Check the Free demo of our JN0-480 Exam Dumps with 67 Questions [Q29-Q46

Check the Free demo of our JN0-480 Exam Dumps with 67 Questions
Clear your concepts with JN0-480 Questions Before Attempting Real exam

To pass the Juniper JN0-480 exam, candidates need to have a good understanding of data center architectures and technologies, as well as hands-on experience with Juniper Networks data center platforms. JN0-480 exam consists of 65 multiple-choice questions that need to be completed within 90 minutes. Candidates who pass the exam will gain a comprehensive understanding of data center technologies and can demonstrate their skills and knowledge to potential employers. The JNCIS-DC certification is a valuable asset for network engineers, network administrators, and other IT professionals who work with data center technologies.

Q29. Exhibit.

Referring to the exhibit, how do you display the IPv6 subnets lot all of the listed VXLANs?
* IPv6 subnets ate shown when each VXLAN is selected individually.
* Select Columns, then select IPv6 Subnet.
* Select all VXLANs. and the IPv6 Subnets column will appear
* An IPv6 Subnets column is not shown, indicating that no VXLAN has an assigned IPv6 subnet

Referring to the exhibit, the image shows a user interface of the Juniper Apstra software application, which is used for network management and configuration. The image shows the Virtual Networks table under the Resources menu, which displays the details of the VLANs and VXLANs in the network. The table has 11 columns, but only 9 are visible in the image. The other two columns are IPv6 Connectivity and IPv6 Subnet, which are hidden by default. To display the IPv6 subnets for all of the listed VXLANs, the user needs to select Columns, then select IPv6 Subnet. This will show the IPv6 Subnet column in the table, which will display the IPv6 addresses assigned to the VXLANs from the IPv6 pools. For more information, see Virtual Networks (Resources). References:

* Virtual Networks (Resources)

* IPv6 Pools (Resources)

* Apstra User Guide

Q30. Using Juniper Apstra. which component is defined in a template?
* the leaf-to-spine interconnection
* the speed of the links between the spine devices and the leaf devices
* the number of spine devices in a topology
* the definition of IP pools

According to the Juniper documentation1, a template is a configuration template that defines a network’s policy intent and structure. A template can be either rack-based or pod-based, depending on the type and number of racks and pods in the network design. A template includes the following details:

* Policies: These are the parameters that apply to the entire network, such as the overlay control protocol, the ASN allocation scheme, and the underlay type.

* Structure: This is the physical layout of the network, such as the type and number of racks, pods, spines, and leaves. The structure also defines the leaf-to-spine interconnection, which is the number and type of links between the leaf and spine devices. The leaf-to-spine interconnection can be either single or dual, depending on the redundancy and bandwidth requirements.

Therefore, the correct answer is A. the leaf-to-spine interconnection. This is a component that is defined in a template, as it determines the physical connectivity of the network. The speed of the links, the number of spine devices, and the definition of IP pools are not components that are defined in a template, as they are either derived from the device profiles, the resource pools, or the blueprint settings. References: Templates Introduction | Apstra 4.2 | Juniper Networks

Q31. Which two statements are correct about Time Voyager? {Choose two.)
* Time Voyager retains all of the blueprint revisions from the last Juniper Apstra backup.
* Time Voyager retains the five most recent blueprint commits.
* Time Voyager retains the last ten blueprint commits.
* Time Voyager retains up to twenty-five saved revisions.

Time Voyager is a feature of Juniper Apstra that allows you to restore previous revisions of a blueprint, which is a logical representation of your network design and configuration. Time Voyager automatically saves the five most recent blueprint commits, which are the changes that you apply to the network. You can also manually save up to twenty-five revisions by keeping them, which prevents them from being overwritten by new commits. Therefore, the correct answer is B and D. Time Voyager retains the five most recent blueprint commits and Time Voyager retains up to twenty-five saved revisions. References: Time Voyager | Apstra 4.1 | Juniper Networks, Time Voyager Introduction | Apstra 4.2 | Juniper Networks, Juniper Apstra at a Glance | Flyer

Q32. You have designed your fabric in Juniper Apstra prior to deploying the network devices.

Which Apstra element in the Staged tab would be used to assist the team that is installing and cabling the devices?
* Connectivity Templates
* Virtual Networks table
* Managed Devices list
* Links table

The Links table in the Staged tab shows the physical connections between the devices in the fabric. It provides information such as the source and destination device names, hostnames, serial numbers, roles, interfaces, and link status. The Links table can be used to assist the team that is installing and cabling the devices by verifying that the devices are connected correctly and that the links are operational. The Links table can also be used to troubleshoot any connectivity issues that may arise during the installation process. For more information, see Links (Staged). References:

* Links (Staged)

* Topology (Staged)

* Staged

Q33. Which fabric type should be chosen in a template to create a five-stage Clos?
* pod-based
* rack-based
* circuit switched
* collapsed

Q34. Exhibit.

You connect two single-homed servers using Juniper Apstra as shown in the exhibit. You are using the ERB design blueprint with two virtual networks in a common routing zone.

In this scenario, which two types of VXLAN tunnels will be automatically created by the EVPN control plane? (Choose two.)
* EVPN signaled route Type-8 VXLAN tunnels
* EVPN signaled route Type-3 VXLAN tunnels
* EVPN signaled route Type-6 VXLAN tunnels
* EVPN signaled route Type-2 VXLAN tunnels

According to the Juniper documentation1, EVPN route Type-3 is used to advertise the IP address of the VTEP and the VNIs that it supports. This allows the VTEPs to discover each other and form VXLAN tunnels for the VNIs that they have in common. EVPN route Type-2 is used to advertise the MAC and IP addresses of the hosts connected to the VTEPs. This allows the VTEPs to learn the MAC-to-IP bindings and the MAC-to-VTEP mappings for the hosts in the same VNI. Therefore, these two types of VXLAN tunnels will be automatically created by the EVPN control plane when using Juniper Apstra with the ERB design blueprint and two virtual networks in a common routing zone. References: Example: Configure an EVPN-VXLAN Centrally-Routed Bridging Fabric

Q35. Exhibit.

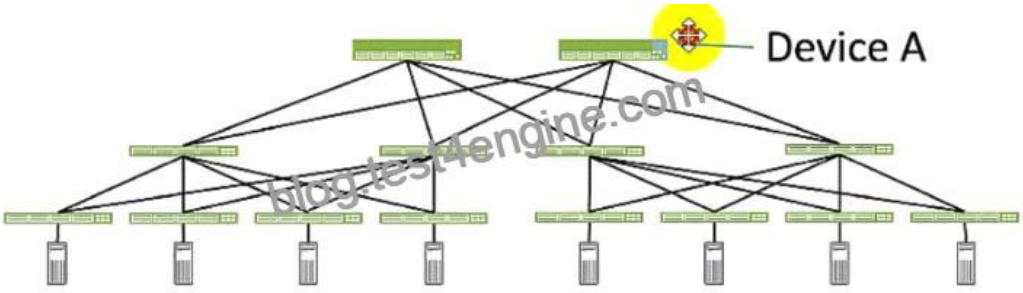

Referring to the exhibit, which role does Device A serve in an IP fabric?
* leaf
* spine

* super spine
* server

Device A serves as a spine in an IP fabric. An IP fabric is a network architecture that uses a spine-leaf topology to provide high performance, scalability, and reliability for data center networks. A spine-leaf topology consists of two layers of devices: spine devices and leaf devices. Spine devices are the core devices that interconnect all the leaf devices using equal-cost multipath (ECMP) routing. Leaf devices are the edge devices that connect to the servers, storage, or other network devices. In the exhibit, Device A is connected to four leaf devices using multiple links, which indicates that it is a spine device. The other options are incorrect because:

* A. leaf is wrong because a leaf device is an edge device that connects to the servers, storage, or other network devices. In the exhibit, Device A is not connected to any servers, storage, or other network devices, but only to four leaf devices, which indicates that it is not a leaf device.

* C. super spine is wrong because a super spine device is a higher-level device that interconnects multiple spine devices in a large-scale IP fabric. A super spine device is typically used when the number of leaf devices exceeds the port density of a single spine device. In the exhibit, Device A is not connected to any other spine devices, but only to four leaf devices, which indicates that it is not a super spine device.

* D. server is wrong because a server device is a compute or storage device that connects to a leaf device in an IP fabric. A server device is typically the end host that provides or consumes data in the network.

In the exhibit, Device A is not connected to any leaf devices, but only to four leaf devices, which indicates that it is not a server device. References:

* IP Fabric Underlay Network Design and Implementation

* IP Fabric Overview

* IP Fabric Architecture

Q36. What are two system-defined user roles that are available in Juniper Apstra? (Choose two.)

* authorized
* root
* viewer
* user

Juniper Apstra provides four system-defined user roles that are available in the Apstra GUI environment. They are: administrator, device_ztp, viewer, and user1. Based on the web search results, we can infer the following statements:

* viewer: This role includes permissions to only view various elements in the Apstra system, such as blueprints, devices, design, resources, external systems, platform, and others. Users with this role cannot create, edit, or delete any element12.

* user: This role includes permissions to view and edit various elements in the Apstra system, such as blueprints, devices, design, resources, external systems, platform, and others. Users with this role cannot create or delete any element12.

* authorized: This is not a system-defined user role in Juniper Apstra. It is a term used to describe users who have been authenticated by an external system, such as LDAP, Active Directory, TACACS+, or RADIUS3.

* root: This is not a system-defined user role in Juniper Apstra. It is a term used to describe the superuser account on a Linux system, which has full access to all commands and files. Creating a user in the Apstra GUI does not provide that user access to the Apstra platform via SSH. To access the Apstra platform via SSH, you must create a local Linux system user4. References:

* User / Role Management Introduction

* User/Role Management (Platform)

* AAA Providers

* User Profile Management

Q37. Exhibit.

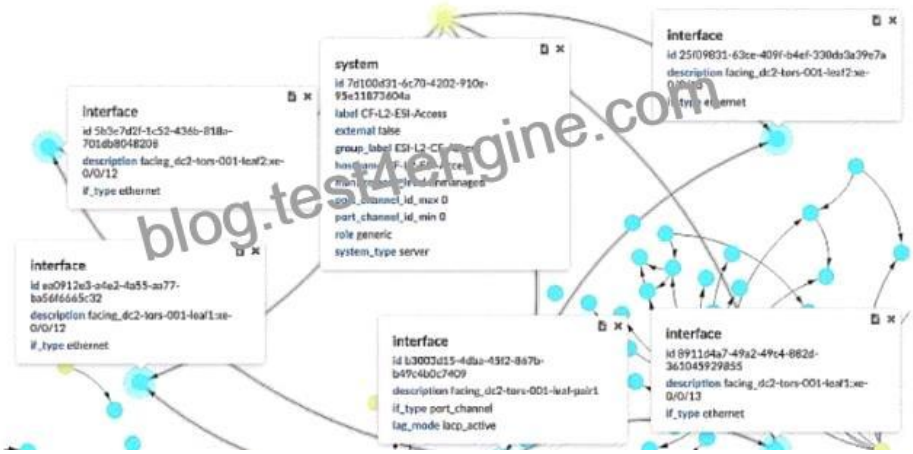

Which two statements ate correct about the information shown in the exhibit? (Choose two.)
* The system is fully managed by Juniper Apstra.
* The device shown is a generic system.
* Four physical interfaces exist in a LAG facing the leaf pair.
* The physical ports are not part of the LAG.

According to the Juniper documentation1, a generic system is a device that is not managed by Juniper Apstra and does not have a specific role or type assigned to it. A generic system can be used to represent a server, a firewall, a load balancer, or any other device that is not part of the fabric. In the exhibit, the device shown is a generic system, as indicated by its role, system type, and management level. Therefore, the correct answer is B: The device shown is a generic system.

According to the Juniper documentation2, a LAG is a link aggregation group that bundles multiple physical interfaces into a single logical interface. A LAG can provide increased bandwidth, redundancy, and load balancing for the network traffic. In the exhibit, the device shown has four physical interfaces that are part of a LAG, as indicated by their description and li_type. The LAG is facing the leaf pair, which are the two switches that connect to the device. Therefore, the correct answer is C. Four physical interfaces exist in a LAG facing the leaf pair. References: Generic Systems (Datacenter Design), Form LAG | Apstra 4.1 | Juniper Networks

Q38. Multilenancy for applications is achieved by creating virtual networks (VNs) within which construct?
* security policy
* routing table
* connectivity template
* routing zone

According to the Juniper documentation1, a routing zone is an L3 domain, the unit of tenancy in multi-tenant networks. You create routing zones for tenants to isolate their IP traffic from one another, thus enabling tenants to re-use IP subnets. In addition to being

in its own VRF, each routing zone can be assigned its own DHCP relay server and external system connections. You can create one or more virtual networks within a routing zone, which means a tenant can stretch its L2 applications across multiple racks within its routing zone. Therefore, the correct answer is D. routing zone. A routing zone is the construct within which you create virtual networks to achieve multitenancy for applications. References: Routing Zones

Q39. What does EVPN use lo identity which remote leaf device advertised the EVPN route?
* a route distinguisher value
* a community tag
* a route target value
* a VRF target value

EVPN uses a route distinguisher (RD) value to identify which remote leaf device advertised the EVPN route.

An RD is a 64-bit value that is prepended to the EVPN NLRI to create a unique VPNv4 or VPNv6 prefix. The RD value is usually derived from the IP address of the PE that originates the EVPN route. By comparing the RD values of different EVPN routes, a PE can determine which remote PE advertised the route and which VRF the route belongs to. The other options are incorrect because:

* B. a community tag is wrong because a community tag is an optional transitive BGP attribute that can be used to group destinations that share some common properties. A community tag does not identify the source of the EVPN route.

* C. a route target value is wrong because a route target (RT) value is an extended BGP community that is used to control the import and export of EVPN routes between VRFs. An RT value does not identify the source of the EVPN route.

* D. a VRF target value is wrong because there is no such thing as a VRF target value in EVPN. A VRF is a virtual routing and forwarding instance that isolates the IP traffic of different VPNs on a PE. A VRF does not have a target value associated with it. References:

* EVPN Fundamentals

* RFC 9136 – IP Prefix Advertisement in Ethernet VPN (EVPN)

* EVPN Type-5 Routes: IP Prefix Advertisement

* Understanding EVPN Pure Type 5 Routes

Q40. Which attribute enables Juniper Apstra to scale and manage thousands of devices with a single server instance?
* Apstra is installed as a cloud resource.
* Apstra is based on NGINX.
* Apstra is available as an OVA.
* Apstra is a distributed state system.

The attribute that enables Juniper Apstra to scale and manage thousands of devices with a single server instance is that Apstra is a distributed state system. This means that Apstra uses a graph database to store the network topology and configuration data in a distributed and replicated manner across multiple server nodes.

This allows Apstra to handle large-scale networks with high performance, reliability, and availability. Apstra also uses a stateful orchestration engine that ensures the network state is always consistent with the intent of the blueprint, which is the logical representation of thenetwork design and behavior. Apstra can automatically detect and resolve any discrepancies between the desired and actual network state, as well as handle any changes or failures in the network. The other options are incorrect because:

* A. Apstra is installed as a cloud resource is wrong because Apstra can be installed either as a cloud resource or as an on-premises resource. Apstra is available as a virtual machine image that can be deployed on various hypervisors, such as VMware ESXi,

QEMU/KVM, Microsoft Hyper-V, or Oracle VirtualBox. Apstra can also be deployed on public cloud platforms, such as Amazon Web Services (AWS) or Microsoft Azure. However, the installation method does not affect the scalability of Apstra, which is determined by the distributed state system architecture.

* B. Apstra is based on NGINX is wrong because Apstra is not based on NGINX, but on Python and Django. NGINX is a web server and reverse proxy that Apstra uses to serve the web user interface and the REST API. However, NGINX is not the core component of Apstra, and it does not affect the scalability of Apstra, which is determined by the distributed state system architecture.

* C. Apstra is available as an OVA is wrong because Apstra is available as an OVF, not an OVA. An

* OVF (Open Virtualization Format) is a standard format for packaging and distributing virtual machine images. An OVA (Open Virtual Appliance) is a single file that contains the OVF and the virtual disk images. Apstra provides an OVF file that can be imported into various hypervisors, such as VMware ESXi, QEMU/KVM, Microsoft Hyper-V, or Oracle VirtualBox. However, the availability of Apstra as an OVF does not affect the scalability of Apstra, which is determined by the distributed state system architecture. References:

* JUNIPER APSTRA ARCHITECTURE

* Apstra Server Requirements/References

* Juniper Networks Apstra 4.0 enhances the experience of users and operators

Q41. In Juniper Apstra. which statement is correct?
* VMware anomaly detection is on by default.
* VMware anomaly detection requires a vCenter server configured under External Systems
* VMware anomaly detection requires a VMware hypervisor with exports enabled.
* VMware anomaly detection requires an Apstra server running on VMware.

VMware anomaly detection is a feature of Apstra that provides visibility and validation of the virtual network settings and the physical network settings in a VMware vSphere environment. To enable this feature, Apstra requires a connection to a vCenter server that manages the ESX/ESXi hosts and the VMs connected to the Apstra-managed leaf switches. The vCenter server must be configured under External Systems in the Apstra web interface, and the vCenter integration must be staged and committed in the blueprint. This allows Apstra to collect information about VMs, ESX/ESXi hosts, port groups, and VDS, and to flag any inconsistencies or mismatches that might affect VM connectivity. The other options are incorrect because:

* VMware anomaly detection is not on by default. It must be enabled by configuring a vCenter server under External Systems and adding a virtual infra to the blueprint.

* VMware anomaly detection does not require a VMware hypervisor with exports enabled. It only requires LLDP transmit to be enabled on the VMware distributed virtual switch to associate host interfaces with leaf interfaces.

* VMware anomaly detection does not require an Apstra server running on VMware. It can run on any supported platform, such as Linux, Windows, or Docker. References:

* VMware vCenter/vSphere Virtual Infra

* Anomalies (Service)

* A Better Experience: VMware + Juniper Apstra

Q42. In the Juniper Apstra design phase, which object dictates port count, port speed, and how the ports would be used?
* logical devices
* rack type
* network devices
* interface map

Interface maps are objects that map interfaces between logical devices and physical hardware devices in the Juniper Apstra design phase. They dictate port count, port speed, and how the ports would be used for achieving the intended network configuration rendering. Interface maps also allow you to select device ports, transformations, and interfaces, provision breakout ports, and disable unused ports. For more information, see Interface Maps (Datacenter Design). References:

* Interface Maps (Datacenter Design)

* Design

* Interface Maps Introduction

Q43. In the Juniper Apstra Ul. which three resources are assigned under the Resources menu? (Choose three.)
* VTEP pools
* ASN pools
* VNI pools
* logical device pools
* IP address pools

In the Juniper Apstra UI, the Resources menu allows you to create and manage global and local resources that are used for various elements of the network design and configuration. The Resources menu includes the following three types of resources that can be assigned to the network devices and virtual networks:

* ASN pools: These are pools of autonomous system numbers (ASNs) that are used for the underlay routing protocol (EBGP) between the leaf and spine devices. You can create ASN pools with either

2-byte or 4-byte ASNs, and assign them to the logical devices in the blueprint.

* VNI pools: These are pools of virtual network identifiers (VNIs) that are used for the overlay network (VXLAN) between the end hosts. You can create VNI pools with a range of VNIs, and assign them to the virtual networks in the blueprint.

* IP address pools: These are pools of IPv4 or IPv6 addresses that are used for various purposes in the network, such as the loopback addresses for the devices, the IP prefixes for the virtual networks, the host

* IP addresses for the end hosts, and the gateway IP addresses for the IRB interfaces. You can create IP address pools with a range of IP addresses, and assign them to the logical devices and virtual networks in the blueprint.

The following two types of resources are not assigned under the Resources menu:

* VTEP pools: These are not resources that can be created or assigned by the user. VTEPs are VXLAN tunnel endpoints that are automatically generated by the Apstra server based on the loopback IP addresses of the devices. VTEPs are used as the source and destination IP addresses for the VXLAN tunnels in the overlay network.

* Logical device pools: These are not resources that can be created or assigned by the user. Logical device pools are groups of logical devices that share the same role, interface map, and resource assignments in the blueprint. Logical device pools are used to simplify the network design and configuration by applying the same settings to multiple devices.

References:

* Resources Introduction

* ASN Pools (Resources)

* VNI Pools (Resources)

* IP Address Pools (Resources)

Q44. You are adding a new switch to Juniper Apstra software. The Managed Devices page shows the “0 OS-Quarantined” status. What is the proper next step to make the device ready for use in a blueprint?
* Acknowledge the device.
* Take the device out of maintenance mode.
* Install the agent for the device.
* Take the device out of drain state.
When a new switch is added to Juniper Apstra software, it initially shows the “0 OS-Quarantined” status, which means that the device is not yet managed by Apstra and has not been assigned to any blueprint. The proper next step to make the device ready for use in a blueprint is to acknowledge the device, which is a manual action that confirms the device identity and ownership. Acknowledging the device changes its status to

“OOS-Ready”, which means that the device is ready to be assigned to a blueprint and deployed12. References

:

* Managing Devices

* AOS Device Configuration Lifecycle

Q45. Using the Juniper Apstra multitenancy capabilities, which approach will allow a tenant to interconnect two different routing zones?
* Interconnection is the default behavior.
* Use interconnection through the fabric spine nodes.
* Interconnection cannot be enabled.
* Use interconnection through an external gateway.
According to the Juniper documentation1, a routing zone is an L3 domain, the unit of tenancy in multi-tenant networks. You create routing zones for tenants to isolate their IP traffic from one another, thus enabling tenants to re-use IP subnets. In addition to being in its own VRF, each routing zone can be assigned its own DHCP relay server and external system connections. You can create one or more virtual networks within a routing zone, which means a tenant can stretch its L2 applications across multiple racks within its routing zone. For virtual networks with Layer 3 SVI, the SVI is associated with a Virtual Routing and Forwarding (VRF) instance for each routing zone isolating the virtual network SVI from other virtual network SVIs in other routing zones. If you’re using multiple routing zones, external system connections must be from leaf switches in the fabric. Routing between routing zones must be accomplished with external systems. Therefore, the correct answer is D. Use interconnection through an external gateway.
References: Routing Zones

Q46. Which statement about Juniper Apstra role-based access control is correct?
* The viewer role is predefined and can be deleted.
* The administrator role can see all permissions.
* The user role can create roles.
* The administrator role is the only predefined role.

Juniper Apstra role-based access control (RBAC) is a feature that allows you to specify access permissions for different users based on their roles. RBAC servers are remote network servers that authenticate and authorize network access based on roles assigned to individual users within an enterprise1. Juniper Apstra has four predefined user roles: administrator, device_ztp, user, and viewer2. The administrator role is the most powerful role, and it can see all permissions and perform all actions in the Apstra software application. The administrator role can also create, clone, edit, and delete user roles, except for the four predefined user roles, which cannot be modified2. Therefore, the statement that the administrator role can see all permissions is correct.

The following three statements are incorrect in this scenario:

* The viewer role is predefined and can be deleted. This is not true, because the viewer role is one of the four predefined user roles, and it cannot be deleted. The viewer role is the most restricted role, and it can only view the network information and configuration, but not make any changes2.

* The user role can create roles. This is not true, because the user role is one of the four predefined user roles, and it cannot create roles. The user role can perform most of the network configuration and management tasks, but it cannot access the platform settings or the user management features2.

* The administrator role is the only predefined role. This is not true, because there are four predefined user roles, not just one. The other three predefined user roles are device_ztp, user, and viewer2.

References:

* Providers – Apstra 3.3.0 documentation

* User/Role Management (Platform)

The JN0-480 exam covers a range of topics related to data center networking, including network architecture, storage, virtualization, and automation. JN0-480 exam is divided into multiple sections, each of which tests a specific area of knowledge. Candidates must demonstrate their proficiency in each section to pass the exam and earn their certification.