



Free Questions for JN0-480 by certscare

Shared by Pittman on 15-04-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

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

Options:

A- authorized

B- root

C- viewer

D- user

Answer:

C, D

Explanation:

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. Reference:

User / Role Management Introduction

User/Role Management (Platform)

AAA Providers

User Profile Management

Question 2

Question Type: MultipleChoice

You want to make a widget appear on the main dashboard in Juniper Apstra

a. In this scenario, which statement is correct?

Options:

- A-** When creating the widget, select the Add to Blueprint Dashboard option.
- B-** On the blueprint dashboard, click on the Add Widget option.
- C-** Widgets automatically appear on the blueprint dashboard.
- D-** Set the Default toggle switch to On for the desired widget.

Answer:

D

Explanation:

In Juniper Apstra, a widget is a graphical element that displays data from an intent-based analytics (IBA) probe. A widget can be used to monitor different aspects of the network and raise alerts to any anomalies. A widget can be viewed by itself or added to an analytics dashboard. A dashboard is a collection of widgets that can be customized and organized according to the user's preference¹.

The main dashboard in Juniper Apstra is the blueprint dashboard, which is the default view that shows the network information and configuration for the active blueprint. A blueprint is a logical representation of the network design and intent. The blueprint dashboard can display the system-generated dashboards, the user-generated dashboards, and the individual widgets that are relevant to the network².

To make a widget appear on the main dashboard in Juniper Apstra, the user needs to set the Default toggle switch to On for the desired widget. This will add the widget to the blueprint dashboard, where it can be viewed along with other network information. The user can also remove the widget from the blueprint dashboard by setting the Default toggle switch to Off for the widget³. Therefore, the statement D is correct in this scenario.

The following three statements are incorrect in this scenario:

When creating the widget, select the Add to Blueprint Dashboard option. This is not true, because there is no such option when creating a widget in Juniper Apstra. The user can only select the widget type, the probe, and the display mode when creating a widget⁴. To add the widget to the blueprint dashboard, the user needs to set the Default toggle switch to On for the widget after creating it³.

On the blueprint dashboard, click on the Add Widget option. This is not true, because there is no such option on the blueprint dashboard in Juniper Apstra. The user can only view, edit, or delete the existing widgets and dashboards on the blueprint dashboard². To add a widget to the blueprint dashboard, the user needs to set the Default toggle switch to On for the widget from the widgets table view³.

Widgets automatically appear on the blueprint dashboard. This is not true, because widgets do not automatically appear on the blueprint dashboard in Juniper Apstra. The user needs to manually add the widgets to the blueprint dashboard by setting the Default toggle switch to On for the widgets that they want to see on the blueprint dashboard³. The only exception is the widgets that are part of the system-generated dashboards, which are automatically created and added to the blueprint dashboard based on the state of the active blueprint².

Widgets Overview

[Blueprint Summaries and Dashboard](#)

[Widgets Introduction](#)

[Create Widget](#)

Question 3

Question Type: MultipleChoice

Using Juniper Apstr

a. which component is defined in a template?

Options:

A- the leaf-to-spine interconnection

B- the speed of the links between the spine devices and the leaf devices

C- the number of spine devices in a topology

D- the definition of IP pools

Answer:

A

Explanation:

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. Reference: [Templates Introduction | Apstra 4.2 | Juniper Networks](#)

Question 4

Question Type: MultipleChoice

Which attribute enables Juniper Apstra to scale and manage thousands of devices with a single server instance?

Options:

- A- Apstra is installed as a cloud resource.
- B- Apstra is based on NGINX.
- C- Apstra is available as an OVA.
- D- Apstra is a distributed state system.

Answer:

D

Explanation:

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 the network 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. Reference:

[JUNIPER APSTRA ARCHITECTURE](#)

[Apstra Server Requirements/Reference](#)

[Juniper Networks Apstra 4.0 enhances the experience of users and operators](#)

Question 5

Question Type: MultipleChoice

Within Managed Devices in the Juniper Apstra UI, you notice that several devices have the OOS-Quarantined status. The devices cannot be added to any blueprint. Which action would solve this problem?

Options:

- A- Acknowledge the device.
- B- Fix the hardware issues with the quarantined devices.
- C- Install the agent, even though connectivity is established.
- D- Upload a new pristine configuration.

Answer:

A

Explanation:

When an agent installation is successful, devices are placed into the Out of Service Quarantined (OOS-QUARANTINED) state using the Juniper Apstra UI. This state means that the device is not yet managed by Apstra and has not been assigned to any blueprint. The device configuration at this point is called Pristine Config. To make the device ready for use in a blueprint, you need to acknowledge the device, which is a manual action that confirms the device identity and ownership. Acknowledging the device changes its status to Out of Service Ready (OOS-READY)¹². Reference:

Managing Devices

AOS Device Configuration Lifecycle

Question 6

Question Type: MultipleChoice

What is the purpose of a Juniper Apstra rack?

Options:

- A-** It stores information on how pods connect to super spines.
- B-** It stores information on how leaf nodes connect to generic devices

C- It stores IP address and ASN pool information.

D- It stores device port data rates and vendor information.

Answer:

B

Explanation:

A Juniper Apstra rack is a physical entity that contains one or more network devices, such as leaf nodes, access switches, or generic systems. A rack is used to organize and manage the network devices in the Apstra software application. A rack has the following characteristics:

It stores information on how leaf nodes connect to generic devices. This is because a rack can include generic systems, which are devices that are not managed by Juniper Apstra, but are connected to the network. A generic system can be a server, a firewall, a load balancer, or any other device that has a network interface. A rack stores the information on how the leaf nodes, which are the devices that provide access to the end hosts, connect to the generic devices, such as the port number, the link speed, the LAG mode, and the roles¹.

It has a rack type, which defines the type and number of leaf devices, access switches, and/or generic systems that are used in the rack. A rack type is a resource that is created in the data center design phase, and it does not specify the vendor or the model of the devices. A rack type can be predefined or custom-made, and it can be used to create multiple racks with the same structure and configuration².

It has a rack build, which assigns the specific vendor and model of the devices to the rack. A rack build is created in the staged phase, and it uses the rack type as a template. A rack build can also assign the resources, such as the IP addresses, the ASNs, and the VNIs, to the devices in the rack³.

It has a rack deployment, which applies the network configuration and services to the devices in the rack. A rack deployment is performed in the active phase, and it uses the rack build as a reference. A rack deployment can also monitor the network performance and compliance of the devices in the rack⁴.

The following three statements are incorrect in this scenario:

It stores information on how pods connect to super spines. This is not true, because a rack does not store any information on the pod or the super spine level of the network. A pod is a cluster of leaf and spine devices that form a 3-stage Clos topology, and a super spine is a device that connects multiple pods in a 5-stage Clos topology. A rack only stores information on the leaf and the access level of the network¹.

It stores IP address and ASN pool information. This is not true, because a rack does not store any information on the IP address and ASN pools. IP address and ASN pools are resources that are created in the data center design phase, and they contain a range of IP addresses and ASNs that can be assigned to the devices and the virtual networks. A rack only uses the IP address and ASN pools to assign the resources to the devices in the rack build².

It stores device port data rates and vendor information. This is not true, because a rack does not store any information on the device port data rates and vendor information. The device port data rates and vendor information are specified in the rack build, which assigns the specific vendor and model of the devices to the rack. A rack only uses the rack build to apply the network configuration and services to the devices in the rack deployment³.

Racks (Staged)

Rack Types (Datacenter Design)

Rack Builds (Staged)

Racks (Active)

To Get Premium Files for JN0-480 Visit

<https://www.p2pexams.com/products/jn0-480>

For More Free Questions Visit

<https://www.p2pexams.com/juniper/pdf/jn0-480>

