



QUESTION & ANSWER

HIGHER QUALITY, BETTER SERVICE

Provide One Year Free Update!

<https://www.passquestion.com>

Exam : **JN0-480**

Title : **Data Center, Specialist
(JNCIS-DC)**

Version : **DEMO**

1. When an agent installation is successful, devices are placed into which state using the Juniper Apstra UI?

- A. IS-MAINT
- B. OOS-READY
- C. OOS-QUARANTINED
- D. IS-ACTIVE

Answer: B


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 changes its state to Out of Service Ready (OOS-READY)¹².

References:

- Managing Devices
- AOS Device Configuration Lifecycle

2. Exhibit.



You are working to build an ESI-LAG for a multihomed server. The ESI-LAG is not coming up as multihomed.

Referring to the exhibit, what are two solutions to this problem? (Choose two.)

- A. The gateway IP addresses on both devices must be different.
- B. The LACP system ID on both devices must be the same.
- C. The loopback IP addresses on both devices must be the same.
- D. The ESI ID on both devices must be the same.

Answer: B D

Explanation:

According to the Juniper documentation¹, an ESI-LAG is a link aggregation group (LAG) that spans two or more devices and is identified by an Ethernet segment identifier (ESI). An ESI-LAG provides redundancy and load balancing for a multihomed server in an EVPN-VXLAN network.

To configure an ESI-LAG, you need to ensure that the following requirements are met:

- The LACP system ID on both devices must be the same. This ensures that the LACP protocol can negotiate the LAG parameters and form a single logical interface for the server.
- The ESI ID on both devices must be the same. This ensures that the EVPN control plane can advertise the ESI-LAG as a single Ethernet segment and synchronize the MAC and IP addresses of the server

across the devices.

- The VLAN ID and VNI on both devices must be the same. This ensures that the server can communicate with other hosts in the same virtual network and that the VXLAN encapsulation and decapsulation can work properly.

In the exhibit, the LACP system ID and the ESI ID on both devices are different, which prevents the ESI-LAG from coming up as multihomed. Therefore, the correct answer is B and D. The LACP system ID on both devices must be the same and the ESI ID on both devices must be the same.

References: ESI-LAG Made Easier with EZ-LAG, Example: Configuring an ESI on a Logical Interface With EVPN-MPLS Multihoming, Introduction to EVPN LAG Multihoming

3.Which statement is correct about the Juniper Apstra Rendered configuration?

- A. It is built at commit time and stored in a MySQL database.
- B. It is stored in a NoSQL database and incrementally updated.
- C. It is dynamically tendered at commit time.
- D. It is rendered from the graph database and stored locally.

Answer: D

Explanation:

The Juniper Apstra Rendered configuration is the configuration that is generated from the staged blueprint and applied to the devices in the network. The Rendered configuration is dynamically rendered at commit time, which means that it is created on the fly based on the latest changes and validations in the blueprint. The Rendered configuration is not stored in any database, but it can be viewed in the Apstra UI or downloaded as a file. The Rendered configuration reflects the desired state of the network as defined by the intent of the blueprint.

The other options are incorrect because:

- A. It is built at commit time and stored in a MySQL database is wrong because the Rendered configuration is not stored in any database, let alone a MySQL database. Apstra uses a graph database to store the network topology and configuration data, not a relational database like MySQL.
- B. It is stored in a NoSQL database and incrementally updated is wrong because the Rendered configuration is not stored in any database, let alone a NoSQL database. Apstra uses a graph database to store the network topology and configuration data, not a non-relational database like NoSQL. The Rendered configuration is not incrementally updated, but dynamically rendered at commit time.
- D. It is rendered from the graph database and stored locally is wrong because the Rendered configuration is not rendered from the graph database, but from the staged blueprint. The graph database stores the network topology and configuration data, but the Rendered configuration is generated from the blueprint, which is a logical representation of the network design and intent. The Rendered configuration is not stored locally, but it can be downloaded as a file if needed.

References:

- Config Rendering in Juniper Apstra
- AOS Device Configuration Lifecycle
- Configlets (Datacenter Design)

4.Which three statements describe intent-based analytics? (Choose three.)

- A. It indicates when device operating versions require updating.
- B. It is a real-time information processing pipeline.

- C. It is used to establish network performance baselines.
- D. It alerts the network operator when network performance moves away from the baseline.
- E. It collects information from vendor websites.

Answer: B C D

Explanation:

Intent-based analytics (IBA) is a feature of Juniper Apstra that allows you to combine intent from the network design with current and historic data from devices to reason about the network at-large¹.

IBA has the following characteristics:

- It is a real-time information processing pipeline. This means that IBA can ingest, process, and analyze large amounts of data from devices in real time, using agents and probes. Agents are software components that collect data from devices and send them to the Apstra server. Probes are user-defined queries that aggregate data across devices and generate advanced data that can be more easily reasoned about¹.
- It is used to establish network performance baselines. This means that IBA can use the advanced data to measure and monitor the network performance against the expected outcomes and service levels. IBA can also use the historic data to create baselines that represent the normal behavior and state of the network².
- It alerts the network operator when network performance moves away from the baseline. This means that IBA can detect and report any anomalies or deviations from the baseline or the intent in the network. IBA can also provide insights and recommendations for troubleshooting and resolving the issues².

The following two statements are incorrect in this scenario:

- It indicates when device operating versions require updating. This is not true, because IBA does not provide any information or guidance about the device operating versions or updates. IBA is focused on the network performance and compliance, not on the device maintenance or upgrade¹.
- It collects information from vendor websites. This is not true, because IBA does not collect any information from vendor websites or external sources. IBA only collects information from the devices in the network, using agents and probes¹.

References:

- Intent-Based Analytics — Apstra 3.3.0 documentation
- What is Intent Based Networking? | Juniper Networks US

5.Exhibit.

Filter selected by all selected only unselected only

<input type="checkbox"/>	Name ↕	Rack Type ↕
<input type="checkbox"/>	borderleaf_001	BorderLeaf 2022-03-02 09:33
<input type="checkbox"/>	serverrack_001	ServerRack 2022-03-01 10:19
<input type="checkbox"/>	serverrack_002	ServerRack 2022-03-01 10:19
<input type="checkbox"/>	serverrack_003	ServerRack 2022-03-01 10:19
<input type="checkbox"/>	serverrack_004	ServerRack 2022-03-01 10:19
<input type="checkbox"/>	serverrack_005	ServerRack 2022-03-01 10:19
<input type="checkbox"/>	serverrack_006	ServerRack 2022-03-01 10:19

Referring to the exhibit, how many tack types ate used in the staged blueprint?

- A. six
- B. three
- C. seven
- D. two

Answer: D

Explanation:

Referring to the exhibit, the image shows the Racks table under the Staged menu in the Juniper Apstra UI. The Racks table displays the details of the racks that are used in the blueprint, such as the name, rack type, and date. The rack type is a resource that defines the type and number of leaf devices, access switches, and/or generic systems that are used in rack builds¹. The image shows seven racks in the table, but only two rack types: BorderLeaf and ServerRack. Therefore, the statement D is correct in this scenario.

The following three statements are incorrect in this scenario:

- A. six. This is not true, because there are not six rack types in the table, but only two. The number six corresponds to the number of racks that have the same rack type: ServerRack.
- B. three. This is not true, because there are not three rack types in the table, but only two. The number three does not correspond to any relevant information in the table or the image.
- C. seven. This is not true, because there are not seven rack types in the table, but only two. The number seven corresponds to the total number of racks in the table, not the rack types.

References:

- Rack Types (Datacenter Design)
- Racks (Staged)