Juniper JN0-481 Practice Questions

Shared by Martin on 19-10-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You are allowed to assign tags for which three objects? (Choose three.)

Options:

A- Virtual networks
B- Interfaces
C- Generic systems
D- Property sets
E- Device profiles

P2P exams

Answer:

A, B, C

Explanation:

In Apstra, tags are an intent-level metadata mechanism used to classify objects and drive automation and reuse. Within a data center blueprint, Apstra supports tagging multiple blueprint objects so operators can apply configuration or policy logic conditionally (for example, applying a connectivity template or a configlet based on a tag match). In this scenario, three valid taggable objects are virtual networks, interfaces, and generic systems.

Virtual network tagging is supported directly from the blueprint's virtual network table, enabling you to label virtual networks (such as ''finance,'' ''pci,'' or ''dev'') and then reference those tags elsewhere in blueprint operations and policy application. Interface tagging is also explicitly supported in the blueprint, allowing you to assign tags to switch interfaces and use those tags to control how templates, assignments, or other intent-driven operations apply to those ports. Finally, generic systems (which are modeled endpoint systems such as servers or external routers represented as ''systems'' in the blueprint) can be tagged so that downstream intent logic can distinguish system roles and apply the correct operations consistently across expansions and changes.

By contrast, property sets are structured data objects used for variable substitution and probe/configlet parameterization, not a primary target for operational tagging in the blueprint UI; and device profiles are catalog artifacts describing hardware/NOS compatibility rather than blueprint objects typically tagged for intent application.

Verified Juniper sources (URLs):

https://www.juniper.net/documentation/us/en/software/apstra5.1/apstra-user-guide/topics/topic-map/tag-interface-add-remove-datacenter.html

https://www.juniper.net/documentation/us/en/software/apstra5.0/apstra-user-guide/topics/task/t

ag-virtual-network-update.html

https://www.juniper.net/documentation/us/en/software/apstra5.0/apstra-user-guide/topics/topic-map/tag-system-add-remove-freeform.html

Question 2

Question Type: MultipleChoice

Which Root Cause Identifier is currently supported in Juniper Apstra software?

Options:

A- Virtual network
B- Connectivity
C- ESI imbalance
D- BGP

Answer:

B

Explanation:

In Juniper Apstra 5.1, Root Cause Identification (RCI) is implemented with a currently supported model focused on connectivity. Practically, this means RCI is designed to take telemetry and state learned from the fabric (for example, interface operational status, LLDP neighbor information, and routing session status) and correlate those signals to determine the most likely underlying cause of a connectivity-impacting event. Within an EVPN-VXLAN IP fabric, many operational symptoms can appear similar at the service layer (endpoints cannot reach each other, routes disappear, overlays degrade), but RCI narrows the problem by correlating evidence across the underlay and control plane.

The ''connectivity'' RCI model targets common failure scenarios that directly break device-to-device reachability, such as a broken link, a miscabled link (wrong LLDP neighbors), or an operator-disabled interface. These conditions often cascade into higher-level symptoms, including BGP sessions dropping over affected links. With Junos v24.4-based leaf-spine fabrics, maintaining stable underlay connectivity is foundational for EVPN signaling and VXLAN forwarding; therefore, Apstra's connectivity-focused RCI helps operators rapidly isolate whether the primary fault lies in physical adjacency, cabling/neighbor correctness, or administrative shutdown---reducing mean time to repair by pointing to the most probable root cause rather than only listing alarms.

Question 3

Question Type: MultipleChoice

You have an EVPN-VXLAN data center IP fabric, with all single-homed hosts/servers. Which two EVPN route types are present in this scenario? (Choose two.)

Options:

A- Type 3
B- Type 7
C- Type 2
D- Type 4

Answer:

A, C

Explanation:

In an EVPN-VXLAN fabric where all hosts are single-homed (each endpoint is attached to only one leaf/VTEP), the EVPN control plane still needs to advertise endpoint reachability and enable BUM handling across the overlay. Two EVPN route types are fundamental in this case: Type 2 and Type 3.

EVPN Route Type 2 (MAC/IP Advertisement) is used to advertise learned MAC addresses and, optionally, associated IP addresses for endpoints connected to the local leaf. This enables remote VTEPs to learn where a given host resides (which VTEP to send unicast traffic to) without relying on data-plane flooding for MAC learning. In Junos v24.4 EVPN-VXLAN deployments, Type 2 routes are the core mechanism for distributing endpoint reachability (MAC and MAC+IP bindings) within the EVPN domain.

EVPN Route Type 3 (Inclusive Multicast Ethernet Tag / IMET) is used to establish the flooding scope for BUM traffic in EVPN-VXLAN. In VXLAN fabrics that use ingress replication (common in data centers), Type 3 routes help build the list of remote VTEPs that should receive replicated BUM traffic for a given segment.

By contrast, Type 4 (Ethernet Segment) routes are associated with EVPN multihoming (ESI-based) and DF election; with only single-homed hosts, Type 4 is not required. Type 7 is not part of the baseline single-homed EVPN-VXLAN host advertisement set in this context.

Verified Juniper sources (URLs):

https://www.juniper.net/documentation/us/en/software/junos/evpn/topics/concept/evpn-bgp-mult

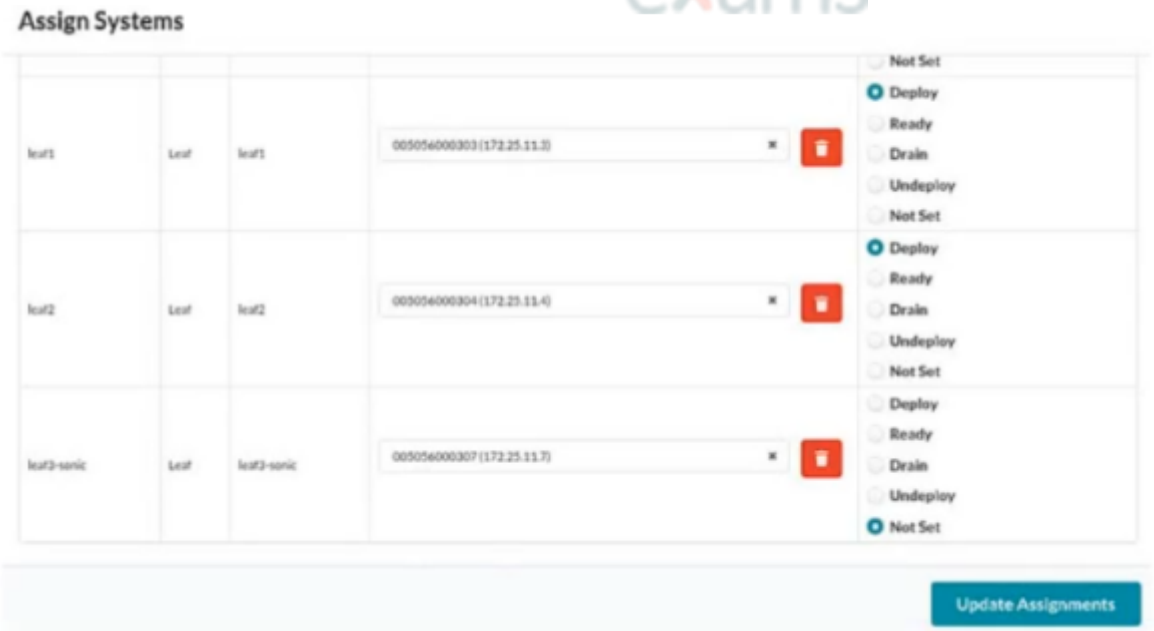
ihoming-overview.html

https://www.juniper.net/documentation/us/en/software/junos/evpn/topics/topic-map/assisted-replication-evpn.html

Question 4

Question Type: MultipleChoice

You have created a blueprint and are in the process of assigning systems. You require the leaf3-sonic device in the blueprint but do not want it to actively participate in the routing of the IP fabric.



In the Juniper Apstra UI, which two modes satisfy this requirement? (Choose two.)

Options:

A- Drain
B- Ready
C- Deploy
D- Undeploy

Answer:

B, D

Explanation:

Apstra deploy modes control how far a device progresses in the configuration lifecycle and whether it becomes active in the fabric. If you must keep leaf3-sonic present in the blueprint (modeled, cabled, and available for future use) but you do not want it to participate in IP-fabric routing, you use modes that keep the device not active.

Ready mode assigns the device to the blueprint and applies only ''Ready (Discovery 2)'' level configuration---hostnames, interface descriptions, and port speed/breakout settings---while explicitly keeping the device out of fabric routing. In this mode, Apstra does not configure routing/BGP or L3 interface addressing for the IP fabric, so the switch is staged and visible for validation (for example, LLDP wiring checks) but does not forward as part of the Clos underlay.

Undeploy mode removes the complete Apstra service configuration from the device. Operationally, this also ensures the device is not active in the fabric. It is commonly used when a device must be retained in the blueprint inventory/topology but should not be participating (for example, temporarily withdrawn, decommission preparation, or held as a spare).

By contrast, Deploy makes the device active (full rendered fabric configuration, including BGP), and Drain is a maintenance state used to gracefully remove traffic from an already-active device rather than a state for keeping it non-participatory from the outset.

Verified Juniper sources (URLs):

https://www.juniper.net/documentation/us/en/software/apstra6.0/apstra-user-guide/topics/topic-map/device-config-lifecycle.html

https://www.juniper.net/documentation/us/en/software/apstra4.2/apstra-user-guide/topics/topic-map/device-config-lifecycle.html

Question 5

Question Type: MultipleChoice

In Juniper Apstra, which statement about resources is correct?

Options:

A- User-defined resources are supported.
B- The scope of a pool is limited to a single blueprint.
C- The scope of a pool can be defined as global or blueprint specific.
D- Only the default resources are supported globally.

Answer:

C

Explanation:

In Apstra 5.1, ''resources'' are the identifier values consumed by the fabric design and rendered into device configuration---examples include ASNs, IP addresses, VNIs, VLAN-related identifiers (where applicable), and similar allocation-driven values. These values are provided through resource pools, which are the authoritative containers Apstra draws from when assigning resources to blueprint roles (for example, leaf ASNs, spine ASNs, loopbacks, point-to-point subnets, and VNI ranges). A key architectural feature is that resource pools are not confined to one blueprint. Apstra supports pools with different scopes to match operational needs: some pools are managed centrally and reused across multiple blueprints, while other pools are created and used within the context of a specific blueprint when you want strict separation and lifecycle alignment with that blueprint.

This is why the correct statement is that a pool's scope can be global or blueprint-specific. Global pools are appropriate when you want consistent allocation policy across fabrics (for example, enterprise-wide ASN ranges). Blueprint-specific pools are appropriate when you want per-fabric independence or when allocations are generated dynamically within the blueprint. This scope behavior is independent of Junos v24.4; Junos receives the final rendered values, but the pool scoping and allocation control are Apstra design-time constructs that ensure deterministic, conflict-free assignments at scale.

Verified Juniper sources (URLs):

https://www.juniper.net/documentation/us/en/software/apstra5.1/apstra-user-guide/topics/concept/resources.html

https://www.juniper.net/documentation/us/en/software/apstra5.1/apstra-user-guide/topics/concept/freeform-resource-management.html

https://www.juniper.net/documentation/us/en/software/apstra5.1/apstra-user-guide/topics/ref/resource-pools-api.html

Question 6

Question Type: MultipleChoice

Which two statements are correct about Juniper Apstra reference designs? (Choose two.)

Options:

A- The Freeform reference design requires knowledge of device CLI.
B- The data center reference design supports only Junos devices.
C- The data center reference design requires knowledge of device CLI.
D- The Freeform reference design supports only Junos devices.

Answer:

A, D

Explanation:

Apstra 5.1 provides multiple reference designs that define how intent is modeled and how configuration is produced. In the Freeform reference design, the network designer is responsible for creating and validating the device configurations using constructs such as device contexts, property sets, and config templates. Because configuration is authored rather than generated from a fixed data-center abstraction set, Freeform inherently requires practical knowledge of the device's configuration model and operational behavior---on Juniper platforms, that means being comfortable with Junos-style configuration and show/verification workflows. This makes statement A correct.

Freeform in Apstra 5.1 is also limited to Juniper devices (Junos and Junos Evolved families). In other words, Freeform blueprints do not provide multi-vendor device onboarding and deployment in that reference design at this software level. That makes statement D correct as written in the question's choices.

By contrast, the data center reference design in Apstra 5.1 is not limited to Junos-only device families; Apstra qualifies multiple NOS families for data center operation via device profiles and packages, and it renders intent accordingly. Therefore statement B is not correct. Finally, while CLI knowledge is always useful for troubleshooting, the data center reference design does not fundamentally require an operator to hand-author device CLI to deploy a standards-based EVPN-VXLAN fabric; that's one of the core benefits of the intent-based data center reference design.

Verified Juniper sources (URLs):

https://www.juniper.net/documentation/us/en/software/apstra5.1/apstra-user-guide/topics/concept/freeform.html

https://www.juniper.net/documentation/us/en/software/apstra4.2/apstra-freeform-4.2.1/apstra-freeform-guide/topics/concept/freeform-overview-and-design.html

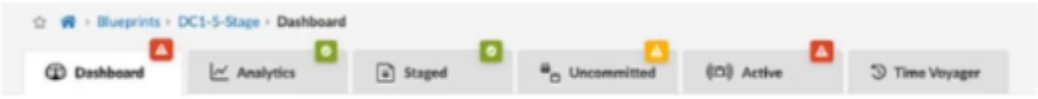
https://www.juniper.net/documentation/us/en/software/apstra5.1/apstra-user-guide/topics/task/device-profile-import-freeform.html

https://www.juniper.net/documentation/us/en/software/apstra5.1/apstra-user-guide/topics/topic-map/devices-qualified.html

Question 7

Question Type: MultipleChoice

You have accessed your deployed blueprint and see the banner shown in the exhibit.



Which two statements are correct in this scenario? (Choose two.)

Options:

A- Devices must be assigned to profiles.
B- There are changes that are not active on the fabric.
C- Resources must be assigned to devices.
D- There are anomalies that must be addressed.

Answer:

B, D

Explanation:

In Apstra 5.1, the top-level blueprint banner uses tab indicators (colored badges) to summarize blueprint status across areas such as Staged, Uncommitted, Active, and Analytics. The presence of an Uncommitted indicator signifies that there are staged modifications that have not yet been committed and therefore are not part of the active, deployed intent. That directly corresponds to the statement that changes exist which are not active on the fabric.

At the same time, the banner shows an Active indicator in an alarm state, which reflects that the running fabric has issues requiring attention---commonly surfaced as anomalies (for example, configuration deviation, interface/link faults, protocol/session issues, or service-impacting conditions). In Apstra's operational model, these issues appear as anomalies that operators should investigate and remediate to restore compliance and health. Therefore, the statement that there are anomalies that must be addressed is also correct.

The remaining options are not implied by this banner alone. Device profile assignment and resource assignment are build-time tasks, but their absence is not what the Uncommitted/Active alert indicators are specifically communicating here. The banner is highlighting uncommitted intent changes and active anomalies that affect the deployed blueprint state and assurance posture.

Verified Juniper sources (URLs):

https://www.juniper.net/documentation/us/en/software/apstra5.1/apstra-user-guide/topics/conce

pt/uncommitted.html

https://www.juniper.net/documentation/us/en/software/apstra5.0/apstra-user-guide/topics/topic-map/anomalies-service-active.html

https://cloudlabs.apstra.com/labguide/Cloudlabs/6.0.0/test-drive-guide/lab1-junos-5_blueprints_.html

Question 8

Question Type: MultipleChoice

What are three port group roles that you are allowed to assign to a logical device? (Select three.)

Options:

A- Leaf
B- Empty
C- Generic
D- Spine
E- Root

Answer:

A, C, D

Explanation:

In Apstra, a logical device abstracts a physical switch's front-panel layout into one or more panels containing port groups. Each port group has a defined speed and one or more roles that describe how those ports are expected to be used in the fabric. These roles are essential because they constrain where ports may be consumed during rack type and template construction (for example, spine-facing vs server-facing vs generic connectivity).

Apstra-supported port group roles include fabric roles such as Spine and Leaf, and endpoint-facing roles such as Generic (commonly used for ports that connect to servers or external generic systems). Assigning Leaf and Spine roles ensures Apstra can correctly validate and render intent for uplinks and interconnects in a three-stage Clos or larger topologies. Assigning Generic indicates ports that can be used for non-fabric connections (such as server links, external routers modeled as generic systems, or other non-managed endpoints).

The options Empty and Root are not valid Apstra port group roles in the logical device model;

Apstra uses other explicit role names (for example, Access, Peer, Unused, Generic, Leaf, Spine, Superspine depending on design type and version). In Junos v24.4 EVPN-VXLAN fabrics, getting these roles correct is foundational because Apstra relies on them to place underlay and overlay configuration onto the right interfaces with predictable results.

Verified Juniper sources (URLs):

https://www.juniper.net/documentation/us/en/software/apstra4.2/apstra-user-guide/topics/concept/logical-devices.html

https://www.juniper.net/documentation/us/en/software/jvd/jvd-collapsed-dc-fabric-juniper-apstra-access-switches/configuration_walkthrough.html

Question 9

Question Type: MultipleChoice

You are considering the bridged overlay EVPN-VXLAN architecture. In this scenario, how many VLANs would be enabled in the VLAN-based service type at the MAC-VRF EVPN instance level?

Options:

A- One VLAN
B- Two VLANs
C- 4000 VLANs
D- Four VLANs

Answer:

A

Explanation:

In Junos EVPN, the service type determines how VLANs (broadcast domains) are mapped into EVPN constructs. With VLAN-based service, the mapping is one-to-one: a single VLAN (single broadcast domain) maps to a single EVPN instance (EVI), resulting in a separate bridge table per VLAN. When you implement EVPN-VXLAN in a bridged overlay (L2 extension over an L3 underlay), leaf devices act as VTEPs and encapsulate VLAN traffic into VXLAN using the associated VNI, but the service-type mapping rule still applies.

At the MAC-VRF hierarchy, you configure the EVPN-VXLAN parameters and choose the service type for that EVPN instance. If the service type selected is VLAN-based, then the EVI represents exactly one VLAN at that EVPN instance level. If you need multiple VLANs under the same

MAC-VRF/EVI construct, Junos provides alternative service types (for example, VLAN-aware or VLAN-bundle) specifically intended to associate multiple VLANs with a single EVPN instance; that is not what VLAN-based service does.

Therefore, in a bridged overlay EVPN-VXLAN design using VLAN-based service at the MAC-VRF EVPN instance level, the correct answer is one VLAN. This remains true regardless of how many total VLANs exist across the fabric; each VLAN-based EVI handles a single VLAN.

Verified Juniper sources (URLs):

https://www.juniper.net/documentation/us/en/software/junos/evpn/topics/concept/evpn-based-services.html

https://www.juniper.net/documentation/us/en/software/junos/evpn/topics/concept/mac-vrf-routing-instance-overview.html

Question 10

Question Type: MultipleChoice

What are three phases of the Juniper Apstra data center life cycle? (Choose three.)

Options:

A- Configuration
B- Design
C- Installation
D- Operational phase
E- Deployment

Answer:

B, D, E

Explanation:

Juniper Apstra describes data center fabric management as a full lifecycle that spans three core phases: Design (Day 0), Deployment (Day 1), and Operations (Day 2). These phases map directly to how Apstra applies intent-based networking to a data center fabric.

In the Design phase, you model the intended architecture---templates (3-stage or 5-stage), rack types, logical devices, interface maps, resource pools, and high-level constructs such as routing zones and virtual networks. The objective is to capture intent in a vendor-agnostic way while

ensuring consistency and validation before anything is pushed.

In the Deployment phase, Apstra turns the modeled intent into device-level implementation. This includes onboarding systems, assigning device profiles, allocating resources, rendering configurations, and pushing the resulting configuration to switches so the IP fabric becomes operational. This is where Junos v24.4 leaf/spine nodes receive underlay and overlay configuration generated from the blueprint.

In the Operational phase, Apstra continuously validates the running network against intent using telemetry and analytics (IBA), detects deviations and anomalies, supports maintenance workflows (such as drain), and provides troubleshooting tools (queries, time-series utilization, and configuration compliance).

''Configuration'' and ''installation'' are activities that occur within the lifecycle, but the lifecycle phases themselves are Design, Deployment, and Operations.

Question 11

Question Type: MultipleChoice

In Juniper Apstra terminology, to which network operating system concept does a routing zone refer?

Options:

A- IRB
B- VRF
C- VLAN
D- Access list

Answer:

B

Explanation:

In Apstra 5.1, a routing zone is the primary construct used to represent an L3 domain for multitenant isolation. In traditional network operating system terms, that maps to a VRF (Virtual Routing and Forwarding instance). Each routing zone is placed ''in its own VRF,'' which provides independent routing tables and isolates IP traffic so that different tenants can reuse overlapping IP subnets without conflict. This is central to modern EVPN-VXLAN data center design, where tenants typically require clean separation of routing and policy boundaries.

Within a routing zone, you can create one or more virtual networks (often mapped to VXLAN segments) that provide L2 extension across racks while still being contained by the tenant's VRF. If L3 gateway services are enabled for those virtual networks, their gateway interfaces (for example, IRB interfaces on Junos v24.4 leaf switches) are associated with the routing zone's VRF so that inter-subnet routing occurs within the tenant boundary.

This terminology distinction is important: an IRB is an interface construct used to provide L3 gateway functionality for a VLAN/VXLAN segment; a VLAN is a Layer 2 segmentation mechanism; and an access list is a policy enforcement tool. A routing zone, however, defines the tenant's L3 routing context, which is precisely what a VRF provides on Junos.

Verified Juniper sources (URLs):

https://www.juniper.net/documentation/us/en/software/apstra5.0/apstra-user-guide/topics/concept/routing-zones.html

https://www.juniper.net/documentation/us/en/software/apstra4.2/apstra-user-guide/topics/concept/routing-zones.html

Question 12

Question Type: MultipleChoice

You are creating a new security policy using Juniper Apstra.



Referring to the exhibit, which application point should you select to allow or deny traffic to or from a particular VRF?

Options:

A- Routing Zone
B- External Endpoint
C- Internal Endpoint
D- Virtual Network

Answer:

A

Explanation:

In Apstra 5.1, multitenancy is modeled using routing zones, which map directly to the network operating system concept of a VRF. A VRF is an isolated Layer 3 routing instance with its own routing table and forwarding context, and Apstra's routing zone is the intent-based abstraction used to define and manage that isolation consistently across the fabric. Therefore, if your goal is to allow or deny traffic to or from a particular VRF, you must select Routing Zone as the security policy application point.

This choice enables you to express policy at the tenant boundary (VRF boundary) rather than at

a single segment boundary. In EVPN-VXLAN data center fabrics, a tenant VRF commonly contains multiple virtual networks (VXLAN segments) and their associated IRB gateways on the leaf switches. Applying policy at the routing-zone level allows Apstra to compile intent and deploy enforcement consistently where traffic enters or exits that VRF context---typically as ACL constructs rendered as Junos firewall filters on the appropriate interfaces (for example, IRB interfaces for east-west controls or border interfaces for north-south controls).

By contrast, selecting Virtual Network targets a single segment (not the whole VRF), and Internal/External Endpoint targets specific endpoints or endpoint groups rather than the VRF-wide policy boundary. Hence, Routing Zone is the correct application point when policy scope is the VRF.

To Get Premium Files for JN0-481 Visit

https://www.p2pexams.com/products/jn0-481

For More Free Questions Visit

https://www.p2pexams.com/juniper/pdf/jn0-481

