoice

Referring to the exhibit, which statement is correct about the 192.168.1.0/24 route?

Options:

A- The 192.168.1.0/24 route is load balanced across both next hops.
B- The 192.168.1.0/24 route will select 10.20.20.254 as the next hop.
C- The 192.168.1.0/24 route will select 10.20.20.253 as the next hop.
D- The 192.168.1.0/24 route is not installed in the routing table.

Answer:

B

Explanation:

The configuration defines a static route for 192.168.1.0/24 with a primary next hop of 10.20.20.254 and a qualified-next-hop of 10.20.20.253 with preference 25. In Junos, qualified next hop is used to create preference-ranked redundancy for a static route, typically primary and backup forwarding. The key behavior is that the best, lowest preference path is selected as active. For static routes, the default preference for the route and its primary next hop is lower than 25 unless you explicitly raise it. Because the qualified-next-hop is configured with preference 25, it is less preferred than the primary next hop, so 10.20.20.254 remains the selected next hop during normal operation.

This is not load balancing. Junos does not automatically install both next hops for forwarding simply because more than one next hop is configured when one is qualified with a higher preference. Instead, the qualified next hop is held in reserve and is used only when the primary next hop becomes unusable, such as when it is no longer resolvable or the associated forwarding condition fails. Therefore, the route is expected to be installed and active using 10.20.20.254 under normal conditions, with 10.20.20.253 acting as a backup path.

Question 6

Question Type: MultipleChoice

You want your Junos OS router to drop packets that it receives for a particular destination IP address. You also want to notify the source of the traffic that the destination IP address is unreachable. How would you accomplish this task?

Options:

A- Create a static route for the destination IP address in the inet.3 routing table.
B- Add the destination IP address to the list of martian IP addresses.
C- Create a static route for the destination IP address with a next hop of reject.
D- Create and apply a firewall filter to the ingress interface that discards traffic destined to the IP address.

Answer:

C

Explanation:

To intentionally drop traffic to a specific destination while also informing the sender that the destination is unreachable, Junos provides a routing action specifically meant for this behavior: a static route with a next hop of reject. A reject route installs a forwarding entry that causes matching packets to be dropped, and it also triggers generation of an ICMP unreachable message back to the traffic source. This is useful in data center routing for controlled blackholing with feedback, for example when you want to signal misrouted traffic, invalid destinations, or to enforce policy with explicit notification.

A discard route also drops packets, but it does so silently without sending ICMP unreachable messages. That makes discard appropriate for traffic-sink use cases such as safe summarization or DDoS mitigation where you do not want to generate return traffic. Adding addresses to martians is intended to block invalid special-use prefixes from being treated as routable, not to create a targeted unreachable response for an arbitrary destination. Using inet.3 is related to MPLS label-switched forwarding and does not address the requirement. A firewall filter that discards traffic will typically drop silently unless explicitly configured to send rejects in a security context, and it is not the standard routing-table-based method when the requirement is an ICMP unreachable response for a destination prefix.

Therefore, a reject static route is the correct solution.

Question 7

Question Type: MultipleChoice

According to Juniper Networks, the bridge table is more commonly known as a _____.

Options:

A- forwarding table

B- forwarding bridge
C- bridging information table
D- forwarding information table

Answer:

A

Explanation:

In Ethernet switching, the bridge table is the data structure that maps MAC addresses to the switch interfaces where those MAC addresses were learned. Juniper commonly describes this function as the Ethernet switching table and also refers to it as the forwarding table in Layer 2 contexts. The concept is the same: the switch learns source MAC addresses from incoming frames, associates them with an ingress port and VLAN or bridge domain, and then uses that learned information to forward future frames to the correct egress port as known unicast.

Calling it a forwarding table is accurate because its primary operational purpose is deciding how to forward Layer 2 frames efficiently. When a destination MAC is present in the table, the switch performs a unicast forward to the learned port. When a destination MAC is not present, the switch treats it as unknown unicast and floods it within the VLAN or bridge domain, while still learning the source MAC for future use.

The term forwarding information table is more strongly associated with Layer 3 routing, where a FIB represents resolved next hops for IP prefixes in the forwarding plane. That is a different structure than the Layer 2 bridge or MAC table. The other options are not standard Juniper terms for this function.

Verification sources from Juniper documentation

https://www.juniper.net/documentation/us/en/software/junos/multicast-l2/topics/topic-map/ethernet-switching-components.html

https://www.juniper.net/documentation/us/en/software/junos/multicast-l2/topics/concept/ethernet-switching-table-understanding.html

Question 8

Question Type: MultipleChoice

Which statement is correct about building an IP fabric?

Options:

A- Each leaf device should have a direct physical connection to every other leaf device.
B- Each leaf device must have two or more physical connections to each spine device.
C- Each leaf device should have a direct physical connection to every spine device.
D- Each leaf device must have two or more physical connections to each connected server.

Answer:

C

Explanation:

A modern IP fabric in a leaf spine topology is built to deliver predictable latency, high bandwidth, and horizontal scale. The defining characteristic is that every leaf switch connects upstream to every spine switch. This full mesh between leaf and spine layers creates multiple equal cost paths between any two endpoints connected anywhere in the fabric. With this design, east west traffic between servers attached to different leaves can traverse the fabric using one spine hop, keeping path length consistent and enabling efficient load sharing across links.

Leaf to leaf direct connections are not used in a standard leaf spine IP fabric because they create irregular topologies, complicate troubleshooting, and undermine the uniform multi path behavior that makes fabrics scalable. Similarly, a leaf does not inherently require two or more physical connections to each spine. Many designs start with a single link per leaf spine pair and increase capacity by adding additional parallel links as needed. Redundancy is achieved by having multiple spines and multiple equal cost paths, not by mandating multiple links to every spine from day one.

Server connectivity requirements also vary. Some servers use single uplinks, others use dual homing for redundancy. That decision is independent of the fundamental requirement that each leaf should connect to every spine.

Question 9

Question Type: MultipleChoice

Which high availability feature permits you to combine multiple Ethernet interfaces into a single link layer interface?

Options:

A- graceful Routing Engine switchover
B- link aggregation groups
C- nonstop bridging

D- nonstop active routing

Answer:

B

Explanation:

Link aggregation groups, implemented on Junos as aggregated Ethernet interfaces, allow multiple physical Ethernet links to operate as one logical Layer 2 interface. This increases available bandwidth and provides link level resiliency because member links can fail without taking down the logical interface, as long as at least one member remains operational. In data center leaf spine designs, link aggregation is commonly used for server dual homing, uplinks to appliances, or inter switch connectivity where parallel links are desired with a single logical adjacency.

From a forwarding perspective, the device distributes traffic across member links using a hashing algorithm based on packet header fields so that individual flows remain in order while the aggregate uses multiple links. Control plane operation can be static or negotiated with LACP. With LACP, both sides exchange protocol information to ensure consistent bundling and to remove failed or miswired members automatically. This makes link aggregation a core high availability building block at the link layer, independent of Routing Engine redundancy features.

Graceful Routing Engine switchover, nonstop bridging, and nonstop active routing are control plane redundancy features. They are designed to minimize disruption during Routing Engine failover or preserve protocol state, but they do not combine multiple physical Ethernet interfaces into one logical link layer interface.

Verification sources from Juniper documentation

https://www.juniper.net/documentation/us/en/software/junos/interfaces-ethernet/topics/topic-map/understanding-lacp.html

https://www.juniper.net/documentation/us/en/software/junos/interfaces-ethernet/topics/topic-map/aggregated-ethernet-interfaces-overview.html

Question 10

Question Type: MultipleChoice

What are two ways to solve the configuration error shown in the exhibit? Choose two.

Options:

A- Connect the xe-0/0/4 interface to the network.
B- Configure the xe-0/0/4.0 interface as a member of a single VLAN.
C- Configure the xe-0/0/4.0 interface as a trunk port.
D- Configure the logical unit number to match the default VLAN ID.

Answer:

B, C

Explanation:

The commit error indicates that the interface is being treated as an access port while the configuration attempts to associate it with more than one VLAN. In Junos Ethernet switching, an access mode interface represents a single untagged VLAN membership. Because access ports accept and transmit frames without 802.1Q tags, the switch must map all ingress untagged traffic to exactly one VLAN. For that reason, Junos enforces the rule that an access interface can be part of only one VLAN, and it rejects configurations that try to add multiple VLAN members under access mode.

There are two valid ways to resolve this, depending on the intended design. First, if the port truly connects to a single endpoint that should live in only one broadcast domain, configure the interface as a member of only one VLAN. This aligns with access port semantics and eliminates the conflict that causes the commit to fail.

Second, if the endpoint or downstream device needs to carry multiple VLANs over the same physical link, change the interface to trunk mode. A trunk port is designed to transport multiple VLANs using 802.1Q tagging, so multiple VLAN members are valid and expected. In data center environments, trunking is common for server virtualization hosts, appliance uplinks, and switch-to-switch links.

Connecting the interface to the network does not affect configuration validation, and logical unit numbering is unrelated to VLAN membership rules for access ports.

Question 11

Question Type: MultipleChoice

Referring to the exhibit, which type of protocol session will be established?

Options:

A- OSPF

B- IS-IS
C- EBGP
D- IBGP

Answer:

D

Explanation:

The configuration shown is under protocols bgp and defines a BGP group named service-pod with the statement type internal. In Junos, the BGP group type determines whether the sessions formed by neighbors in that group are treated as internal BGP or external BGP. When the group is set to type internal, the router establishes an iBGP session to the configured neighbor address. That means both peers are expected to be in the same autonomous system, and the session is used to exchange routes within that single AS.

The local-address statement sets the source address that the router will use when initiating the TCP connection for BGP. In data center designs, this is commonly a loopback address so the session can be resilient and independent of any single physical interface. The neighbor statement identifies the remote peer address. With type internal, the session is iBGP even if the neighbor is directly connected, because the relationship is defined by AS scope rather than topology.

OSPF and IS-IS are different routing protocols entirely and are not established through BGP configuration. EBGP would require type external and typically involves different AS numbers across the peering. Therefore, the only correct answer is IBGP.

Question 12

Question Type: MultipleChoice

What are two characteristics of EBGP? Choose two.

Options:

A- EBGP sessions do not require an IGP.
B- EBGP sessions require loopback IP address peering.
C- EBGP does not support sessions with non-directly connected peers.
D- EBGP sessions are typically established between directly connected peers.

Answer:

A, D

Explanation:

EBGP is the BGP session type formed between different autonomous systems. In Juniper data center IP fabrics, EBGP is frequently used for the underlay because it can provide all required reachability without an additional interior gateway protocol. The fabric can advertise loopbacks and point-to-point link subnets directly in BGP, then use ECMP to install multiple equal-cost next hops. This is why EBGP sessions do not require an IGP as a fundamental dependency. Some designs still add an IGP for other reasons, but EBGP itself can carry the underlay routes needed for full fabric connectivity.

EBGP sessions are also typically established between directly connected peers. In a standard leaf-spine underlay, each leaf peers with each spine over the routed physical links between them, using the interface IP addresses on those point-to-point subnets. This matches EBGP default behavior, including a one-hop TTL expectation and straightforward operational troubleshooting.

Loopback peering is not required for EBGP. It is possible, but it usually needs additional configuration such as multihop and a routing method to ensure reachability to the remote loopback before the BGP session can form. EBGP also supports sessions with non-directly connected peers when multihop is configured, so it is incorrect to claim EBGP does not support that capability.

To Get Premium Files for JN0-281 Visit

https://www.p2pexams.com/products/jn0-281

For More Free Questions Visit

https://www.p2pexams.com/juniper/pdf/jn0-281

