PassQueen
a.k.a dumpsqueen.com

Implementing Cisco Application Centric Infrastructure (300-620 DCACI)

Cisco 300-620

Version Demo

Total Demo Questions: 35

Total Premium Questions: 350

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Architecture	14
Topic 2, Fabric Infrastructure	93
Topic 3, Fabric Endpoint Connectivity	134
Topic 4, Automation and Orchestration	11
Topic 5, Security and External Connectivity	98
Total	350

QUESTION NO: 1

Which type of profile needs to be created to deploy an access port policy group?

A. attachable entity

B. Pod

C. module

D. leaf interface

ANSWER: A

Explanation:

An attachable entity profile is required when deploying an access port policy group because it provides the association between the interface policy configuration and the relevant domain context. In Cisco ACI, an attachable access entity profile (AAEP) is linked to a physical domain and then associated with the access port policy group. This relationship tells APIC which domain and VLAN-pool resources are valid for endpoints connected through ports using that policy group. The AAEP therefore acts as the policy attachment point that makes the access policy usable for physical access interfaces.

After the attachable entity profile is associated with the access port policy group, the policy group can be selected through interface selectors in a leaf interface profile to apply the resulting policy to specific leaf-switch ports. This design separates reusable access-policy definitions from their eventual assignment to concrete interfaces, while preserving the domain and encapsulation scope required for endpoint deployment. Cisco documents AAEPs as the object used to associate domains with interface policy groups in its Cisco APIC Operating ACI guide. See also Cisco's APIC Access Policy Configuration Guide for the access-policy workflow.

QUESTION NO: 2

A Cisco ACI bridge domain and VRF are configured with a default data-plane learning configuration. Which two endpoint attributes are programmed in the leaf switch when receiving traffic? (Choose two.)

A. Remote MAC, IP

B. Remote Subnet

C. Local IP, not MAC

D. Local MAC, IP

E. Local Subnet

F. Remote IP

ANSWER: D F

Explanation:

With the default Cisco ACI data-plane learning behavior, a leaf switch directly connected to an endpoint learns and programs the endpoint as **Local MAC, IP**. The local MAC address identifies the endpoint's Layer 2 attachment and forwarding location, while the learned IP-to-MAC association supports integrated host tracking and routing functions in the bridge domain and VRF. This local endpoint information is advertised through the fabric control-plane mechanisms so that other leaves can forward traffic toward the correct endpoint location.

For an endpoint attached to another leaf, the receiving leaf programs the destination as a **Remote IP** entry. This allows the ingress leaf to route or bridge traffic for the remote IP toward the appropriate fabric destination, while endpoint location information is maintained and distributed by ACI's endpoint directory mechanisms. The combination of locally learned MAC/IP bindings and remotely programmed IP reachability supports distributed forwarding while retaining endpoint mobility and efficient fabric-wide endpoint resolution. Cisco documents endpoint learning and bridge-domain forwarding behavior in its ACI Layer 2 networking guide and describes VRF-based IP forwarding in the APIC configuration documentation: Cisco APIC Layer 2 Networking Configuration Guide and Cisco APIC VRF Configuration Guide.

QUESTION NO: 3

```
<fvTenant name="ACILab">
   <fvCtx name="pvn1"/>
   <fvBD name="bd1">
      <fvRsCtx tnFvCtxName="pvn1"/>
      <fvSubnet ip="10.1.100.1/24"/>
   </fvBD>
</fvTenant>
```

Refer to the exhibit. Which two objects are created as a result of the configuration? (Choose two.)

A. application profile

B. attachable AEP

C. bridge domain

D. endpoint group

E. VRF

ANSWER: C E

Explanation:

The configuration creates a **bridge domain** and a **VRF**. In the Cisco ACI object model, a VRF is represented by the `fvCtx` managed object. It provides the tenant-level Layer 3 routing and forwarding context, including route-table separation. When the configuration defines an `fvCtx` object, APIC creates the corresponding VRF in the specified tenant.

A bridge domain is represented by an `fvBD` managed object. A bridge domain defines the Layer 2 forwarding scope for endpoints and VLANs, and it can be associated with the VRF created by the `fvCtx` configuration. The resulting tenant configuration therefore contains both objects: the VRF for Layer 3 context and the bridge domain for Layer 2 forwarding behavior. These are independent ACI managed objects, even when the bridge domain references the VRF as part of its configuration. Cisco documents the ACI REST object model and the relevant managed-object classes in the Cisco APIC REST API Configuration Guide. Additional context on tenant networking constructs, including VRFs and bridge domains, is available in the Cisco APIC Basic Configuration Guide.

QUESTION NO: 4

Refer to the exhibit. An engineer configures the Cisco ACI fabric for VMM integration with ESXi servers that are to be connected to the ACI leaves. The server team requires the network switches to initiate the LACP negotiation as opposed to the servers. The LAG group consists of two 10 Gigabit Ethernet links. The server learn also wants to evenly distribute traffic across all available links. Which two enhanced LAG policies meet these requirements? (Choose two.)

A. LACP Mode: LACP Standby

B. LB Mode: Destination IP Address and TCP/UDP Port

C. LB Mode: Source and Destination MAC Address

D. LB Mode: Source IP Address and TCP/UDP Port

E. LACP Mode: LACP Active

ANSWER: D E

Explanation:

LACP Mode: LACP Active meets the requirement for the Cisco ACI leaf switches to initiate LACP negotiation. In active mode, the leaf sends LACP data units to begin and maintain negotiation with the connected ESXi host. This is the appropriate enhanced LAG policy setting when the infrastructure side, rather than the server side, must initiate the LACP exchange.

LB Mode: Source IP Address and TCP/UDP Port provides a granular hashing input for traffic sent over the two-member 10-Gigabit EtherChannel. By incorporating both the source IP address and the Layer 4 TCP or UDP port, separate flows from a host can be distributed across available member links more effectively than with a less granular hash. This supports the goal of using the available links evenly while preserving the normal port-channel behavior that keeps an individual flow on one physical member link. Cisco ACI enhanced LAG policies define both the LACP negotiation mode and the load-balancing algorithm used for the port channel. Cisco documentation describes LACP active operation and port-channel policy configuration in the Cisco APIC Virtual Edge Configuration Guide and the Cisco APIC Basic Configuration Guide.

QUESTION NO: 5

Refer to the exhibit.

An engineer wants to initiate an ICMP ping from Server1 to Server2. The requirement is for the BD1 to enforce ICMP replies that follow the expected path. The packets must be prevented from taking the direct path from Leaf1 to Server1. Which action must be taken on BD1 to meet these requirements?

A. Set L2 Unknown Unicast to Flood.

B. Set L2 Unknown Unicast to Hardware Proxy.

C. Disable Unicast Routing.

D. Enable ARP Flooding.

ANSWER: B

Explanation:

Set L2 Unknown Unicast to Hardware Proxy. In Cisco ACI, the hardware-proxy behavior uses the spine proxy and the fabric endpoint directory to resolve the appropriate destination leaf for a Layer 2 destination that is not locally known. The ingress leaf forwards the frame toward the proxy, which consults endpoint-location information and steers traffic to the leaf where the destination endpoint is attached. This provides controlled, endpoint-aware forwarding through the fabric rather than permitting the frame to be delivered through an unintended local or direct path. In this situation, applying Hardware Proxy to BD1 ensures that ICMP reply traffic requiring unknown-unicast handling is resolved by the ACI fabric's proxy mechanism and follows the intended fabric forwarding path to Server1. This behavior is particularly important while endpoint information is being learned, updated, or recovered, because the fabric can use its centralized endpoint mapping to maintain the desired forwarding decision. Cisco documents bridge-domain Layer 2 unknown-unicast behavior and hardware proxy operation in its ACI configuration guidance: Cisco APIC Basic Configuration Guide — Configuring Bridge Domains. Additional ACI forwarding and endpoint-directory concepts are described in the Cisco ACI Fundamentals.

QUESTION NO: 6

A network engineer must configure a Cisco ACI system to detect network loops for untagged and tagged traffic The loop must be detected and slopped by disabling an interface within 4 seconds Which configuration must be used?

A. Option A

B. Configure an ACI L2 Loop Protection policy with a detection period of no more than 4 seconds and the Port Disable action.

C. Option C

ANSWER: B

Explanation:

Configure an ACI L2 Loop Protection policy with the port-disable action and timer settings that provide a detection time of no more than four seconds. L2 Loop Protection is designed to identify Layer 2 loops on interfaces handling either tagged or untagged frames. It sends loop-detection packets and evaluates their return according to the configured detection interval and multiplier. When the calculated detection period is within the required four-second limit, the loop is identified quickly enough for the policy to take corrective action.

The required remediation is *Port Disable*. With this action, APIC administratively disables the affected interface after a loop is detected, stopping continued circulation of frames and limiting the impact of the loop. The policy must then be associated with the relevant interface policy group so that it is applied to the physical ports where external Layer 2 connectivity is present. Cisco documents ACI access-policy and interface-policy configuration in its APIC support documentation; Cisco's ACI L2 Loop Protection documentation search provides the corresponding release-specific configuration guides.

QUESTION NO: 7

A Solutions Architect is asked to design two data centers based on Cisco ACI technology that can extend L2/L3, VXLAN, and network policy across locations. ACI Multi-Pod has been selected. Which two requirements must be considered in this design? (Choose two.)

A. ACI underlay protocols, i.e. COOP, IS-IS and MP-BGP, spans across pods. Create QoS policies to make sure those protocols have higher priority.

B. A single APIC Cluster is required in a Multi-Pod design. It is important to place the APIC Controllers in different locations in order to maximize redundancy and reliability.

C. ACI Multi-Pod requires an IP Network supporting PIM-Bidir.

D. ACI Multi-Pod does not support Firewall Clusters across Pods. Firewall Clusters should always be local.

E. Multi-Pod requires multiple APIC Controller Clusters, one per pod. Make sure those clusters can communicate to each other through a highly available connection.

ANSWER: B C

Explanation:

A single APIC Cluster is required in a Multi-Pod design. It is important to place the APIC Controllers in different locations in order to maximize redundancy and reliability. ACI Multi-Pod is one logical ACI fabric composed of multiple pods, so it is managed through one APIC cluster rather than independent controller clusters per pod. Distributing APIC controllers across locations can improve resilience to a pod or site failure, provided the deployment meets Cisco's APIC-cluster connectivity, latency, and quorum-design guidance. This centralized policy and management model is what allows consistent tenant policy, endpoint groups, bridge domains, and VRFs to extend across the pods.

ACI Multi-Pod requires an IP Network supporting PIM-Bidir. The Inter-Pod Network transports the VXLAN-based traffic between pods and must provide the multicast capability used for handling broadcast, unknown-unicast, and multicast traffic across the Multi-Pod fabric. Cisco specifies bidirectional PIM in the IPN design, along with resilient routed connectivity and the required multicast reachability. These IPN capabilities are therefore foundational design requirements before extending L2/L3 connectivity and policy between data centers. See Cisco's ACI Multi-Pod Configuration Guide and the Cisco ACI Multi-Pod overview.

QUESTION NO: 8

An engineer is configuring ACI VMM domain integration with Cisco UCS-B Series. Which type of port channel policy must be configured in the vSwitch policy?

A. LACP Active

B. MAC Pinning

C. LACP Passive

D. MAC Pinning-Physical-NIC-load

ANSWER: B

Explanation:

MAC Pinning is the required port-channel policy for Cisco UCS B-Series VMM domain integration. In this design, the Cisco UCS virtual interface card presents multiple virtual network interface cards to the hypervisor, while the UCS fabric interconnect manages the underlying uplink connectivity. MAC pinning provides the forwarding behavior needed for this architecture: traffic associated with a source MAC address is pinned to a specific uplink, maintaining a stable forwarding path while allowing the vSwitch to use the available physical adapters. This configuration aligns the VMware vSwitch uplink behavior with the UCS B-Series fabric and adapter design used by ACI VMM integration. When creating the vSwitch policy that is associated with the VMM domain, the port-channel policy must therefore be set to MAC Pinning so that the hosts and their virtual interfaces are connected consistently through the UCS infrastructure. Cisco's UCS and ACI VMM integration guidance identifies MAC pinning as the vSwitch policy setting for this UCS B-Series deployment model. See Cisco's ACI and UCS VMM integration configuration guide and the Cisco UCS Manager configuration documentation.

QUESTION NO: 9

Which routing protocol is supported between Cisco ACI spines and IPNs in a Cisco ACI Multi-Pod environment?

A. OSPF

B. ISIS

C. BGP

D. EIGRP

ANSWER: C

Explanation:

BGP is supported between Cisco ACI spine switches and Inter-Pod Network (IPN) devices in a Cisco ACI Multi-Pod deployment. The IPN provides the routed Layer 3 transport that connects pods, while ACI uses BGP-based routing adjacencies across this infrastructure to exchange reachability information required for inter-pod communication. Each spine establishes external BGP (eBGP) peerings with the connected IPN routers; this design provides scalable route exchange and supports the required separation and resiliency of the multi-pod fabric. Cisco's Multi-Pod guidance specifies that the IPN must provide an IP transport network between pods and identifies BGP as the routing protocol used for the spine-to-IPN connections. The underlay transport must also meet Multi-Pod requirements such as end-to-end IP reachability, multicast support where applicable, and an MTU sufficient for encapsulated traffic. Therefore, BGP is the supported routing protocol for the routing adjacency between the ACI spines and IPN in this topology. See Cisco's ACI Multi-Pod Configuration Guide and the Cisco Application Centric Infrastructure documentation.

QUESTION NO: 10

A customer implements RBAC on a Cisco APIC using a Windows RADIUS server that is configured with network control policies. The APIC is as follows:

The customer requires User X to have access to TenantX only, without any extra privilege in the Cisco ACI fabric domain. Which Cisco AV pair must be implemented on the RADIUS server to meet these requirement?

A. shell:domains = TenantX-SD/fabric-admin/,common//read-all

B. shell:domains = TenantX-SD/tenant-admin

C. shell:domains = TenantX-SD/tenant-ext-admin/,common//read-all

D. shell:domains = TenantX-SD/tenant-admin/,common//read-all

ANSWER: D

Explanation:

`shell:domains = TenantX-SD/tenant-admin/,common//read-all` is the required RADIUS Cisco AV pair. Cisco APIC uses the `shell:domains` attribute returned during RADIUS authentication to assign the authenticated user to an APIC security domain and an RBAC role. Here, `TenantX-SD` identifies the tenant-specific security domain and `tenant-admin` grants administrative permissions scoped to that tenant. This provides User X the privileges needed to manage objects within TenantX while avoiding a fabric-level administrative role or access to other tenant security domains.

The `,common//read-all` portion supplies read-only access to the common tenant. APIC tenant-scoped users commonly require this limited visibility because shared objects, such as common infrastructure policy objects, can be referenced by tenant configuration. The empty role field before `read-all` indicates that no write-capable role is assigned in the common tenant; the user receives only the explicitly stated read-only privilege. This combination therefore preserves the required TenantX administration scope without extending fabric-domain privileges. Cisco documents the RADIUS `shell:domains` format and its security-domain/role mapping in the Cisco APIC Security Configuration Guide and the Cisco APIC Security Configuration Guide, Release 6.1(x).

QUESTION NO: 11

The engineer notices frequent MAC and IP address moves between different leaf switch ports. Which action prevents this problem from occurring?

A. Disable enforce subnet check.

B. Enable endpoint loop protection.

C. Enable rogue endpoint control.

D. Disable IP bridge domain enforcement.

ANSWER: B

Explanation:

Enable endpoint loop protection. In Cisco ACI, repeated endpoint moves—where the same MAC address or IP address is learned on different leaf ports in rapid succession—are commonly symptomatic of a Layer 2 loop, such as an improperly connected downstream switch or bridged endpoints. Endpoint loop protection monitors endpoint learning and detects excessive moves within its configured detection interval. When the configured threshold is exceeded, ACI identifies the condition as a loop event and takes protective action on the affected interface, preventing continued endpoint flapping and limiting the associated flooding, instability, and control-plane churn.

This feature is intended specifically for access-facing endpoint connectivity and is configured through endpoint security settings in the relevant ACI policy scope. It provides an automated safeguard against accidental physical loops and similar wiring faults without relying on manual intervention after every endpoint move event. Administrators should still investigate the alerted port and correct the cabling or downstream switching condition before restoring normal connectivity. Cisco documents endpoint loop protection as an endpoint-security mechanism for detecting frequent endpoint moves and mitigating loop behavior. See the Cisco APIC Layer 2 Configuration Guide and Cisco APIC Security Configuration Guide.

QUESTION NO: 12

An engineer plans a Cisco ACI firmware upgrade. The ACI fabric consists of three Cisco APIC controllers, two spine switches, and four leaf switches. Two leaf switches have 1-Gb copper s for bare metal servers, and the other two leaf switches have 10-Gb SFP ports to connect storage. Which set of actions accomplishes an upgrade with minimal disruptions?

A. Divide the switches into two upgrade groups: spines and leaves.

B. Divide the switches into three upgrade groups: spines, 1-Gb switches, and 10-Gb switches.

C. Divide the switches into two upgrade groups with one spine, one 1-Gb switch, and one 10-Gb switch per group.

D. Divide the switches into four upgrade groups with one switch per group.

ANSWER: C

Explanation:

Divide the switches into two upgrade groups with one spine, one 1-Gb switch, and one 10-Gb switch per group. This design keeps a forwarding path available while one member of each redundant layer is rebooted for its firmware installation. Each group contains only one of the two spine switches, so the other spine remains operational and continues to provide fabric connectivity during that group's upgrade. Likewise, separating the two 1-Gb access leaf switches and the two 10-Gb storage-facing leaf switches across the groups preserves service through the corresponding peer leaf while the first leaf is unavailable. This staging is particularly important when endpoints, storage, or external connectivity use redundant attachments across leaf switches. After the first group has completed successfully and its switches have rejoined the fabric in the target version, the second group can be upgraded using the same controlled process. Cisco ACI supports organizing switches into upgrade groups so that upgrades can be sequenced in a way that maintains redundancy and minimizes traffic impact. The APIC upgrade workflow and fabric-switch upgrade-group procedures are documented in Cisco's Cisco APIC Upgrade Guide. Cisco's APIC support documentation also provides release-specific upgrade guidance and compatibility information.

QUESTION NO: 13

An engineer must configure Policy-Based Redirect for a firewall that is connected in one-arm mode. Which two objects are required in the service graph configuration? (Choose two.)

A. Service graph template

B. Device selection policy

C. Bridge domain subnet with the Shared between VRFs scope

D. BGP route reflector policy

E. L3Out route control profile

ANSWER: A B

Explanation:

A service graph template defines the logical service topology and contains the function node that represents the firewall. A device selection policy is also required to associate the logical device with the actual firewall interfaces and endpoint attachments that can be used for redirection.

The redirect policy is associated with the service graph contract to identify the redirect destinations and provide flow-consistency behavior. A bridge-domain subnet or an L3Out route control profile is not a substitute for the service graph device-selection configuration. Cisco documents service graphs and policy-based redirect configuration in the Cisco APIC Layer 4 to Layer 7 Services Deployment Guide.

QUESTION NO: 14

In the context of ACI Multi-Site, when is the information of an endpoint (MAC/IP) that belongs to site 1 advertised to site 2 using the EVPN control plane?

A. Endpoint information is not exchanged across sites unless COOP protocol is used.

B. Endpoint information is not exchanged across sites unless a policy is configured to allow communication across sites.

C. Endpoint information is exchanged across sites as soon as the endpoint is discovered in one site.

D. Endpoint information is exchanged across sites when the endpoints are discovered in both sites.

ANSWER: B

Explanation:

Endpoint information is not exchanged across sites unless a policy is configured to allow communication across sites. In Cisco ACI Multi-Site, each site maintains its own local endpoint database, while the Multi-Site control plane uses EVPN to communicate required remote endpoint reachability between sites. Endpoint advertisements are policy driven: the system distributes the MAC/IP reachability information needed to support permitted intersite connectivity, rather than indiscriminately sharing every locally learned endpoint with every other site.

This behavior is an important scalability characteristic of the Multi-Site architecture. It limits the intersite control-plane state to endpoint information relevant to enabled communication, while still allowing the fabric to establish the appropriate VXLAN forwarding toward the remote site when policy permits the traffic. The policy requirement is therefore the condition that enables EVPN-based advertisement of endpoint information from site 1 to site 2. Cisco describes this policy-driven endpoint-learning and control-plane behavior in its ACI Multi-Site overlay design documentation: Cisco ACI Multi-Site Architecture White Paper. Additional implementation guidance is available in the Cisco ACI Multi-Site Configuration Guide.

QUESTION NO: 15

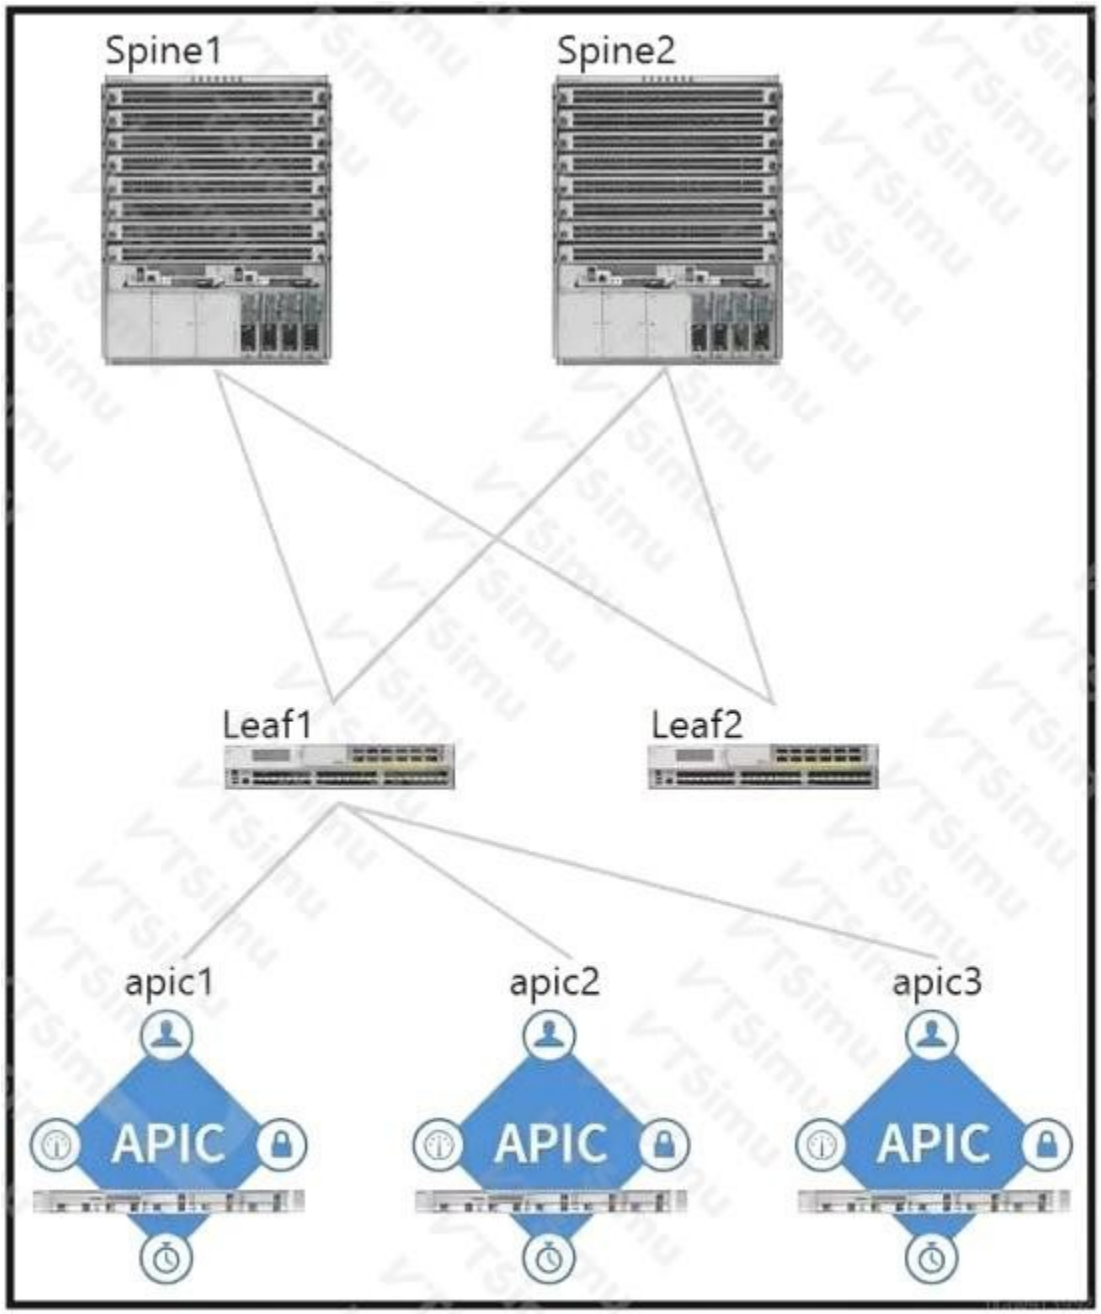

Refer to the exhibit. Which two components should be configured as route reflectors in the ACI fabric? (Choose two.)

A. Spine1

B. apic1

C. Spine2

D. Leaf1

E. Leaf2

F. apic2

ANSWER: A C

Explanation:

In a Cisco ACI fabric, the route-reflector function is placed on spine switches. Therefore, **Spine1** and **Spine2** should be configured as the route reflectors. ACI uses Multiprotocol BGP EVPN control-plane signaling to distribute endpoint reachability information across the fabric. The spines act as the scalable, centrally located BGP route-reflector nodes that

reflect EVPN route advertisements between leaf switches, eliminating the need for a full BGP mesh among all leaves. Selecting two spine nodes provides route-reflector redundancy and supports continued control-plane operation if one route reflector becomes unavailable. The APIC cluster centrally defines and deploys the route-reflector policy, including the selected spine-node IDs, but the BGP route-reflector service itself operates on the chosen spine switches. This placement aligns with the ACI leaf-and-spine architecture: leaves connect endpoints and enforce policy, while the spine layer provides the fabric's transit and control-plane route-reflection role. Cisco's ACI Layer 3 configuration documentation describes configuring route-reflector nodes from the spine-switch inventory, and Cisco's ACI architecture documentation explains the spine layer's role in scalable fabric forwarding and control-plane operation. See the Cisco APIC Layer 3 Configuration Guide and the Cisco ACI architecture white paper.

QUESTION NO: 16

The Application team reports that a previously existing port group has disappeared from vCenter. An engineer confirms that the VMM domain association for the EPG is no longer present. Which action determines which user is responsible for the change?

A. Check the EPG audit logs for the "deletion" action and compare the affected object and user.

B. Evaluate the potential faults that are raised for that EPG.

C. Examine the health score and drill down to an object that affects the EPG combined score.

D. Inspect the server logs to see who was logging in to the APIC during the last few hours.

ANSWER: A

Explanation:

Check the EPG audit logs for the "deletion" action and compare the affected object and user. Cisco ACI maintains audit records for configuration changes made through the APIC, including the identity of the authenticated user, the operation performed, the affected managed object, and the time of the change. Since the missing vCenter port group corresponds to the removed VMM domain association on the EPG, the relevant audit entry is the deletion/modification record for that EPG's VMM-domain relation. Reviewing that record identifies the username associated with the transaction and confirms precisely which object was changed. This is the appropriate evidence source because audit logging is intended to provide configuration-change traceability and accountability, rather than merely operational state information. The engineer can use the APIC GUI audit-log views or query the corresponding audit managed objects through the APIC API, then correlate the recorded affected object and timestamp with the reported disappearance of the port group. Cisco documents APIC audit records and their role in tracking administrative activity in its APIC audit-log documentation search. The APIC object model and API mechanisms used to inspect managed-object records are documented in the Cisco APIC Management Information Model Reference.

QUESTION NO: 17

Refer to the exhibit.

An engineer configures the Cisco ACI fabric for VMM integration with ESXi servers that are to be connected to the ACI leaves. The server team requires the network switches to initiate the LACP negotiation as opposed to the servers. The LAG group consists of two 10 Gigabit Ethernet links. The server learn also wants to evenly distribute traffic across all available links. Which two enhanced LAG policies meet these requirements? (Choose two.)

A. LACP Mode: LACP Standby

B. LB Mode: Destination IP Address and TCP/UDP Port

C. LB Mode: Source and Destination MAC Address

D. LB Mode: Source IP Address and TCP/UDP Port

E. LACP Mode: LACP Active

ANSWER: B E

Explanation:

LACP Mode: LACP Active meets the requirement for the ACI leaf switches to initiate link-aggregation negotiation. In active mode, the leaf sends LACP protocol packets to establish and maintain the port channel with the ESXi host. This is the appropriate enhanced LAG policy setting when the network side, rather than the server side, must initiate LACP.

LB Mode: Destination IP Address and TCP/UDP Port provides a flow-based hashing input with substantially more entropy than a MAC-only hash. The destination IP address and Layer 4 TCP or UDP port are considered when selecting a member link. With multiple application flows, this allows the two 10-Gigabit member links to be used effectively and distributes eligible flows across the available links. As with any Ethernet link aggregation, an individual flow remains on one physical member link to preserve packet ordering; “evenly” refers to distribution of multiple hashed flows, not splitting one flow across both links.

Cisco ACI enhanced LAG policies define both the LACP operational mode and the load-balancing algorithm used by leaf interfaces for VMM-attached hosts. See Cisco’s APIC Virtual Networking Configuration Guide and the APIC Layer 2 Configuration Guide.

QUESTION NO: 18

Refer to the exhibit.

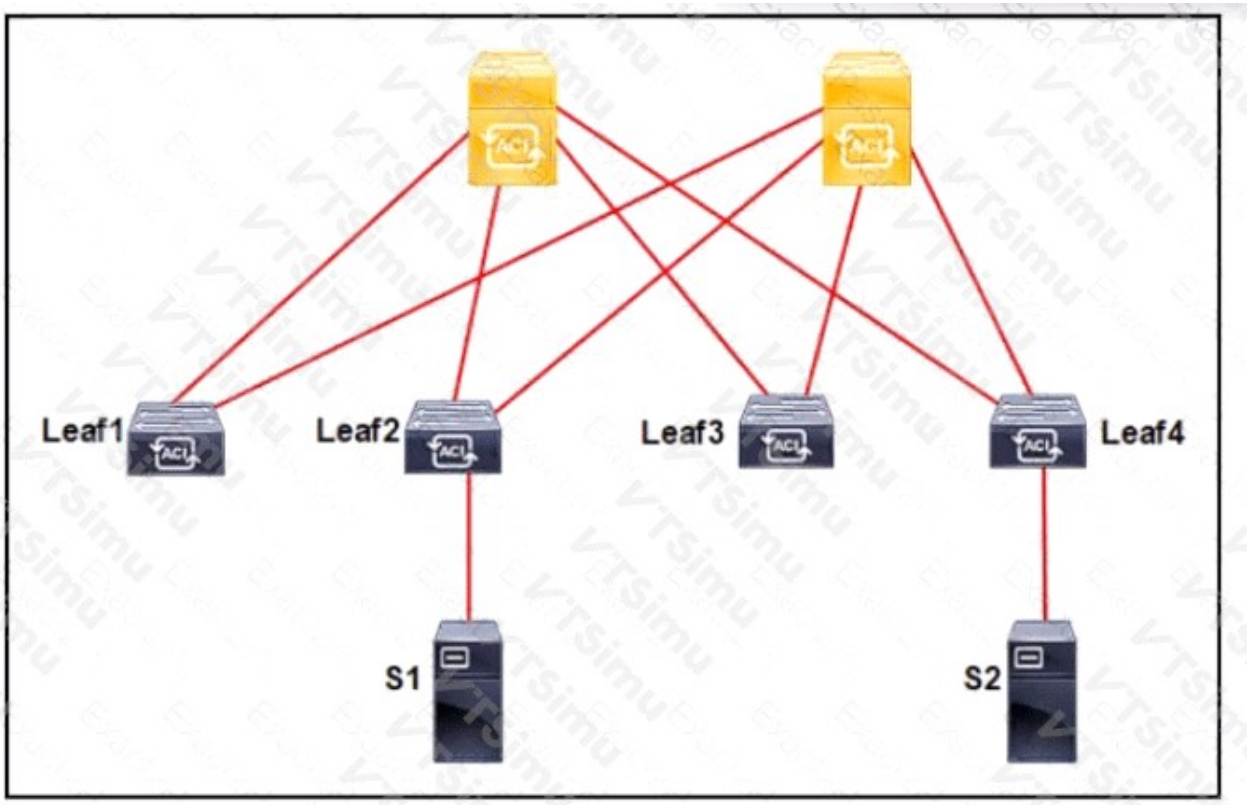

An application called App_1 is hosted on the server called S1. A silent host application. App_2. is hosted on S2. Both applications use the same VLAN encapsulation, which action forces Cisco ACI fabric to learn App_2 on ACI leaf 2?

A. Set Multi-Destination Flooding to Drop.

B. Set Unicast Routing to Hardware Proxy.

C. Set L2 Unknown Unicast to Flood.

D. Set L3 Unknown Multicast to Optimized flood.

ANSWER: C

Explanation:

Set L2 Unknown Unicast to Flood. is correct because App_2 is a silent endpoint: it does not independently originate traffic, so the fabric has no source frame from which to learn its MAC address and location on Leaf 2. When a frame is sent to App_2's as-yet unknown destination MAC address, configuring L2 unknown-unicast handling to flood causes the bridge domain to distribute that frame across the applicable Layer 2 flood domain, including the interface toward S2. App_2 can then receive the traffic and respond. Its response arrives on Leaf 2 with App_2's MAC address as the source, allowing Cisco ACI to create the endpoint entry and associate it with the correct leaf, port, VLAN encapsulation, and endpoint group.

This behavior is specifically useful where a workload cannot be proactively learned because it is quiet until it receives traffic. The setting applies to unknown unicast destination traffic in the bridge domain and provides the initial reachability event needed to elicit traffic from the silent endpoint. Cisco documents unknown-unicast flooding as a bridge-domain forwarding behavior and describes endpoint learning from the source MAC information observed by leaf switches. See the Cisco APIC Layer 2 Configuration Guide and the Cisco APIC Basic Configuration Guide.

QUESTION NO: 19

What are two requirements for the IPN network when implementing a Multi-Pod ACI fabric? (Choose two.)

A. EIGRP routing

B. PIM ASM multicast routing

C. BGP routing

D. VLAN ID 4

E. OSPF routing

ANSWER: B E

Explanation:

For a Cisco ACI Multi-Pod fabric, the Inter-Pod Network (IPN) must provide routed IP connectivity between the spine switches in separate pods and must transport the multicast traffic used by the ACI overlay. **OSPF routing** is required on the IPN-facing spine interfaces: the IPN provides the external Layer 3 transit network, while OSPF establishes the required routing adjacencies and advertises the infrastructure loopback addresses that are used for inter-pod reachability. The IPN must also provide **PIM ASM multicast routing**. ACI uses multicast group addresses for forwarding bridge-domain BUM traffic and for other overlay control or data-plane functions across pods. PIM Any-Source Multicast, together with an available rendezvous point, enables this multicast traffic to be forwarded through the IPN between pods. The IPN must be designed as a highly available Layer 3 network with sufficient MTU for encapsulated traffic, but the routing and multicast capabilities named here are the two protocol requirements represented by the choices. Cisco documents the Multi-Pod IPN as an OSPF-based routed transit network that requires multicast support for the inter-pod overlay. See the Cisco APIC Layer 3 Networking Configuration Guide and the Cisco ACI Multi-Pod white paper.

QUESTION NO: 20

An ACI administrator notices a change in the behavior of the fabric. Which action must be taken to determine if a human intervention introduced the change?

A. Inspect event records in the APIC UI to see all actions performed by users.

B. Inspect /var/log/audit_messages on the APIC to see a record of all user actions.

C. Inspect audit logs in the APIC UI to see all user events.

D. Inspect the output of show command history in the APIC CLI.

ANSWER: C

Explanation:

Inspect audit logs in the APIC UI to see all user events. APIC audit logging is intended to provide an accountability trail for administrative activity in the Cisco ACI fabric. Audit records identify changes made through APIC and include useful context such as the affected managed object, the user identity, the time of the action, and the nature of the modification. Reviewing these records lets an administrator correlate the time at which the fabric behavior changed with a configuration action performed by a human operator. This is particularly important because ACI policy changes can affect endpoint learning, contracts, routing, bridge domains, interfaces, and other fabric behavior immediately after the policy is deployed. In the APIC GUI, audit logs provide a centralized, searchable view of these user-generated events, allowing the administrator to identify who made a change and investigate its scope. Cisco documents audit logs as the source for tracking configuration changes and user activity in APIC. See the Cisco APIC audit-log configuration guide and the Cisco APIC troubleshooting guide.

QUESTION NO: 21

Which two types of interfaces are supported on border leaf switches to connect to an external router? (Choose two.)

A. subinterface with VXLAN tagging

B. subinterface with 802.1Q tagging

C. FEX host interface

D. out of band interface

E. Switch Virtual Interface

ANSWER: B E

Explanation:

In Cisco ACI, a border leaf connects the fabric to external Layer 3 networks through a Layer 3 Out configuration. An external router can be attached using a subinterface with 802.1Q tagging. This design allows a single physical link between the border leaf and router to carry traffic for multiple VLANs, with each tagged subinterface mapped to the appropriate external routed network and routing configuration. It is a common choice when the router connection is a VLAN trunk.

A Switch Virtual Interface is also supported for an external-router connection. An SVI provides the Layer 3 gateway interface for a VLAN on the border leaf, allowing the external router to peer across that VLAN. The SVI approach is appropriate when external connectivity is delivered through a switched VLAN and the Layer 3 Out requires an IP interface associated with that VLAN. Cisco documents routed interfaces, subinterfaces, and SVIs as Layer 3 Out interface deployment models; among the listed choices, the tagged subinterface and SVI models are the supported interface types requested. See the Cisco APIC Layer 3 Configuration Guide and Cisco's Application Centric Infrastructure documentation.

QUESTION NO: 22

An engineer must securely export Cisco APIC configuration snapshots to a secure, offsite location The exported configuration must be transferred using an encrypted tunnel and encoded with a platform-agnostic data format that provides namespace support Which configuration set must be used?

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: B

Explanation:

Option B is correct because it represents the configuration set that combines Secure File Transfer Protocol (SFTP) for the remote backup destination with XML as the export format. APIC snapshot and configuration exports can be sent to a configured remote location, and SFTP provides an SSH-protected, encrypted transport channel. This satisfies the requirement that the offsite transfer occur through an encrypted tunnel while protecting the snapshot contents in transit.

XML is the appropriate platform-neutral encoding when namespace support is required. XML namespaces qualify element and attribute names with a namespace URI, allowing structured configuration data from separate schemas to be represented without naming ambiguity. This makes XML particularly suitable for portable, structured management data exchanged between systems. In APIC, the export policy identifies the remote location and export behavior, while the remote-location

configuration supplies the secure transfer details. Using the SFTP destination and XML encoding together therefore meets both the secure-transfer and namespace-aware data-format requirements for the exported configuration snapshot.

See Cisco's APIC REST API documentation for APIC's XML and JSON data representations, and the Cisco APIC support documentation for APIC configuration and backup guidance.

QUESTION NO: 23

Cisco ACI fabric has three different endpoints S1, S2. and S3. These endpoints must communicate with each other without contracts. These objects have been created in APIC:

•Two EPGs named DNS_EPG and Database_EPG

•Two application profiles. PROD_App and Data_App

•Two bridge domains DNS_BD and Database_BD

•PROD_APP and Database_BD mapped to Tenant PROD

•Data_App and DNS_BD mapped to Tenant Data

Which set of actions completes the fabric configuration?

A. Add S1, S2, S3 under Database_EP

B. Add S1, S2, S3, under DNS_EP

C. Add S1, S2, S3 under DNS_EP

D. Add S1, S2, S3 under Database_EP

E. Place S1, S2, and S3 in DNS_EPG and associate DNS_EPG with DNS_BD in tenant Data.

ANSWER: E

Explanation:

Place S1, S2, and S3 in **DNS_EPG**, and associate that EPG with **DNS_BD** in tenant **Data**. This produces a valid ACI logical configuration because Data_App and DNS_BD already belong to the same tenant, Data. The EPG should be created under the application profile in that tenant and linked to its bridge domain; the bridge-domain association provides the Layer 2 forwarding context for the endpoint group.

Endpoints within the same EPG are permitted to communicate by default without a contract. Therefore, assigning all three endpoints to one EPG meets the requirement for unrestricted communication among S1, S2, and S3 while avoiding a contract configuration. This assumes that intra-EPG isolation has not been enabled, because that feature intentionally prevents endpoint-to-endpoint communication inside an EPG. The endpoint attachment configuration must also use the VLAN and interface or virtual-port-group settings that place each endpoint into DNS_EPG.

Cisco describes an EPG as the policy grouping used to apply connectivity and security rules, and documents that EPGs are associated with bridge domains for forwarding. See the Cisco APIC EPG configuration guide and the Cisco APIC bridge-domain configuration guide.

QUESTION NO: 24

Which two protocols are used for fabric discovery in ACI? (Choose two.)

A. LLDP

B. OSPF

C. CDP

D. DHCP

E. ISIS

ANSWER: A D

Explanation:

LLDP and DHCP are used during Cisco Application Centric Infrastructure fabric discovery and bootstrap. LLDP provides the neighbor-discovery mechanism that enables switches to identify directly connected fabric devices and determine the physical leaf-and-spine topology. This topology information is reported to the controller so that it can validate cabling, identify node adjacencies, and build the fabric inventory. LLDP must therefore be operational on the relevant fabric-facing interfaces for automatic discovery to work correctly.

DHCP supports the initial bootstrap workflow for discovered switches. A newly connected switch uses DHCP to obtain initial network information and to learn the controller-related parameters needed to join and register with the fabric. After the node is discovered and provisioned, the controller can assign the appropriate fabric identity and configuration. Together, LLDP supplies physical-neighbor and topology discovery, while DHCP supplies the initial addressing and bootstrap information that allows a discovered device to become part of the managed fabric. Cisco describes these fabric bring-up and connectivity processes in its APIC Fabric Connectivity Configuration Guide and provides ACI architecture context at Cisco Application Centric Infrastructure.

QUESTION NO: 25

The unicast routing feature is enabled on the bridge domain. Which two conditions enable the Cisco ACI leaf to learn a source IP as a local endpoint? (Choose two.)

A. Through Ethernet traffic received in a bridge domain.

B. IP traffic routed through an SVI.

C. Through VXLAN traffic received on the uplink.

D. IP traffic routed through a Layer 3 Out.

E. Through ARP received on an SVI.

ANSWER: B E

Explanation:

With unicast routing enabled for a bridge domain, the leaf switch can associate an IP address with a locally learned endpoint when it sees traffic that establishes the endpoint's Layer 3 presence on the bridge-domain SVI. **IP traffic routed through an SVI** provides a source IP address in a packet that is being processed by the local leaf as the bridge domain's routed interface. The leaf can therefore bind that source IP to the local endpoint's MAC address, VLAN/encapsulation, and access port information.

Through ARP received on an SVI also enables local IP endpoint learning. ARP carries the sender protocol address and sender hardware address, and receipt on the SVI provides the bridge-domain context needed to create or update the IP-to-MAC endpoint association. This learning is important to ACI's endpoint tracking and host-mobility handling because the fabric maintains endpoint location and IP binding information separately from ordinary routing-prefix knowledge. These two cases specifically describe locally observed Layer 3 information at the SVI for the routing-enabled bridge domain. Cisco's ACI documentation discusses bridge-domain unicast routing and endpoint learning behavior in the Cisco APIC Basic Configuration Guide and provides endpoint-learning architecture background in the Cisco ACI white paper.

QUESTION NO: 26

An engineer must advertise a bridge domain subnet out of the ACI fabric to an OSPF neighbor. Which two configuration steps are required? (Choose two.)

A. Configure Subnet scope to Advertised Externally

B. Add External Subnet for External EPG flag under External EPG.

C. Create Route Control Profile with the export direction under External EPG.

D. Add L3Out profile to the bridge domain using Associated L3Outs section

E. Configure the Subnet under the EPG level.

ANSWER: A D

Explanation:

To originate a bridge-domain subnet into an external OSPF routing domain, the subnet must first be explicitly marked with the *Advertised Externally* scope. This scope instructs APIC that the subnet is eligible to be exported from the fabric through Layer 3 external connectivity, rather than being used only for internal tenant routing and endpoint reachability. The subnet can be configured at the bridge-domain level, which is the appropriate location when it represents the BD's routed gateway subnet.

The bridge domain must also be associated with the relevant L3Out. That association identifies the external routed connection through which the bridge-domain prefix can be advertised. Once the BD is linked to the L3Out and the subnet has the external-advertisement scope, the ACI policy model can advertise the prefix to the configured OSPF peer, subject to the L3Out's routing configuration and any applicable route-control policy. Cisco describes bridge-domain subnet scopes and external routing behavior in its APIC Layer 3 Networking Configuration Guide. Additional ACI L3Out design and configuration resources are available from Cisco APIC documentation.

QUESTION NO: 27 - (DRAG DROP)

An engineer must configure RADIUS authentication with Cisco ACI for remote authentication with out-of-band management access. Drag and drop the RADIUS configuration steps from the left into the required implementation order on the right. Not all steps are used.

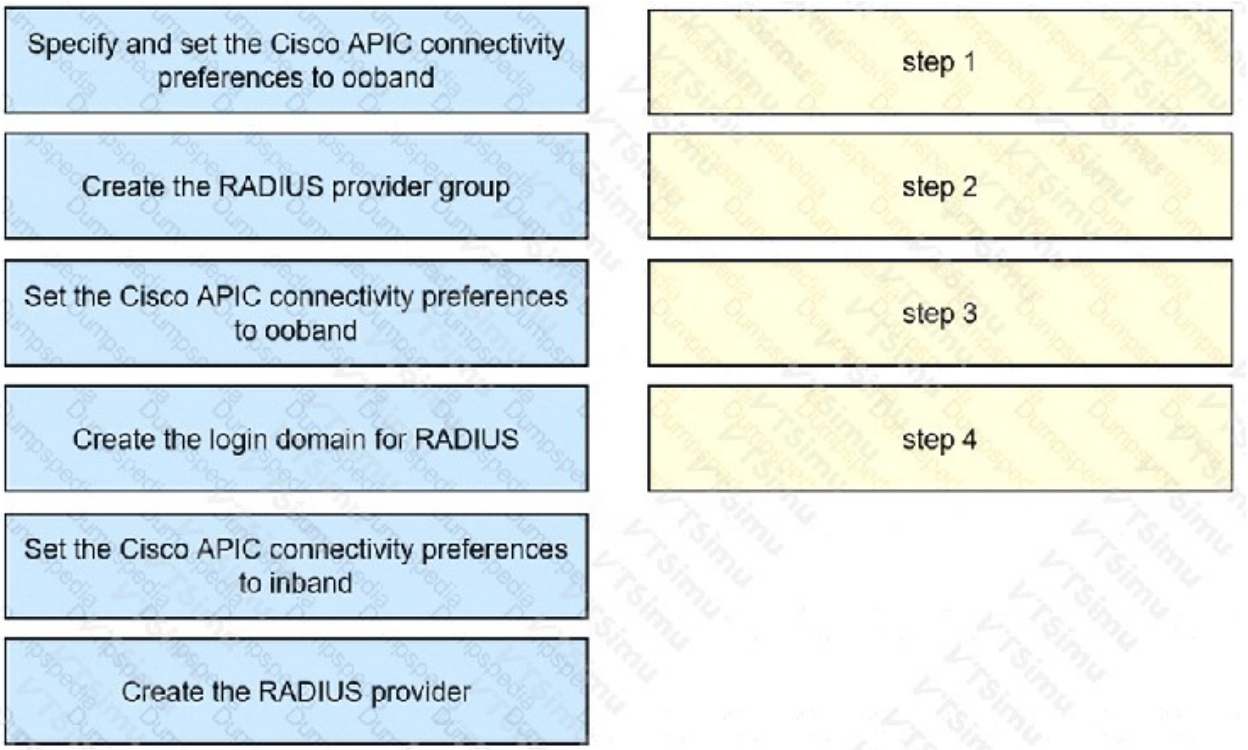

ANSWER:

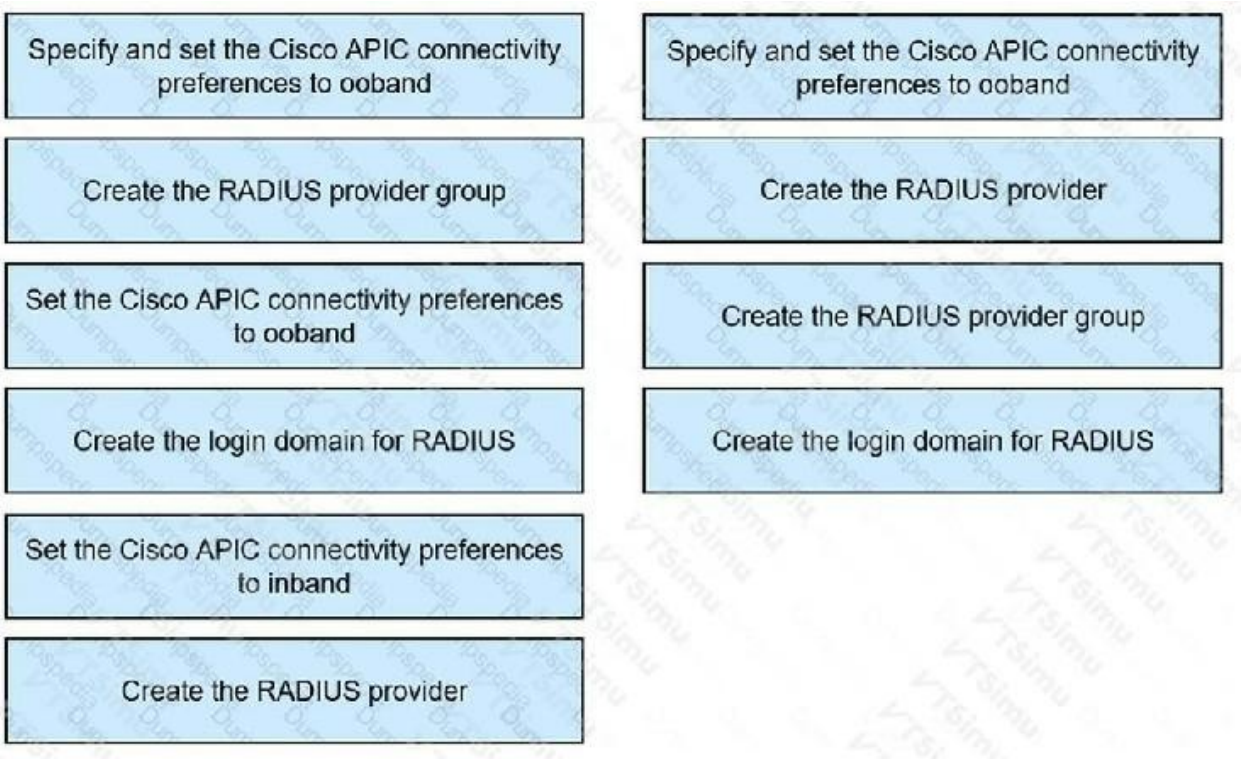

Explanation:

For RADIUS authentication over the out-of-band management network, start by specifying and setting the Cisco APIC connectivity preferences to ooband. This establishes the management-path context APIC uses when it communicates with an external AAA service. In other words, APIC must be told that the RADIUS server is reached through the out-of-band management infrastructure before the authentication objects are built around that communication path. Cisco's ACI documentation describes external authentication configuration as an APIC AAA configuration that uses defined RADIUS providers and provider groups; see the Cisco APIC Security Configuration Guide.

After the connectivity preference is in place, create the RADIUS provider. A provider represents one actual RADIUS server and contains the server's addressing and shared-secret information. Next, create the RADIUS provider group. The group is the logical AAA object that contains one or more configured RADIUS providers and is what allows APIC to apply the remote-server definition as an authentication source. This grouping also provides a scalable design, since additional RADIUS servers can be incorporated without changing the basic authentication-domain model.

Finally, create the login domain for RADIUS. The login domain ties the external authentication method to the RADIUS provider group and is the domain users select or enter when logging in. At that point, the domain can direct authentication requests to the established provider group through the configured out-of-band path. The resulting required order is: specify and set the Cisco APIC connectivity preferences to ooband, create the RADIUS provider, create the RADIUS provider group, and create the login domain for RADIUS.

QUESTION NO: 28

An engineer must ensure that Cisco ACI flushes the appropriate endpoints when a topology change notification message is received in an MST domain. Which three steps are required to accomplish this goal? (Choose three.)

A. Enable the BPDU interface controls under the spanning tree interface policy.

B. Configure a new STP interface policy.

C. Bind the spanning tree policy to the switch policy group.

D. Associate the STP interface policy to the appropriate interface policy group.

E. Create a new region policy under the spanning tree policy.

F. Map VLAN range to MAT instance number.

ANSWER: A C D

Explanation:

Cisco ACI can use received spanning-tree topology-change information to age or flush affected endpoint entries, preventing traffic from being forwarded toward a stale location after an external Layer 2 topology change. **Enable the BPDU interface controls under the spanning tree interface policy.** enables the leaf interface to process the required BPDU-related signaling. **Associate the STP interface policy to the appropriate interface policy group.** applies those interface-level spanning-tree controls to the actual leaf ports connected to the MST domain. Finally, **Bind the spanning tree policy to the switch policy group.** applies the switch-level MST configuration to the participating leaf switches, allowing the fabric to interpret the MST topology-change notification in the applicable spanning-tree context. Together, the interface policy, interface-policy-group association, and switch-policy-group association provide the policy attachment chain required for the leaf to receive and act on MST topology-change events. Cisco documents STP configuration in ACI as a combination of interface-level STP settings and switch-level MST policy configuration and associations. See the Cisco APIC Layer 2 Networking Configuration Guide: Configuring Spanning Tree and the Cisco APIC Basic Configuration Guide.

QUESTION NO: 29

An engineer is implementing a connection that represents an external bridged network. Which two configurations are used? (Choose two.)

A. Layer 2 remote fabric

B. Layer 2 outside

C. Layers 2 internal

D. Static path binding

E. VXLAN outside

ANSWER: A B

Explanation:

Layer 2 outside and **Layer 2 remote fabric** are the configurations used to represent external bridged Layer 2 connectivity in Cisco ACI. A Layer 2 outside connection, commonly called an L2Out, provides a bridge-domain extension from an ACI endpoint group to an external Layer 2 network. It is appropriate when endpoints in the ACI fabric must exchange Ethernet frames with devices or VLANs that reside beyond the fabric, while remaining in the same Layer 2 domain.

Layer 2 remote fabric connectivity is used when the bridged Layer 2 network is extended to a remote ACI fabric. This design allows the same Layer 2 service and associated endpoint connectivity to span fabric locations through the supported remote-fabric mechanism. Both constructs therefore model an external bridged network rather than an internally attached endpoint group or a routed external connection. Cisco's Layer 2 networking documentation describes L2Out as the construct for external Layer 2 connectivity and documents remote-fabric Layer 2 extension as a distinct ACI capability. See the Cisco APIC Layer 2 Networking Configuration Guide and Cisco Application Centric Infrastructure documentation.

QUESTION NO: 30

The customer is looking for redundant interconnection of the existing network to the new ACI fabric. Unicast and multicast traffic must be routed between the two networks. Which L3Out implementation meets these requirements?

A.

B.

C.

D.

E. A routed-interface L3Out deployed on two leaf switches, with PIM enabled toward redundant external routers and a unicast routing protocol configured.

ANSWER: E

Explanation:

A routed-interface L3Out deployed across two leaf switches, with PIM enabled toward redundant external routers and a unicast routing protocol configured, meets the requirement. The two independent leaf-to-external-router paths provide physical and control-plane redundancy, so routing can continue if a leaf, link, or external-router path fails. The L3Out's unicast routing configuration—such as OSPF or BGP—exchanges reachability between the ACI tenant VRF and the existing network. PIM must be enabled on the external-facing L3Out interfaces to establish multicast-routing adjacencies and allow multicast traffic to be forwarded between the ACI fabric and the external multicast domain. The external network and ACI multicast design must also use compatible multicast settings, including the required PIM mode and rendezvous-point design where applicable. In ACI, multicast routing is enabled at the VRF and L3Out interface level, while the use of separate leaf switches for the routed attachments avoids a single attachment-point failure. Cisco documents L3Out routed-interface connectivity and dynamic routing in the Cisco APIC Layer 3 Configuration Guide, and multicast configuration requirements in the Cisco APIC Multicast Configuration Guide.

QUESTION NO: 31

Refer to the exhibit.

Which two objects are created as a result of the configuration? (Choose two.)

A. application profile

B. attachable AEP

C. bridge domain

D. endpoint group

E. VRF

ANSWER: C E

Explanation:

The configuration creates a bridge domain and a VRF. In the Cisco ACI managed-object model, the `fvBD` object represents a bridge domain. A bridge domain defines the forwarding and flooding scope for endpoints and provides the association point for a subnet and a routing context. The configured `fvBD` instance named `bd1` therefore creates the bridge domain object.

The `fvCtx` object represents a VRF, also called a context in the ACI object model. A VRF establishes an isolated routing and address space within a tenant. The configured `fvCtx` instance named `pvn1` consequently creates the VRF object. These tenant-level networking objects are commonly related by associating the bridge domain with its routing context, allowing endpoints in the bridge domain to use the VRF's routing table when routing is enabled.

Cisco documents bridge domains as tenant networking constructs and describes VRFs as the routing contexts used for traffic isolation and routing in an ACI fabric. See the Cisco APIC Basic Configuration Guide: Tenants and the Cisco APIC REST API Configuration Guide: Managed Object Model.

QUESTION NO: 32

Refer to the exhibit. An engineer must configure an L3Out peering with the backbone network. The L3Out must forward unicast and multicast traffic over the link. Which two methods should be used to configure L3Out to meet these requirements? (Choose two.)

A. Layer 3 routed port

B. VPC with SVI

C. port channel with SVI

D. Layer 3 routed subinterface

E. Layer 3 floating SVI

ANSWER: A D

Explanation:

Layer 3 routed port and Layer 3 routed subinterface are the appropriate L3Out interface methods when the external backbone connection must carry both unicast routing and multicast traffic. A routed port provides a direct Layer 3 adjacency on a physical leaf-interface port. A routed subinterface provides the same Layer 3 function while using an IEEE 802.1Q VLAN tag, which is appropriate when the backbone handoff is logically segmented on a trunk. In either case, the L3Out can establish unicast routing adjacencies and can be enabled for multicast routing by configuring the required multicast settings, such as Protocol Independent Multicast, on the external Layer 3 interface and associated routing infrastructure. Cisco ACI multicast documentation identifies routed interfaces and routed subinterfaces as supported external interface types for connecting multicast-capable L3Outs. These interface models provide the direct Layer 3 forwarding path needed for traffic exchanged between ACI bridge domains and the external multicast domain. See Cisco's APIC Multicast Configuration Guide and the Cisco ACI Basic Configuration Guide for L3Out interface and multicast configuration concepts.

QUESTION NO: 33

An engineer must connect Cisco ACI fabric using Layer 2 with external third-party switches. The third-party

switches are configured using 802.1s protocol. Which two constructs are required to complete the task?

(Choose two.)

A. spanning tree policy for mapping MST Instances to VLANs

B. MCP policy with PDU per VLAN enabled

C. MCP instance policy with administrative state disabled

D. dedicated EPG for native VLAN

E. static binding of native VLAN in all existing EPGs

ANSWER: A C

Explanation:

IEEE 802.1s is Multiple Spanning Tree Protocol (MSTP). For an ACI leaf to participate correctly in an MSTP-based Layer 2 connection, a spanning tree policy must define the MST region parameters and, critically, the mapping of VLANs to MST instances. The ACI-side MST configuration must align with the external switches' MST region configuration so that the relevant VLANs are evaluated in the intended spanning-tree instance and the Layer 2 topology can converge consistently.

An MCP instance policy with the administrative state disabled is also required for this external Layer 2 attachment. Mis-Cabling Protocol is an ACI fabric loop-detection mechanism, whereas the adjacent third-party network is using standards-based MSTP for loop prevention. Disabling MCP on this attachment prevents MCP operation from interfering with the external switching domain while MSTP supplies the spanning-tree behavior at the interconnection.

These settings are applied through the appropriate ACI policy objects and attached to the interface or interface-policy group used for the external switch connection. Cisco documents Layer 2 and spanning-tree configuration concepts in its Cisco APIC Layer 2 Configuration Guide; additional ACI switching design guidance is available in the Cisco Live ACI Switching presentation.

QUESTION NO: 34

A Cisco ACI environment consists of multiple silent hosts that are often relocated between leaf switches. When the host is relocated, the bridge domain takes more than a few seconds to relearn the host's new location. The requirement is to minimize the relocation impact and make the ACI fabric relearn the new location of the host faster. Which action must be taken to meet these requirements?

A. Set Unicast Routing to Enabled.

B. Configure ARP Flooding to Enabled.

C. Set L2 Unknown Unicast to Hardware Proxy.

D. Configure IP Data-Plane Learning to No.

ANSWER: B

Explanation:

Configure ARP Flooding to Enabled. is the appropriate bridge-domain setting for frequently moved silent endpoints. A silent host does not originate enough traffic after relocation for the fabric to promptly learn its new attachment point from normal data-plane learning. If ARP flooding is disabled, the fabric can proxy an ARP reply based on its existing endpoint information. That behavior can retain use of the endpoint's previously learned location until the endpoint database is refreshed by another event or aging process.

Enabling ARP flooding causes ARP requests to be flooded within the bridge domain rather than answered solely through the fabric's proxy information. The relocated endpoint receives the ARP request on its new leaf-facing interface and sends an ARP reply. That reply supplies a source MAC and IP learning event at the endpoint's current location, allowing the fabric to update the endpoint record much more quickly. This is specifically useful where endpoint mobility occurs but hosts can remain otherwise quiet, because ARP activity provides the traffic needed to trigger timely relearning.

Cisco documents ARP flooding as a bridge-domain Layer 2 forwarding setting in the Cisco APIC Layer 2 Configuration Guide. Additional ACI configuration documentation is available from the Cisco APIC product documentation page.

QUESTION NO: 35

A customer must deploy three Cisco ACI based data centers. Each site must be separated from the others. Which characteristic of Cisco ACI Multi-Pod makes it unsuitable for this deployment?

A. creates a virtual pod in the remote location

B. requires all pods to share the same Cisco APIC cluster

C. has distance and scale limitations

D. places leaf switches in the remote site that belong to the same fabric as at the headquarters site

ANSWER: B

Explanation:

requires all pods to share the same Cisco APIC cluster is correct because Cisco ACI Multi-Pod extends one ACI fabric across multiple pods, rather than creating separately managed ACI fabrics at each location. All pods are part of the same logical fabric and are managed by a single, shared APIC cluster. Consequently, policy, fabric inventory, and the fault domain are not independent per data center in the way required when sites must be separated from one another. A shared APIC cluster also requires the intersite connectivity and latency characteristics necessary for APIC cluster communication, so the sites cannot operate as autonomous ACI domains. For geographically or operationally separate data centers, Cisco ACI Multi-Site is the appropriate architectural model: each site retains its own independent APIC cluster and ACI fabric, while an intersite policy and connectivity framework provides controlled communication between sites. Cisco describes Multi-Pod as a single fabric consisting of multiple pods under one APIC cluster in its ACI Multi-Pod Configuration Guide. Cisco's ACI Multi-Site overview describes the independent-site model used to connect multiple separate ACI fabrics.