



**Vendor:** Juniper

**Exam Code:** JN0-480

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

**Version:** 24.031

### 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 for 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 UI, 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, which statement is correct?



- 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.

#### 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 BGP. 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-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: 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 to 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.

#### QUESTION 12

Which protocol is used to advertise EVPN routes?

- A. OSPF
- B. BGP
- C. IS-IS
- D. RIP

**Answer:** B

**Explanation:**

BGP is the protocol used to advertise EVPN routes. EVPN routes are a new type of BGP network layer reachability information (NLRI) that carry MAC address and IP prefix information for Ethernet

VPNs. EVPN routes are exchanged between PEs using BGP multiprotocol extensions (MP-BGP) over MPLS, VXLAN, SR, or SRv6 tunnels. EVPN routes enable PEs to learn the reachability of MAC addresses and IP prefixes of different sites within the same EVPN instance. EVPN routes also support various features such as fast convergence, redundancy, aliasing, and inter-subnet routing.

**QUESTION 13**

You use Juniper Apstra to enable a new VXLAN virtual network. Which two components would be automatically derived in this situation? (Choose two.)

- A. IP subnet
- B. VLAN-ID
- C. VXLAN VNI
- D. Route Zone

**Answer:** AC

**Explanation:**

A VXLAN virtual network is a collection of Layer 2 forwarding domains that span multiple racks in a fabric. A VXLAN virtual network requires a name and a VXLAN network identifier (VNI), which is a 24-bit number that identifies the virtual network. The VNI can be either explicitly assigned or auto-assigned from a resource pool. A VXLAN virtual network can also have Layer 3 connectivity, which enables routing between different VNIs within a routing zone. A routing zone is an L3 domain that isolates the IP traffic of different tenants. A routing zone can have one or more VNIs associated with it. To enable Layer 3 connectivity, a VXLAN virtual network needs an IP subnet, which is a range of IP addresses that can be assigned to the hosts in the virtual network. The IP subnet can be either explicitly assigned or auto-assigned from a resource pool. Therefore, the correct answer is A and C. IP subnet and VXLAN VNI are two components that would be automatically derived when enabling a new VXLAN virtual network using Juniper Apstra.

**QUESTION 14**

What is the purpose of a Juniper Apstra rack?

- 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 roles.

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 configuration.



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 rack. 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 rack.

#### **QUESTION 15**

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?

- 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).

#### **QUESTION 16**

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

- 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.

#### **QUESTION 17**

Using Juniper Apstra, which component is defined in a template?



- 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:**

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.

**QUESTION 18**

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

- A. In this scenario, which statement is correct?
- B. When creating the widget, select the Add to Blueprint Dashboard option.
- C. On the blueprint dashboard, click on the Add Widget option.
- D. Widgets automatically appear on the blueprint dashboard.
- E. 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 preference. 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 network.

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 widget.

**QUESTION 19**

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

- A. authorized
- B. root

- C. viewer
- D. user

**Answer:** CD

**Explanation:**

Juniper Apstra provides four system-defined user roles that are available in the Apstra GUI environment. They are: administrator, device\_ztp, viewer, and user. 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 element.

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 element.

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 RADIUS.

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 user.

**QUESTION 20**

You have a virtual network that needs controlled access to other virtual networks in the same routing zone. Using the Juniper Apstra UI, which feature would be used to accomplish this task?

- A. interface policy
- B. anti-affinity policy
- C. routing policy
- D. security policy

**Answer:** D

**Explanation:**

A security policy is the feature that would be used to accomplish the task of controlling access to other virtual networks in the same routing zone using the Juniper Apstra UI. A security policy allows you to define rules that specify which traffic is allowed or denied between different virtual networks, IP endpoints, or routing zones. A security policy can be applied to one or more virtual networks in the same routing zone, and it can use various criteria to match the traffic, such as source and destination IP addresses, protocols, ports, or tags. A security policy can also support DHCP relay, which enables the forwarding of DHCP requests from one virtual network to another.

**QUESTION 21**

What is the purpose of an interface map in Juniper Apstra?

- A. An interface map associates a logical device with a device profile.
- B. An interface map specifies a connection between the interfaces of two devices.
- C. An interface map specifies the number of ports and the port speeds of a logical device
- D. An interface map specifies the connections between racks in a template.

**Answer:** B


**Explanation:**

An interface map is a configuration template that maps interfaces between logical devices and

physical hardware devices (represented with device profiles) while adhering to vendor specifications. An interface map specifies a connection between the interfaces of two devices, such as a leaf and a spine, a leaf and a server, or a leaf and an external gateway. An interface map can also specify port transformations, such as breaking out a 40 GbE port into four 10 GbE ports, or disabling unused ports. An interface map can be used to achieve the intended network configuration rendering and to enable features such as LAG, ESI-LAG, or MLAG.

**QUESTION 22**

Referring to the exhibit, how many broadcast domains will an Ethernet frame pass through when traversing the IP fabric from Server A to Server B?



- A. 1
- B. 4
- C. 2
- D. 3

**Answer: C**

**Explanation:**

Referring to the exhibit, the image shows a simplified diagram of an IP fabric network connecting two servers, labeled as Server A and Server B. The IP fabric is a network architecture that uses a Clos topology to provide high bandwidth, low latency, and scalability for data center networks. The IP fabric consists of spine and leaf devices that use BGP as the routing protocol and VXLAN as the overlay technology.

A broadcast domain is a logical portion of a network where any device can directly transmit broadcast frames to other devices at the data link layer (OSI Layer 2). A broadcast frame is a frame that has a destination MAC address of all ones (FF:FF:FF:FF:FF:FF), which means that it is intended for all devices in the same broadcast domain. A broadcast domain is usually bounded by a router, which does not forward broadcast frames to other networks.

In the exhibit, there are two broadcast domains that an Ethernet frame will pass through when

traversing the IP fabric from Server A to Server B. The first broadcast domain is the one that contains Server A and the leaf device that it is connected to. The second broadcast domain is the one that contains Server B and the leaf device that it is connected to. The IP fabric itself is not a broadcast domain, because it uses IP routing and VXLAN encapsulation to transport the Ethernet frames over the Layer 3 network.

### QUESTION 23

Which two actions are required during Juniper Apstra's deploy phase? (Choose two.)

- A. Assign device profiles to the blueprint.
- B. Assign user roles to the blueprint.
- C. Assign interlace maps to the blueprint.
- D. Assign resources to the blueprint.

**Answer:** AD

#### **Explanation:**


The deploy phase is the final step in the Juniper Apstra data center fabric design and deployment process. In this phase, you apply the Apstra-rendered configuration to the devices and verify the intent of the blueprint. Based on the web search results, we can infer the following actions are required during the deploy phase:

Assign device profiles to the blueprint. This action associates a specific vendor model to each logical device in the blueprint. 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. Assign resources to the blueprint. This action allocates the physical devices, IP addresses, VLANs, and ASNs to the logical devices, networks, and routing zones in the blueprint. Resources can be assigned manually or automatically by the Apstra system. Assigning resources ensures that the blueprint has all the necessary elements to generate the configuration and deploy the fabric5 . Assign user roles to the blueprint. This action is not required during the deploy phase. User roles are defined at the system level, not at the blueprint level. User roles determine the permissions and access levels of different users in the Apstra system. User roles can be system-defined or custom-defined .

Assign interface maps to the blueprint. This action is not required during the deploy phase. Interface maps are defined at the design phase, not at the deploy phase. Interface maps are objects that map the logical interfaces of a logical device to the physical interfaces of a device profile. Interface maps enable the Apstra system to generate the correct interface configuration for each device in the fabric.

### QUESTION 24

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



- A. leaf
- B. spine
- C. super spine
- D. server

**Answer: B**

**Explanation:**

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.

**QUESTION 25**

A member of your organization made changes to a predefined interface map using Juniper Apstra. Which two statements are correct in this scenario? (Choose two.)

- A. Changes to interface maps in the global catalog do not affect interface maps that have already been imported into blueprint catalogs
- B. Any changes made to predefined interface maps are discarded when Apstra is upgraded.
- C. Changes made to predefined interface maps will not have an impact on the Apstra software.
- D. Changes to interface maps in the global catalog will raise anomalies that may need to be addressed at the next commit.

**Answer: AB**

**Explanation:**

An interface map is a configuration template that maps interfaces between logical devices and physical hardware devices (represented with device profiles) while adhering to vendor specifications. An interface map can be either predefined or custom. A predefined interface map is one that ships with Apstra software and supports most qualified Juniper devices. A custom interface map is one that is created by the user to meet specific requirements. An interface map can be stored in either the global catalog or the blueprint catalog. The global catalog contains all the interface maps that are available for use in any blueprint. The blueprint catalog contains the interface maps that are imported from the global catalog and used in a specific blueprint.

**QUESTION 26**

Which statement is true when onboarding a Juniper Networks device using a Juniper Apstra ZTP server?

- A. The Device Key to be used can be set In the dhcpd.conf file on the ZTP server.
- B. The State can be set In the ztp.Json file on the ZTP server.
- C. The Management IP address cannot be predetermined.
- D. The Hostname will be the serial-number of the device.

**Answer: B**

**Explanation:**

The ztp.Json file on the Apstra ZTP server contains the configuration parameters for each device that is onboarded using ZTP. One of the parameters is the State, which can be one of the following values: init, ready, in\_progress, done, error, or disabled. The State indicates the current status of the device in the ZTP process. For example, if the State is ready, it means that the device is ready to be onboarded by the Apstra ZTP server. If the State is done, it means that the device has completed the ZTP process and is managed by the Apstra server. The State can be manually set or changed in the ztp.Json file to control the behavior of the device during ZTP.

**QUESTION 27**

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?

- A. Connectivity Templates
- B. Virtual Networks table
- C. Managed Devices list
- D. Links table

**Answer: D**

**Explanation:**

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.

**QUESTION 28**

When editing a device configuration to install some manual changes, which procedure should be followed?

- A. Edit the configuration on the device directly by the CLI; the changes will automatically be adjusted in the Juniper Apstra configuration
- B. Edit the pristine configuration of the device.
- C. Add a persistent change to a device configuration with a configlet.
- D. Delete the device from the Juniper Apstra system, change the configuration, then re-import the device.

**Answer: C**

**Explanation:**

A configlet is a small piece of configuration that can be applied to a device or a group of devices to make persistent changes that are not overwritten by Apstra. Configlets can be used to install manual changes that are not part of the Apstra rendered configuration, such as custom commands, scripts, or features. Configlets can be created, edited, and deleted from the Apstra GUI or CLI.

**QUESTION 29**

In Juniper Apstra, which statement is correct?

- A. VMware anomaly detection is on by default.
- B. VMware anomaly detection requires a vCenter server configured under External Systems
- C. VMware anomaly detection requires a VMware hypervisor with exports enabled.
- D. VMware anomaly detection requires an Apstra server running on VMware.


**Answer: B**

**Explanation:**

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.

**QUESTION 30**

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.)



- A. EVPN signaled route Type-8 VXLAN tunnels
- B. EVPN signaled route Type-3 VXLAN tunnels
- C. EVPN signaled route Type-6 VXLAN tunnels
- D. EVPN signaled route Type-2 VXLAN tunnels



**Answer:** BD

**Explanation:**

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.

**QUESTION 31**

In the Juniper Apstra design phase, which object dictates port count, port speed, and how the ports would be used?

- A. logical devices
- B. rack type
- C. network devices
- D. interface map

**Answer:** D

**Explanation:**

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.

**QUESTION 32**

Juniper Apstra has indicated an anomaly with respect to cabling. What are two ways to remediate the issue? (Choose two.)

- A. Manually edit the cabling map.
- B. Redeploy the errant device.
- C. Set the invalid ports to a disabled state.
- D. Have Apstra autoremediate the cabling map using LLDP.

**Answer:** AD

**Explanation:**

A cabling anomaly is an issue that occurs when the physical connections between the devices in the data center fabric do not match the expected connections based on the Apstra Reference Design.

A cabling anomaly can cause problems such as incorrect routing, suboptimal traffic flow, or device isolation.

To remediate the issue, you can use one or both of the following methods:

Manually edit the cabling map. This allows you to override the Apstra-generated cabling and specify the correct connections between the devices.

You can use the Apstra UI or the Apstra CLI to edit the cabling map and apply the changes to the fabric.

Have Apstra autoremediate the cabling map using LLDP. This allows Apstra to collect LLDP data from the devices 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 the LLDP data to detect and correct any cabling errors in the fabric.

**QUESTION 33**

You are installing a Juniper Apstra server in your data center. You have multiple users that will be expected to configure, manage, and carry out operational tasks in your data center. You have decided to implement remote user authentication for the role-based access control of your Apstra server. In this scenario, which three methods are supported? (Choose three.)

- A. TACACS+
- B. LDAP
- C. RADIUS
- D. SAML
- E. Auth0

**Answer:** ABC

**Explanation:**

To implement remote user authentication for the role-based access control of your Apstra server, you can use one of the following methods: TACACS+, LDAP, or RADIUS. These are the protocols that Juniper Apstra supports to authenticate and authorize users based on roles assigned to individual users within an enterprise. You can configure the Apstra server to use one or more of these protocols as the authentication sources and specify the order of preference. You can also configure the Apstra server to use local user accounts as a fallback option if the remote authentication fails.

**QUESTION 34**

Which statement is correct about making changes to a predefined device profile in Juniper Apstra?

- A. The changes you make to a predefined device profile will be discarded and overwritten when upgrading the Apstra server version.
- B. The processor load on the Apstra server will be negatively impacted when changes are made to a predefined device profile.
- C. Changing the predefined device profile will also affect deployed devices that use the same profile.
- D. You must reboot the Apstra server after changes are made to a predefined device profile before they will become active.

**Answer:** A

**Explanation:**

A predefined device profile is a configuration template that is shipped with Apstra software and supports most qualified Juniper devices. A predefined device profile cannot be changed, since any changes would be discarded and overwritten when you upgrade the Apstra server version. If you want to customize a predefined device profile, you can clone and edit it instead.

**QUESTION 35**

Which two statements are correct regarding a pristine configuration in Juniper Apstra? (Choose two.)

- A. It is the configuration file placed on the device when decommissioning the device.
- B. It is the device's currently active configuration.
- C. It is the configuration file on a device before acknowledgment in Apstra.
- D. It is the device's previously active configuration.

**Answer:** AC

**Explanation:**

A pristine configuration in Juniper Apstra is the configuration file that is used to onboard a device into the Apstra software application. A pristine configuration contains the minimum settings that are required for the device to communicate with the Apstra server, such as the hostname, management IP address, username, password, and SSH key. A pristine configuration has the following characteristics:

It is the configuration file placed on the device when decommissioning the device. This is because when a device is decommissioned from the Apstra software application, it is reverted back to its pristine configuration, which removes all the network configuration and services that were applied by the Apstra software application. This allows the device to be reused or repurposed for another network.

It is the configuration file on a device before acknowledgment in Apstra. This is because when a device is onboarded into the Apstra software application, it is initially in the discovery state, which means that the device is discovered by the Apstra server, but not yet acknowledged by the user. In the discovery state, the device has the pristine configuration, which can be viewed and edited by the user. Once the user acknowledges the device, the device moves to the deployed state, which means that the device is ready to receive the network configuration and services from the Apstra software application.

**QUESTION 36**

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: C**

**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).

**QUESTION 37**

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 rendered at commit time.
- D. It is rendered from the graph database and stored locally.

**Answer: C**

**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.

**QUESTION 38**

Multitenancy for applications is achieved by creating virtual networks (VNs) within which construct?

- A. security policy
- B. routing table
- C. connectivity template
- D. routing zone

**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.

**QUESTION 39**

Which two statements about VXLAN VNIs are correct? (Choose two.)

- A. VNIs can have over 16 million unique values.
- B. VNIs identify a collision domain.
- C. VNIs identify a broadcast domain
- D. VNIs are alphanumeric values.

**Answer: AC**

**Explanation:**

VXLAN VNIs are virtual network identifiers that are used to identify and isolate Layer 2 segments in the overlay network. VXLAN VNIs have the following characteristics:

VNIs can have over 16 million unique values. This is because VXLAN VNIs are 24-bit fields that can range from 4096 to 16777214, according to the VXLAN standard. This allows VXLAN to support a large number of Layer 2 segments and tenants in the network. VNIs identify a broadcast domain. This is because VXLAN VNIs are used to group the end hosts that belong to the same Layer 2 segment and can communicate with each other using VXLAN tunnels. The VXLAN tunnels are established using the VTEP information that is distributed by EVPN. The VTEPs are VXLAN tunnel endpoints that perform the VXLAN encapsulation and decapsulation. The VXLAN tunnels preserve the Layer 2 semantics and support the broadcast, unknown unicast, and multicast traffic within the same VNI.

**QUESTION 40**

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?

- A. Acknowledge the device.
- B. Take the device out of maintenance mode.
- C. Install the agent for the device.
- D. Take the device out of drain state.


**Answer: A**

**Explanation:**

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 deployed.

**QUESTION 41**

Which two statements about ESI values are correct for the server connections to the fabric shown in the exhibit? (Choose two.)



- A. A valid ESI value for Server A is 0x00.00.00.00.00.00.00.00.00.
- B. A valid ESI value for Server B is 0x00.20.20.20.20.20.20.20.20.
- C. A valid ESI value for Server A is 0x00.10.10.10.10.10.10.10.10.
- D. A valid ESI value for Server B is 0x00.00.00.00.00.00.00.00.00.

**Answer:** CD

**Explanation:**

To answer this question, we need to understand the concept of ESI values in EVPN LAGs. An ESI is a 10-byte value that identifies an Ethernet segment, which is a set of links that connect a multihomed device (such as a server) to one or more PE devices (such as leaf switches) in an EVPN network. The same ESI value must be configured on all the PE devices that connect to the same Ethernet segment. This allows the PE devices to form an EVPN LAG, which supports active-active or active-standby multihoming for the device. The ESI value can be manually configured (type 0) or automatically derived from LACP (type 1) or other methods. In the exhibit, Server A is connected to two leaf switches (QFX 5210) using a LAG with LACP enabled. Server B is connected to three leaf switches (QFX 5120) using a LAG with LACP enabled.

**QUESTION 42**

What is the function of the Revert button in the Juniper Apstra UI?

- A. The Revert button will rollback to the previous configuration of a specified device.
- B. The Revert button deletes any uncommitted changes within Apstra.
- C. The Revert button performs a rollback 0 on Juniper devices in the fabric.
- D. The Revert button will undo the last change made to the configuration.

**Answer: B**

**Explanation:**

The Revert button is located on the Uncommitted tab of the blueprint page. The Uncommitted tab shows the changes that have been staged but not yet committed to the network. The Revert button allows you to discard any uncommitted changes and revert to the last committed state of the blueprint. This is useful if you want to cancel the changes that you have made or if you want to start over with a fresh slate.

**QUESTION 43**

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: BCD**

**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-large. 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 about.

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 network.

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 issues.

**QUESTION 44**

When working with logical devices, you specify where each port group is connected. In this scenario, which two Juniper Apstra UI options are available to the operator? (Choose two.)

- A. router
- B. unused
- C. generic
- D. firewall

**Answer:** BC

**Explanation:**

When working with logical devices, you specify where each port group is connected by selecting the port group layout and the port speed and role (s) for each port group. The Juniper Apstra UI offers two options to the operator for the port group role: unused and generic.

Unused: This option means that the port group is not configured or used by Apstra. This can be useful for ports that are faulty, reserved, or not part of the data center fabric.

Generic: This option means that the port group is configured with a generic role that is not specific to any device type or function. This can be useful for ports that are used for testing, troubleshooting, or custom purposes.

#### **QUESTION 45**

You must configure a static route for traffic to exit a configured routing zone. In the Juniper Apstra UI, where would you accomplish this task?

- A. under Active -> Virtual -> Routing Zones
- B. under Staged -> Virtual -> Routing Zones
- C. under Active -> Connectivity Templates
- D. under Staged -> Connectivity Templates

**Answer:** D

**Explanation:**

To configure a static route for traffic to exit a configured routing zone, you need to use the Connectivity Templates feature in the Juniper Apstra UI. A Connectivity Template is a set of configuration parameters that can be applied to a device or a group of devices in a blueprint. You can use Connectivity Templates to configure static routes, BGP, OSPF, and other network services. To create a Connectivity Template, you need to go to the Staged tab and select Connectivity Templates from the left menu. Then, you can click on the + icon to create a new template. You can specify the name, description, and scope of the template. The scope determines which devices or device groups the template will be applied to. You can also specify the order of the template, which determines the priority of the template when multiple templates are applied to the same device. After creating the template, you can add configuration items to the template. To add a static route, you need to select Static Route from the drop-down menu and enter the destination network, subnet mask, and next-hop IP address. You can also specify the administrative distance and the track object for the static route. After adding the configuration items, you need to save the template and commit the changes to the blueprint.

#### **QUESTION 46**

Which fabric type should be chosen in a template to create a five-stage Clos?

- A. collapsed
- B. circuit switched
- C. rack-based
- D. pod-based

**Answer:** D

**Explanation:**

A five-stage Clos architecture allows for large-scale topologies with an additional aggregation layer that interconnects multiple pods into a single fabric. A pod is a group of racks that share the same spine devices. A rack is a group of leaf devices that connect to the same servers. To create a five-stage Clos network using Juniper Apstra, you need to choose the pod-based fabric type in the template creation wizard. This will allow you to specify the number of pods, planes, spines, and



leaves for your network design.

**QUESTION 47**

Referring to the exhibit, how do you display the IPv6 subnets for all of the listed VXLANs?

| Name                        | Routing Zone | Type  | VN ID | Assigned to | DHCP Service | IPv4 Connectivity | IPv4 Subnet  | Actions |
|-----------------------------|--------------|-------|-------|-------------|--------------|-------------------|--------------|---------|
| vlan_30_leaf3_v4            | default      | VLAN  | 30    | 1 nodes     | Enabled      | Enabled           | 10.1.3.0/24  |         |
| red_vxlan_42_v4_one_ep_mlag | red          | VXLAN | 30011 | 2 nodes     | Enabled      | Enabled           | 10.1.15.0/24 |         |
| red_vxlan_41_v4_one_ep      | red          | VXLAN | 30010 | 2 nodes     | Enabled      | Enabled           | 10.1.14.0/24 |         |

- A. IPv6 subnets are shown when each VXLAN is selected individually.
- B. Select Columns, then select IPv6 Subnet.
- C. Select all VXLANs, and the IPv6 Subnets column will appear
- D. An IPv6 Subnets column is not shown, indicating that no VXLAN has an assigned IPv6 subnet

**Answer: B**

**Explanation:**

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).

**QUESTION 48**

You want to apply a configlet to a specific device using Juniper Apstra. Which two parameters would be used to accomplish this task? (Choose two.)

- A. form factor
- B. hostname
- C. port group
- D. tags

**Answer: BD**

**Explanation:**

To apply a configlet to a specific device using Juniper Apstra, you need to specify the device's

hostname and tags. The hostname is the unique identifier of the device in the Apstra system, and the tags are the labels that you can assign to the device to group it with other devices that share the same characteristics. You can use the hostname and tags to filter the devices that you want to apply the configlet to in the blueprint catalog.

**QUESTION 49**

What does EVPN use to identify which remote leaf device advertised the EVPN route?

- A. a route distinguisher value
- B. a community tag
- C. a route target value
- D. a VRF target value

**Answer: A**


**Explanation:**

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.

**QUESTION 50**

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: BD**

**Explanation:**

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.

#### **QUESTION 51**

In the case of IP Clos data center five-stage fabric design, what are two roles of the super spines? (Choose two.)

- A. Super spines are used to interconnect two different data center pods.
- B. Super spines connect to all spine devices within the five-stage architecture.
- C. Super spines are used to connect leaf nodes within a data center pod.
- D. Super spines are always connected to an external data center gateway.

**Answer:** AB

**Explanation:**

In the case of IP Clos data center five-stage fabric design, the super spines are the devices that provide the highest level of aggregation in the network. They have two main roles:

Super spines are used to interconnect two different data center pods. A pod is a cluster of leaf and spine devices that form a 3-stage Clos topology. A 5-stage Clos topology consists of multiple pods that are connected by the super spines. This allows for scaling the network to support more devices and bandwidth.

Super spines connect to all spine devices within the five-stage architecture. The spine devices are the devices that provide the second level of aggregation in the network. They connect to the leaf devices, which are the devices that provide access to the end hosts. The super spines connect to all the spine devices in the network, regardless of which pod they belong to. This provides any-to-any connectivity between the pods and enables optimal routing and load balancing.

#### **QUESTION 52**

IBA probes analyze telemetry data from specified devices within a blueprint. Which component identifies devices that supply data for a specific probe?

- A. data selector
- B. processor
- C. search engine
- D. graph query

**Answer:** D

**Explanation:**

A graph query is a component that identifies devices that supply data for a specific probe. A graph query is an expression that matches nodes in the Apstra graph database based on their attributes, such as device name, role, type, or tag. A graph query can be used to select the source devices for the input processors of a probe, as well as to filter the data by device attributes in the subsequent processors of a probe.

#### **QUESTION 53**

You want to keep virtual networks isolated from each other within the Juniper Apstra system. In this scenario, what are three ways to accomplish this task? (Choose three.)

- A. Disable IPv4 connectivity when creating the virtual network within the same Routing Zone.
- B. Enable Security Policy for virtual networks in the same Routing Zone.
- C. Disable Route Target exports when creating the Routing Zones.
- D. Use Connectivity Templates to block access within the same Routing Zone.
- E. Put each network in different Routing Zones.

**Answer:** BDE

**Explanation:**

To keep virtual networks isolated from each other within the Juniper Apstra system, you can use one or more of the following methods:

Enable Security Policy for virtual networks in the same Routing Zone. This allows you to define rules that control the traffic flow between different virtual networks within the same routing zone. You can specify the source and destination virtual networks, the protocol, the port, and the action (allow or deny) for each rule. The security policy is applied on the ingress interface of the leaf devices.

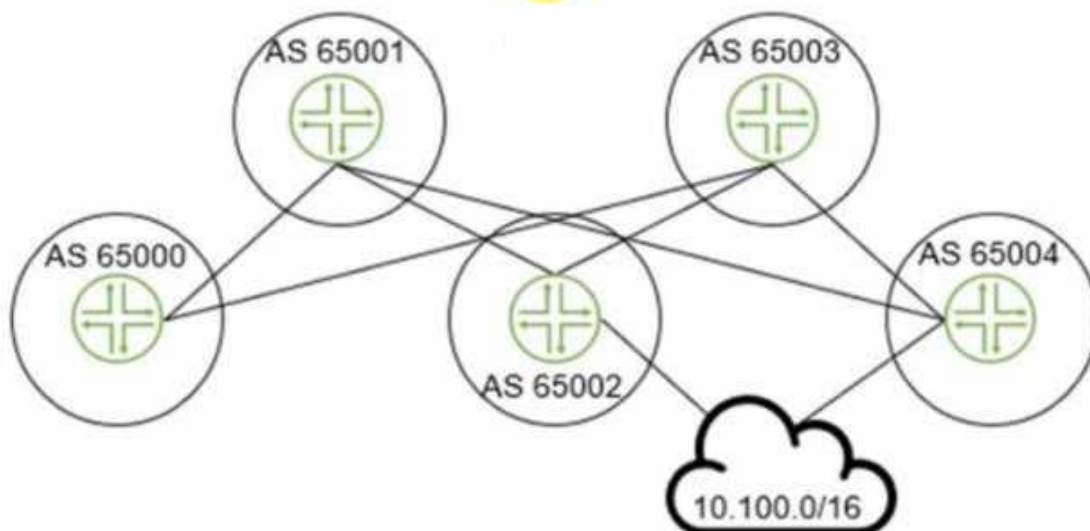
Use Connectivity Templates to block access within the same Routing Zone. This allows you to customize the connectivity between different racks within the same routing zone. You can create templates that define the link type, the routing protocol, and the access control list (ACL) for each rack pair. The ACL can be used to filter the traffic based on the source and destination IP addresses, the protocol, and the port.

Put each network in different Routing Zones. This allows you to create logical boundaries between different virtual networks based on the route target (RT) values. A routing zone is a collection of virtual networks that share the same RT for importing and exporting routes. Virtual networks in different routing zones do not exchange routes with each other, unless you configure remote EVPN gateways to connect them.

**QUESTION 54**

The 10.100.0.0/16 route is being advertised into your BGP IP fabric. ECMP load balancing has been properly enabled on all devices.

In this scenario, how many routes will the leaf device in AS 65000 receive for the 10.100.0.0/16 prefix?



- A. 3
- B. 1
- C. 2
- D. 4

**Answer:** A

**Explanation:**

The leaf device in AS 65000 will receive three routes for the 10.100.0.0/16 prefix, one from each spine device in AS 65001, AS 65002, and AS 65003. Since ECMP load balancing is enabled, the leaf device will install all three routes in its routing table and distribute the traffic among them.

#### **QUESTION 55**

Using the Juniper Apstra multitenancy capabilities, which approach will allow a tenant to interconnect two different routing zones?

- A. Interconnection is the default behavior.
- B. Use interconnection through the fabric spine nodes.
- C. Interconnection cannot be enabled.
- D. Use interconnection through an external gateway.


**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. 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.

#### **QUESTION 56**

In the EVPN-VXLAN data center fabric bridged overlay architecture shown in the exhibit, the servers are connected to Lead and Leaf6 using the same virtual network identifier (VNI). Which two statements are correct in this scenario? (Choose two.)



- A. The underlay must use IRB interfaces.
- B. The underlay must be provisioned with PIMv2.
- C. Loopback IPv4 addresses must be advertised into the EBGp underlay from leaf and spine devices.
- D. The underlay EBGp peering's must be established between leaf and spine devices.

**Answer:** CD

**Explanation:**

In the EVPN-VXLAN data center fabric bridged overlay architecture shown in the exhibit, the servers are connected to Leaf1 and Leaf6 using the same virtual network identifier (VNI). This means that the servers belong to the same Layer 2 domain and can communicate with each other using VXLAN tunnels across the fabric. The underlay network provides the IP connectivity between the leaf and spine devices, and it uses EBGp as the routing protocol. Therefore, the following two statements are correct in this scenario:

Loopback IPv4 addresses must be advertised into the EBGp underlay from leaf and spine devices. This is because the loopback addresses are used as the source and destination IP addresses for the VXLAN tunnels, and they must be reachable by all the devices in the fabric. The loopback addresses are also used as the router IDs and the BGP peer addresses for the EBGp sessions. The underlay EBGp peering's must be established between leaf and spine devices. This is because the EBGp sessions are used to exchange the underlay routing information and the EVPN routes for the overlay network. The EBGp sessions are established using the loopback addresses of the devices, and they follow a spine-and-leaf topology, where each leaf device peers with all the spine devices, and each spine device peers with all the leaf devices.

**QUESTION 57**

You are receiving cable, interface, and BGP anomalies from several devices within the data center fabric. In Juniper Apstra, how would you troubleshoot these types of errors?

- A. In the UI, go to Time Voyager and revert to the last working version.
- B. In the UI, access the console to the devices and review the interface states.
- C. In the UI, go to Devices and confirm that agent connectivity is fine.
- D. In the UI, verify device connectivity by consulting the cable map.

**Answer:** D

**Explanation:**

The cable 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 cable map to identify and troubleshoot any cable, interface, or BGP anomalies that may occur in the fabric. You can also filter the cable map by device name, device type, device role, device state, cable state, interface state, or BGP state.

**QUESTION 58**

Referring to the exhibit, what is the minimum information you must add to create a new routing zone?

The screenshot shows a web interface for creating a routing zone. The title is "Create Routing Zone". There are three required text input fields: "VRF Name\*", "VLAN ID\*", and "VNI". Below these is a "Routing Policies" dropdown menu with "Select..." as the current value. Underneath is a "Route Target Policies" section with two sub-sections: "Import Route Targets" containing an "Add Import Route Target" button, and "Export Route Targets" containing an "Add Export Route Target" button. At the bottom right, there is a "Create Another?" checkbox and a "Create" button.

- A. VRF Name only
- B. VRF Name and Routing policies
- C. VRF Name, VLAN ID. And VNI
- D. VRF Name, VLAN ID, VNI, Routing Policies

**Answer: C**

**Explanation:**

To create a new routing zone, you must specify the VRF Name, VLAN ID, and VNI for the routing zone. These are the mandatory fields in the user interface shown in the exhibit. The VRF Name is the name of the L3 domain that isolates the IP traffic of the routing zone from other routing zones. The VLAN ID is the identifier for the VLAN tagged Layer 3 links on external connections. The VNI is the VxLAN Network Identifier associated with the routing zone. The Routing Policies are optional fields that allow you to configure import and export route targets for the routing zone. These are only applicable for EVPN routing zones, which use MP-EBGP as the overlay control protocol.

**QUESTION 59**

Which two statements are correct about Time Voyager? {Choose two.}

- A. Time Voyager retains all of the blueprint revisions from the last Juniper Apstra backup.



- B. Time Voyager retains the five most recent blueprint commits.
- C. Time Voyager retains the last ten blueprint commits.
- D. Time Voyager retains up to twenty-five saved revisions.

**Answer:** BD

**Explanation:**

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.

**QUESTION 60**

In the Juniper Apstra UI, which three resources are assigned under the Resources menu? (Choose three.)

- A. VTEP pools
- B. ASN pools
- C. VNI pools
- D. logical device pools
- E. IP address pools

**Answer:** BCE

**Explanation:**

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.

**QUESTION 61**

In Juniper Apstra, which three modes are available for devices? (Choose three.)

- A. Deploy
- B. Active
- C. Stopped
- D. Drain
- E. Ready

**Answer:** ADE

**Explanation:**

Juniper Apstra supports three deploy modes for devices: Deploy, Drain, and Ready. These modes determine the configuration and state of the devices in the data center fabric.


**Deploy:** This mode applies the full Apstra-rendered configuration to the device, according to the Apstra Reference Design. The device state becomes IS-ACTIVE and the device is ready to carry traffic in the fabric.

**Drain:** This mode adds a "drain" configuration to the device, which prevents any new traffic from entering the device. The device state becomes IS-READY and the device is prepared for maintenance or decommissioning.

**Ready:** This mode removes the Apstra-rendered configuration from the device, leaving only the basic configuration such as device hostname, interface descriptions, and port speed/breakout. The device state becomes IS-READY and the device is not part of the fabric.

**QUESTION 62**

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 63**

Which two statements are correct about the graph query output shown in the exhibit? (Choose two.)

```
{
  "count": 1,
  "items": [
    {
      "interface": {
        "id": "0be7ea4e-d691-44ff-aec7-da4fb2836555",
        "type": "interface",
        "label": null,
        "description": "to.cf-l2-esi-access",
        "evpn_esi_mac": null,
        "if_name": "ae2",
        "if_type": "port_channel",
        "ipv4_addr": null,
        "ipv4_addr_type": null,
        "ipv4_enabled": null,
        "ipv6_addr": null,
        "ipv6_addr_type": null,
        "ipv6_enabled": null,
        "lag_mode": "lACP_active",
        "loopback_id": null,
        "mlag_id": null,
        "mode": null,
        "po_control_protocol": null,
        "port_channel_id": 2,
        "property_set": null,
        "protocols": null,
        "ref_count": null,
        "subintf_id": null,
        "tags": null,
        "vlan_id": null
      }
    }
  ]
}
```

- A. The interface has an IP address assigned to it.
- B. The output shows a LAG connection.
- C. The switch in the output is a Juniper device.
- D. The interface has tags assigned to it.

**Answer:** BC

**Explanation:**

The graph query output shown in the exhibit is a JSON representation of an interface node and its properties in the Apstra graph database. Based on the output, we can infer the following statements: The output shows a LAG connection. This is true because the interface node has a property called

lag\_mode which is set to lacp\_active, indicating that the interface is part of a link aggregation group (LAG) that uses the Link Aggregation Control Protocol (LACP) to negotiate the link state and parameters.

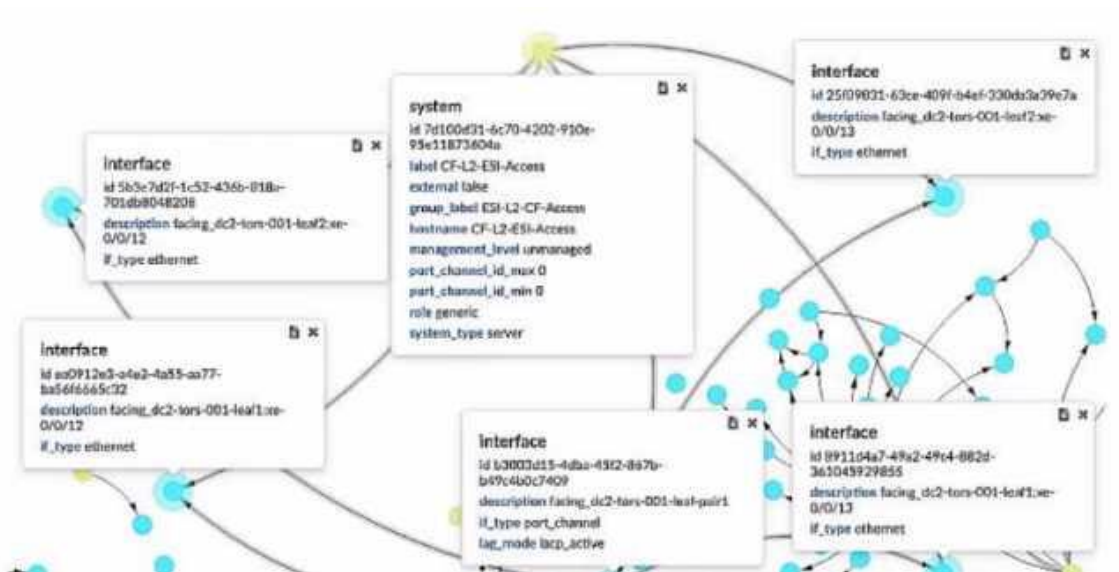
The switch in the output is a Juniper device. This is true because the interface node has a property called if\_name which is set to ae, indicating that the interface name follows the Juniper naming convention for aggregated Ethernet interfaces.

The interface has an IP address assigned to it. This is false because the interface node has properties called ipv4\_addr and ipv6\_addr which are both set to null, indicating that the interface does not have any IPv4 or IPv6 address configured.

The interface has tags assigned to it. This is false because the interface node has a property called tags which is set to null, indicating that the interface does not have any tags associated with it.

**QUESTION 64**

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



- 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: BC**

**Explanation:**

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. 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 if\_type. The LAG is facing the leaf pair, which are the two switches that connect to the device.

**QUESTION 65**

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

Filter selected by  all  selected only  unselected only

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

- 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 builds. The image shows seven racks in the table, but only two rack types: BorderLeaf and ServerRack.