



Free Questions for JN0-480 by certsinside

Shared by Goodwin on 22-07-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

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?

Options:

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.

[Rack Types \(Datacenter Design\)](#)

[Racks \(Staged\)](#)

Question 2

Question Type: MultipleChoice

What is the purpose of using a routing zone inside Juniper Apstra software?

Options:

A- A routing zone is used to enable L4-L7 inspection inside the fabric.

B- A routing zone is defined to secure the routing protocols.

- C- A routing zone defined at the Apstra manager level requires firewalls to be deployed.
- D- A routing zone is used to enable the communication between two VNIs within a VRF.

Answer:

D

Explanation:

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. Therefore, the correct answer is D. A routing zone is used to enable the communication between two VNIs within a VRF. A routing zone is not used for L4-L7 inspection, securing routing protocols, or requiring firewalls. Those are not the purposes of a routing zone in Juniper Apstra software. Reference: Routing Zones

Question 3

Question Type: MultipleChoice

You are working with a three-stage IP fabric using EBGW for peering.

In this scenario, which two actions are required to implement ECMP? (Choose two.)

Options:

- A- Use a load balancing policy applied to the forwarding table as an export policy.
- B- Use a load balancing policy applied to BGP as an export policy.
- C- Use the multipath multiple-as BGP parameter.
- D- Use a load balancing policy applied to BGP as an import policy.

Answer:

B, C

Explanation:

To implement ECMP in IP fabric using EBGW, you need to enable BGP to install multiple equal-cost paths in the routing table and to advertise them to the peers. The following actions are required to achieve this:

B) Use a load balancing policy applied to BGP as an export policy. This is true because you need to apply a load balancing policy to BGP as an export policy to allow BGP to advertise multiple paths to the same destination to the peers. By default, BGP only advertises

the best path to the peers, which prevents ECMP. A load balancing policy can be configured to match the desired routes and set the multipath attribute to true. This will enable BGP to advertise up to the maximum number of paths configured by the maximum-paths command. For example, the following configuration applies a load balancing policy to BGP as an export policy for the neighbor 10.10.10.1:

```
policy-statement load-balance { term 1 { from { route-filter 192.168.0.0/16 exact; } then { multipath; accept; } } } protocols { bgp { group ebgp { type external; neighbor 10.10.10.1 { export load-balance; } } } }
```

C) Use the multipath multiple-as BGP parameter. This is true because you need to enable the multipath multiple-as BGP parameter to allow BGP to install multiple paths from different autonomous systems in the routing table. By default, BGP only installs multiple paths from the same autonomous system, which limits ECMP. The multipath multiple-as parameter can be configured under the BGP group or neighbor level. This will enable BGP to install up to the maximum number of paths configured by the maximum-paths command. For example, the following configuration enables the multipath multiple-as parameter for the BGP group ebgp:

```
protocols { bgp { group ebgp { type external; multipath multiple-as; } } }
```

The following options are incorrect because:

A) Use a load balancing policy applied to the forwarding table as an export policy is wrong because applying a load balancing policy to the forwarding table does not affect the BGP advertisement or installation of multiple paths. A load balancing policy applied to the forwarding table only affects how the traffic is distributed among the multiple paths in the forwarding table. It does not enable ECMP in BGP.

D) Use a load balancing policy applied to BGP as an import policy is wrong because applying a load balancing policy to BGP as an import policy does not affect the BGP advertisement of multiple paths. A load balancing policy applied to BGP as an import policy only affects how the BGP routes are accepted or rejected from the peers. It does not enable ECMP in BGP. Reference:

IP Fabric Underlay Network Design and Implementation

Use ECMP to distribute traffic between two paths, one learned by eBGP and one learned by iBGP on a Cisco NX-OS switch

Example: Configure an EVPN-VXLAN Centrally-Routed Bridging Fabric Using EBGP

Question 4

Question Type: MultipleChoice

Which two statements are correct about repairing a Juniper Apstra cabling map before deploying your blueprint? (Choose two.)

Options:

- A- You must manually change the cabling map to update spine-to-leaf fabric links.
- B- Apstra can use LLDP data from the spine-to-leaf fabric devices to update the connections in the cabling map.
- C- Apstra can use LLDP data from the leaf devices to update the leaf-to-generic connections in the cabling map.
- D- You must manually change the cabling map to update leaf-to-generic links.

Answer:

B, C

Explanation:

The cabling map is a graphical representation of the physical connections between the devices in the data center fabric. It shows the status of the cables, interfaces, and BGP sessions for each device. You can use the cabling map to verify and repair the cabling before deploying your blueprint. Based on the web search results, we can infer the following statements:

Apstra can use LLDP data from the spine-to-leaf fabric devices to update the connections in the cabling map. This is true because Apstra can collect LLDP data from the devices using the Generic Graph Collector processor and use it to update the cabling map automatically. LLDP is a protocol that allows devices to exchange information about their identity, capabilities, and neighbors¹².

Apstra can use LLDP data from the leaf devices to update the leaf-to-generic connections in the cabling map. This is true because Apstra can also collect LLDP data from the leaf devices and use it to update the connections to the generic devices, such as routers, firewalls, or servers. Generic devices are devices that are not managed by Apstra but are part of the data center fabric²³.

You must manually change the cabling map to update spine-to-leaf fabric links. This is false because Apstra can use LLDP data to update the spine-to-leaf fabric links automatically, as explained above. However, you can also manually change the cabling map to override the Apstra-generated cabling, if needed²⁴.

You must manually change the cabling map to update leaf-to-generic links. This is false because Apstra can use LLDP data to update the leaf-to-generic links automatically, as explained above. However, you can also manually change the cabling map to override the Apstra-generated cabling, if needed²⁴. Reference:

[LLDP Overview](#)

[Edit Cabling Map \(Datacenter\)](#)

[Generic Devices](#)

[Import / Export Cabling Map \(Datacenter\)](#)

Question 5

Question Type: MultipleChoice

You have recently committed a change after creating a new blueprint in Juniper Apstr

a. In the main dashboard, you see a number of anomalies related to BGR What is a likely cause of these anomalies?

Options:

A- You have misconfigured ASNs.

B- The fabric has not converged yet.

- C- Spine-leaf links are incorrectly set.
- D- A generic system has not been configured.

Answer:

B

Explanation:

In Juniper Apstra, a blueprint is a logical representation of the network design and configuration. When you create a new blueprint, you need to commit the changes to apply them to the network devices. However, committing the changes does not mean that the network is immediately updated and operational. It may take some time for the network to converge and reflect the new state of the blueprint. During this time, you may see some anomalies related to BGP in the main dashboard, which indicate that the BGP sessions are not established or stable between the devices. These anomalies are usually temporary and will disappear once the network converges and the BGP sessions are up and running. Therefore, the statement B is the most likely cause of these anomalies in this scenario.

The following three statements are less likely causes of these anomalies in this scenario:

You have misconfigured ASNs. This is possible, but not very likely, because Juniper Apstra provides ASN pools that can be automatically assigned to the devices based on their roles. You can also manually specify the ASNs for the devices, but you need to ensure that they are unique and consistent with the network design. If you have misconfigured ASNs, you may see some anomalies related to BGP, but they will not disappear after the network converges. You will need to fix the ASNs and commit the changes again to resolve the anomalies.

Spine-leaf links are incorrectly set. This is possible, but not very likely, because Juniper Apstra provides connectivity templates that can be used to define the spine-leaf links based on the interface maps. You can also manually specify the spine-leaf links, but you need to ensure that they are correct and match the physical cabling. If you have incorrectly set the spine-leaf links, you may see some anomalies related to BGP, but they will not disappear after the network converges. You will need to fix the spine-leaf links and commit the changes again to resolve the anomalies.

A generic system has not been configured. This is not relevant, because a generic system is a device that is not managed by Juniper Apstra, but is connected to the network. A generic system does not affect the BGP sessions between the devices that are managed by Juniper Apstra. If you have a generic system in your network, you need to configure it manually and ensure that it is compatible with the network design. A generic system does not cause any anomalies related to BGP in the main dashboard.

[Blueprint Summaries and Dashboard](#)

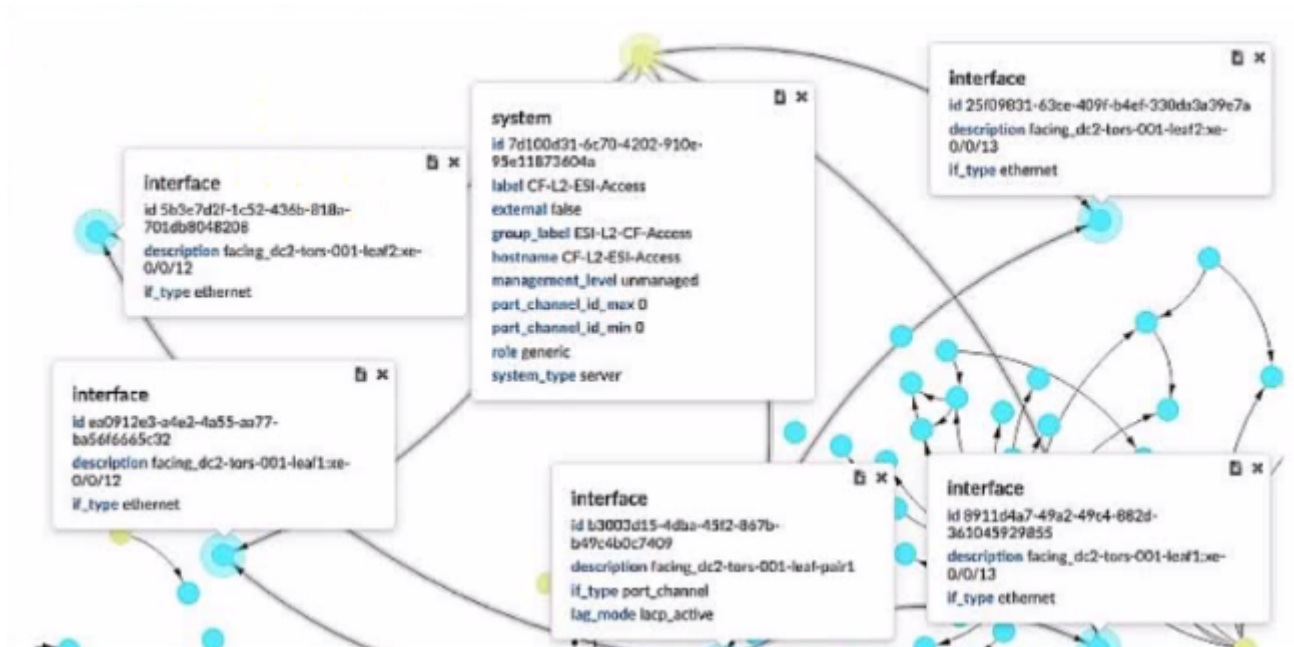
[BGP Session Flapping Probe](#)

[Probe: BGP Session Monitoring](#)

Question 6

Question Type: MultipleChoice

Exhibit.



Which two statements are correct about the information shown in the exhibit? (Choose two.)

Options:

- A- The system is fully managed by Juniper Apstra.
- B- The device shown is a generic system.
- C- Four physical interfaces exist in a LAG facing the leaf pair.
- D- The physical ports are not part of the LAG.

Answer:

B, C

Explanation:

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. Reference: Generic Systems (Datacenter Design), Form LAG | Apstra 4.1 | Juniper Networks

Question 7

Question Type: MultipleChoice

Exhibit.

Create Connectivity Template

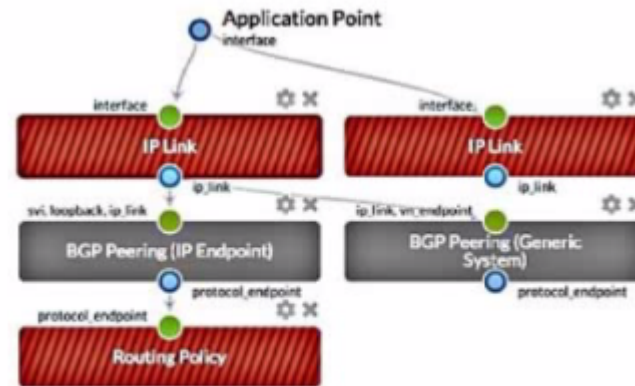
Parameters Primitives User-defined **Pre-defined**

Quick Search

[BGP Dynamic over L3 connectivity](#)

[BGP over L2 connectivity](#)

[BGP over L3 connectivity](#)



Referring to the exhibit, which statement is correct?

Options:

- A- The gray-solid primitives indicate further configuration is required.
- B- The gray-solid primitives indicate that they are incompatible with the connectivity template design.
- C- The red-striped primitives indicate that they are incompatible with the connectivity template design.
- D- The red-striped primitives indicate that further configuration is required.

Answer:

D

Explanation:

A connectivity template is a set of configuration parameters that can be applied to a device or a group of devices in a blueprint. A blueprint is a logical representation of the network design and intent. A primitive is a basic unit of configuration that can be added to a connectivity template. A primitive can be a link, a peering, a policy, or a service. In the exhibit, the red-striped primitives indicate that further configuration is required for them to be compatible with the connectivity template design. The red stripes mean that the primitive is incomplete or invalid, and it needs to be edited or deleted. For example, the IP Link primitive needs to have the interface name and IP address specified for each end of the link. The other options are incorrect because:

A) The gray-solid primitives indicate further configuration is required is wrong because the gray-solid primitives indicate that they are compatible with the connectivity template design. The gray color means that the primitive is valid and complete, and it does not need any further configuration.

B) The gray-solid primitives indicate that they are incompatible with the connectivity template design is wrong because the gray-solid primitives indicate that they are compatible with the connectivity template design, as explained above.

C) The red-striped primitives indicate that they are incompatible with the connectivity template design is wrong because the red-striped primitives indicate that further configuration is required, not that they are incompatible. The red stripes mean that the primitive is incomplete or invalid, but it can be fixed by editing or deleting it. Reference:

[Connectivity Templates](#)

[Data Center Automation Using Juniper Apstra](#)

[Config Rendering in Juniper Apstra](#)

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>

