Vendor: Juniper

Exam Code: JN0-480

Exam Name: Data Center, Specialist (JNCIS-DC)

Version: DEMO

QUESTION 1
You want to add a configuration that is not supported by Juniper Apstra reference architecture using a configlet.
Which two configurations would be applicable in this scenario? (Choose two.)

A. static route configuration
B. policy configuration
C. syslog configuration
D. NTP configuration

Answer: CD
Explanation:
A configlet is a configuration template that augments Apstra's reference design with non-native device configuration. They consist of one or more generators. Each generator specifies a NOS type (config style), when to render the configuration, and CLI commands (and file name as applicable).
Some applications for configlets include the following:
- Syslog
- SNMP access policy
- TACACS / RADIUS
- Management ACLs
- Control plane policing
- NTP
- Username / password

QUESTION 2
Which statement about Juniper Apstra role-based access control is correct?

A. The viewer role is predefined and can be deleted.
B. The administrator role can see all permissions.
C. The user role can create roles.
D. The administrator role is the only predefined role.

Answer: B
Explanation:
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 enterprise. Juniper Apstra has four predefined user roles: administrator, device_ztp, user, and viewer. 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 modified. Therefore, the statement that the administrator role can see all permissions is correct.

QUESTION 3
You are using Juniper Apstra to design a data center fabric. In this scenario, which object type associates a specific vendor model to a logical device?

A. templates
B. device profiles
C. agent profiles
D. interface map

Answer: B
Explanation:
Device profiles are objects that associate a specific vendor model to a logical device in Juniper Apstra. Device profiles contain extensive hardware model details, such as form factor, ASIC, CPU, RAM, ECMP limit, and supported features. Device profiles also define how configuration is generated, how telemetry commands are rendered, and how configuration is deployed on a device. Device profiles enable the Apstra system to render and deploy the configuration according to the Apstra Reference Design.

QUESTION 4
Which two statements are correct about probes? (Choose two.)

A. Default probes can be cloned, modified, and saved.
B. Only the variable parameters tor default probes can be edited and saved.
C. All default probes are enabled for all blueprints.
D. Default probes are enabled, based on the intent for a blueprint.

Answer: AD
Explanation:
Probes are the basic unit of abstraction in Intent-Based Analytics (IBA). They are used to collect, process, and analyze data from the network and raise anomalies based on specified conditions. Probes are composed of processors and stages that form a directed acyclic graph (DAG) of data flow.

QUESTION 5
In the Juniper Apstra Ul. which two resource types would be created in the Resources menu? (Choose two.)

A. bridge domain identifier (BDI)
B. DHCP pools
C. ASN pools
D. IP pools

Answer: CD
Explanation:
The Resources menu in the Juniper Apstra UI allows you to create and manage various types of resources that are assigned to different elements of the network.
Resources include the following types:

```
- IPv4 (including Host IPv4)
- IPv6 (including Host IPv6)
- ASN (autonomous system number)
- VNI (virtual network identifier)
- VLAN (virtual local area network)
```

Integer (used for pool type VLAN in local pools in Freeform blueprints) Therefore, the correct answer is C and D. ASN pools and IP pools are two types of resources that can be created in the Resources menu. Bridge domain identifier (BDI) and DHCP pools are not applicable in this scenario, because they are not part of the resources types supported by Juniper Apstra.

QUESTION 6

In the Juniper Apstra UI. you are creating a VNI pool for virtual networks.
In this scenario, which VNI range is acceptable?

A. Any range is acceptable for the VNI pool.
B. The valid VNI range is 4096 through 16777214.
C. The valid VNI range is 2 through 4096.
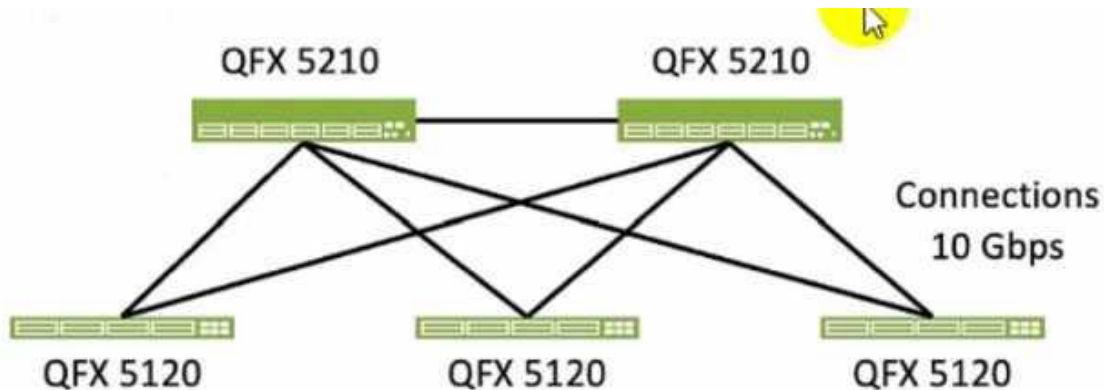D. The valid VNI range is 1 through 10000.

Answer: B
Explanation:
In the Juniper Apstra UI, you can create VNI pools for virtual networks that use VXLAN
encapsulation in the overlay network. A VNI pool is a resource pool that contains a range of VNIs
that can be assigned to the virtual networks. The valid VNI range for a VNI pool is 4096 through
16777214, according to the VXLAN standard. Therefore, the statement B is correct in this
scenario. The following three statements are incorrect in this scenario:
Any range is acceptable for the VNI pool. This is not true, because the VNI range has a lower and
upper limit defined by the VXLAN standard. The lower limit is 4096, and the upper limit is
16777214. Any VNI outside this range is invalid and cannot be used for VXLAN encapsulation.
The valid VNI range is 2 through 4096. This is not true, because the VNI range does not start
from 2, but from 4096. The VNIs from 2 to 4095 are reserved and cannot be used for VXLAN
encapsulation. The valid VNI range is 1 through 10000. This is not true, because the VNI range
does not include 1, which is also reserved and cannot be used for VXLAN encapsulation. The
VNI range also does not end at 10000, but at 16777214, which is the maximum possible value for
a 24-bit VNI field.

QUESTION 7
Referring to the exhibit, what needs to change in the IP fabric to make it a valid IP fabric?

A. The IP fabric must consist of only one device model throughout the fabric.
B. The connection between the two spine nodes must be increased to 40 Gbps.
C. The IP fabric connections must be increased to a speed greater than 10 Gbps.
D. The connection between the two spine nodes must be removed.

Answer: D
Explanation:
To make the IP fabric a valid IP fabric, the connection between the two spine nodes must be
removed. This is because an IP fabric is a network topology that uses a spine-leaf architecture,
where the spine devices are only connected to the leaf devices, and the leaf devices are only

connected to the spine devices. This creates a non-blocking, high-performance, and scalable network that supports Layer 3 routing protocols such as BGP or OSPF. The connection between the two spine nodes in the exhibit violates the spine-leaf design principle and introduces unnecessary complexity and potential loops in the network.

QUESTION 8

You have recently committed a change after creating a new blueprint in Juniper Apstra. In the main dashboard, you see a number of anomalies related to BGR What is a likely cause of these anomalies?

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.

QUESTION 9

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

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-lea! 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: BC
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 neighbors. 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 fabric.
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 needed.

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 needed.

QUESTION 10
You are working with a three-stage IP fabric using EBGP for peering. In this scenario, which two actions are required to implement ECMP? (Choose two.)

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 lo BGP as an import policy.

Answer: BC
Explanation:
To implement ECMP in IP fabric using EBGP, you need to enable BGP to install multiple equal-cost paths in the routing table and to advertise them to the peers.

QUESTION 11
What is the purpose of using a routing zone inside Juniper Apstra software?

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:
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.