PassQueen
a.k.a dumpsqueen.com

Implementing Cisco Application Centric Infrastructure - Advanced (600-660 DCACIA)

Cisco 600-660

Version Demo

Total Demo Questions: 10

Total Premium Questions: 100

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Cisco ACI Advanced Packet Forwarding	48
Topic 2, Cisco ACI Security	15
Topic 3, Cisco ACI External Network Connectivity	21
Topic 4, Cisco ACI Management	4
Topic 5, Cisco ACI Integration	12
Total	100

QUESTION NO: 1

Where are STP BPDUs flooded in Cisco ACI fabric?

A. in the bridge domain VLAN

B. in the native VLAN ID

C. in the access encapsulation VLAN part of different VLAN pools

D. in the VNID that is assigned to the FD VLAN

ANSWER: A

Explanation:

STP BPDUs are flooded in the bridge domain VLAN. Cisco ACI does not participate in Spanning Tree Protocol as a conventional switch would; instead, the fabric transparently carries received BPDUs so that connected external switches can maintain their STP topology. For this purpose, ACI uses the bridge domain's internal VLAN context to flood the BPDU frames to the relevant attachment points in that bridge domain. This preserves Layer 2 control-plane connectivity among endpoints attached to the same bridge domain while maintaining ACI's policy-based forwarding model.

The bridge domain defines the Layer 2 forwarding and flooding scope in ACI, including the handling required for broadcast and control traffic. Consequently, BPDU propagation follows that bridge-domain scope rather than relying on a native VLAN, an access VLAN-pool encapsulation, or a user-facing VLAN assignment. Understanding this behavior is important when integrating ACI with traditional Layer 2 networks, particularly where STP remains active on externally connected switches. Cisco documents ACI's transparent handling of STP BPDUs and its bridge-domain-based Layer 2 behavior in its Cisco ACI Layer 2 networking white paper.

QUESTION NO: 2

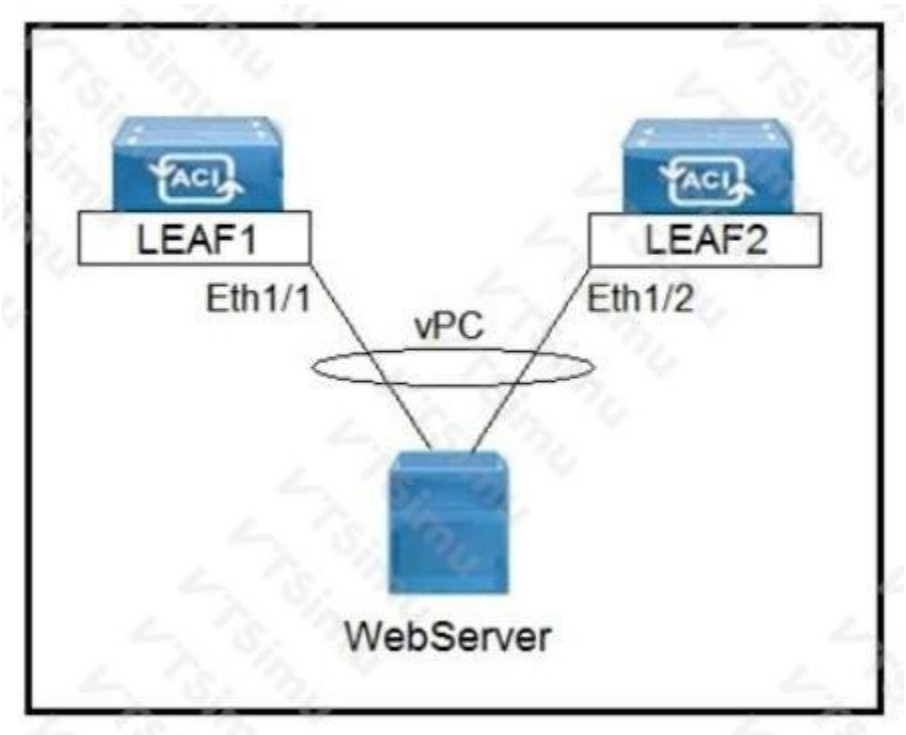

Refer to the exhibit. Which three actions should be taken to implement the vPC in the Cisco ACI fabric? (Choose three.)

A. Select a common vPC interface policy group

B. Select individual interface profiles

C. Select common interface profiles

D. Select individual switch profiles

E. Select common switch profiles

ANSWER: A B E

Explanation:

A Cisco ACI vPC connects a single downstream device to a pair of leaf switches while presenting the two links as one logical port channel to that device. The two member links require a common vPC interface policy group. This shared policy group carries the vPC-specific configuration and ensures that both leaf-facing connections use consistent port-channel behavior, including the vPC identity and associated interface policies.

Individual interface profiles are required because an interface profile defines the physical port selectors applied on a leaf switch. Separate profiles permit the relevant ports on each vPC peer leaf to be selected independently, while both selectors can reference the same vPC interface policy group. A common switch profile is then used to group the two vPC peer leaf switches and associate the applicable interface profiles with that leaf pair. This shared switch-profile scope is what establishes the common access-policy context for the vPC deployment. Cisco documents vPC as a pair-based leaf configuration using shared vPC policy settings and the appropriate leaf/interface profile associations. See the Cisco APIC Virtual Networking Configuration Guide and Cisco's Application Centric Infrastructure documentation portal.

QUESTION NO: 3

What is the purpose of the Forwarding Tag (FTAG) in Cisco ACI?

A. FTAG is used in Cisco ACI to add a label to the iVXLAN traffic in the fabric to apply the correct policy.

B. FTAG is used in Cisco ACI to add a label to the VXLAN traffic in the fabric to apply the correct policy.

C. FTAG trees in Cisco ACI are used to load balance unicast traffic.

D. FTAG trees in Cisco ACI are used to load balance multi-destination traffic.

ANSWER: D

Explanation:

FTAG trees in Cisco ACI are used to load balance multi-destination traffic. In the ACI fabric, multi-destination traffic includes broadcast, unknown unicast, and multicast traffic that must be replicated and delivered to multiple endpoint locations. The fabric constructs multiple loop-free FTAG forwarding trees across the spine-and-leaf topology. Each tree provides a deterministic forwarding path for replicated traffic while avoiding loops and enabling use of available fabric links.

When a leaf switch encapsulates eligible multi-destination traffic for transport across the VXLAN fabric, it assigns an FTAG value that selects one of the available forwarding trees. Switches use that tag to make consistent forwarding decisions for the packet as it traverses the fabric. Having multiple FTAG trees distributes multi-destination flows over different paths, improving link utilization and scale rather than concentrating replicated traffic on one tree. This function is part of the fabric's underlay forwarding behavior and supports efficient delivery of VXLAN-encapsulated multi-destination traffic. Cisco's ACI Fundamentals documentation describes FTAGs as the mechanism used for multi-destination forwarding trees; see Cisco ACI Fundamentals and Cisco Application Centric Infrastructure.

QUESTION NO: 4

Which two configurations are required to redirect traffic through a Layer 4 to Layer 7 service node by using Cisco ACI policy-based redirect? (Choose two.)

A. Apply a service graph to a contract subject

B. Configure a redirect policy that identifies the service nodes

C. Configure a bridge domain subnet for each service node

D. Associate a VLAN pool directly with the contract subject

E. Disable contract policy enforcement in the VRF

ANSWER: A B

Explanation:

A **service graph must be applied to a contract subject** so that traffic matching the contract relationship is steered through the intended service chain. The service graph defines how the service device is inserted between the consumer and provider sides of the policy relationship. A **redirect policy that identifies the service nodes** is also required for policy-based redirect. The redirect policy specifies the service-node destinations to which matching traffic is redirected.

A redirect policy can use health tracking to remove unavailable service nodes from redirection decisions, but health tracking alone does not insert the service into a contract path. A bridge domain subnet supplies gateway addressing and a VLAN pool supplies VLAN encapsulations; neither defines redirect behavior. See the Cisco ACI Layer 4 through Layer 7 Services Configuration Guide.

QUESTION NO: 5

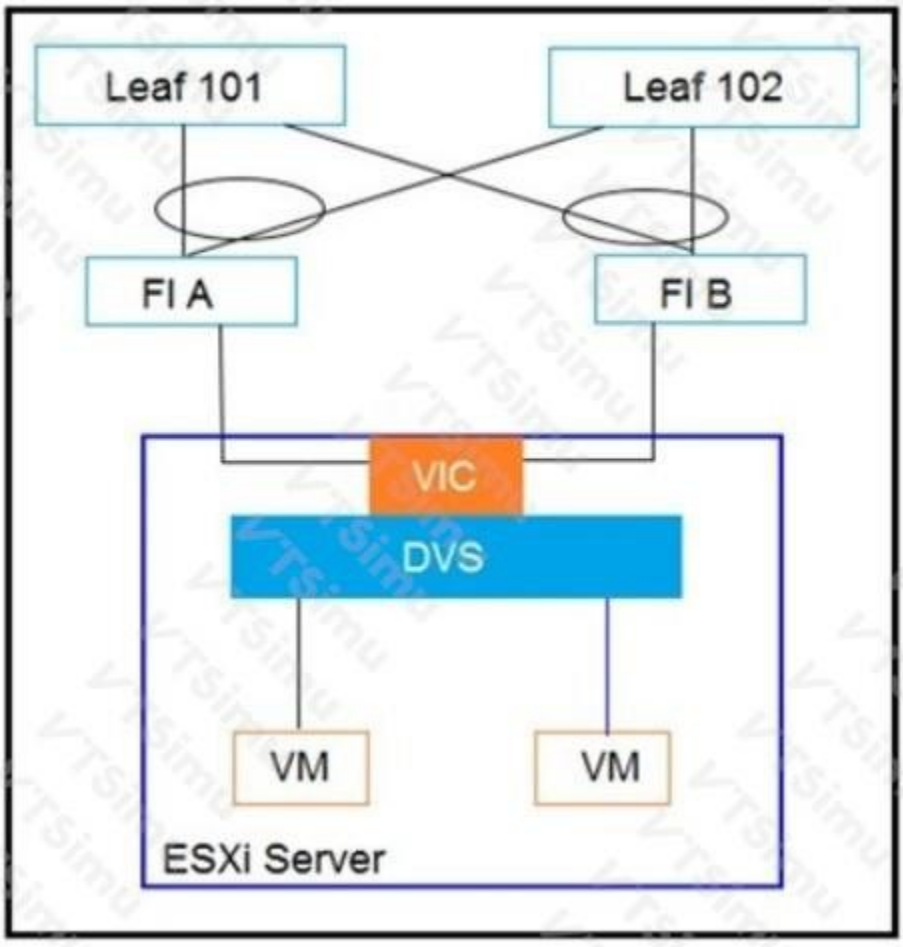

Refer to the exhibit. Between Cisco UCS FIs and Cisco ACI leaf switches, CDP is disabled, the LLDP is enabled, and LACP is in Active mode. Which two discovery protocols and load-balancing mechanism combinations can be implemented for the DVS? (Choose two.)

A. CDP enabled, LLDP disabled, and LACP Active

B. CDP disabled, LLDP enabled, and MAC Pinning

C. CDP enabled, LLDP disabled, and MAC Pinning

D. CDP enabled, LLDP enabled, and LACP Active

E. CDP enabled, LLDP disabled, and LACP Passive

F. CDP disabled, LLDP enabled, and LACP Passive

ANSWER: B F

Explanation:

CDP disabled, LLDP enabled, and MAC Pinning is valid because it preserves the required discovery-protocol state for the UCS Fabric Interconnect-to-ACI leaf adjacency: LLDP is the enabled discovery protocol and CDP remains disabled. MAC pinning is a supported UCS traffic-distribution approach in designs where uplinks are individually pinned rather than bundled from the server-facing virtual-switch perspective. It provides deterministic forwarding while retaining the required LLDP-based fabric discovery behavior.

CDP disabled, LLDP enabled, and LACP Passive is also valid. LACP forms a channel when one endpoint is active and the peer is passive; therefore, a passive participant can successfully negotiate with the ACI-facing active LACP endpoint. This selection also retains the stated LLDP-enabled and CDP-disabled discovery configuration. In an ACI/UCS design, protocol settings must be consistent on the physical adjacency, while LACP mode selection must permit negotiation across the port-channel peers. Cisco documents the ACI virtualization and external connectivity requirements in the Cisco APIC Virtualization Guide and UCS uplink/channel configuration concepts in the Cisco UCS Manager GUI Configuration Guide.

QUESTION NO: 6

Refer to the exhibit.

Two application profiles are configured in the same tenant and different VRFs. Which contract scope is configured to allow communication between the two application profiles?

A. global

B. VRF

C. application

D. tenant

ANSWER: D

Explanation:

The correct contract scope is **tenant**. In Cisco ACI, a contract defines the permitted communication between a provider and a consumer EPG, while its scope determines the policy domain in which that contract can be consumed and enforced. A tenant-scoped contract is designed for use within one tenant and can be applied to EPGs associated with different VRFs in that same tenant. This is the required scope when the application profiles reside under the same tenant but their EPGs use separate VRF contexts.

Although application profiles organize EPGs into application-centric logical groupings, contract applicability is determined by the tenant, VRF, or global policy scope rather than by the application-profile boundary. Setting the contract to tenant scope makes the contract available throughout the tenant, allowing the required provider/consumer relationship to be established across the two VRFs while retaining the tenant's administrative and policy isolation. Cisco documents that contract scope controls the visibility and applicability of contracts, including tenant-level use across VRFs within the tenant. See Cisco's Operating Cisco ACI documentation and the Cisco APIC Configuration Guide.

QUESTION NO: 7

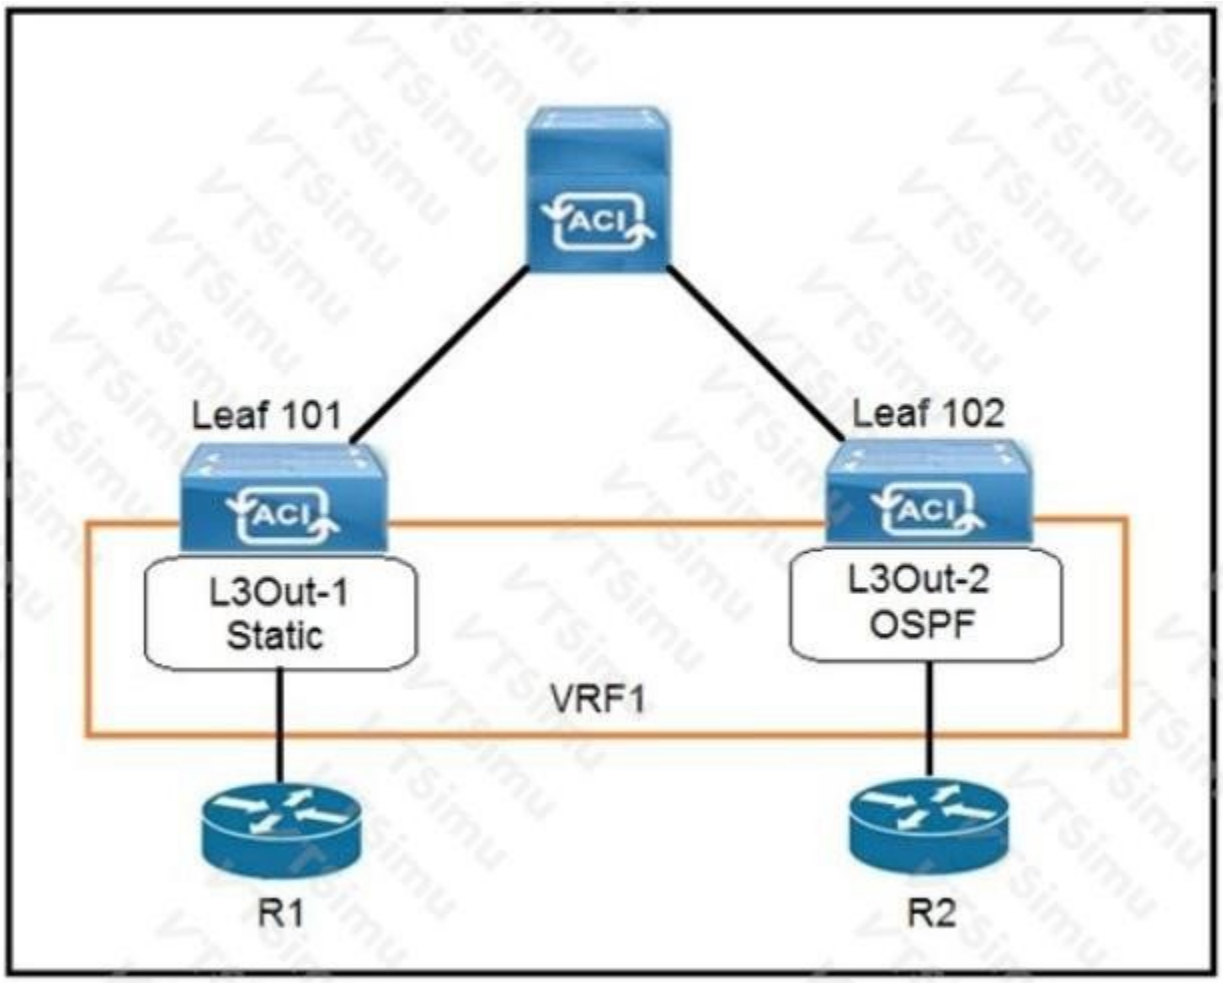

Refer to the exhibit. The 0.0.0.0/0 is configured as a default static route on L3Out-1. Which action should be taken for the 0.0.0.0/0 prefix to advertise out on L3Out-2 OSPF?

A. Enable Shared Security Import Subnet

B. Enable Aggregate Export Subnet

C. Enable Shared Route Control Subnet

D. Enable Export Route Control Subnet

ANSWER: B

Explanation:

Enable Aggregate Export Subnet is correct. In Cisco ACI, external EPG subnet scopes determine how prefixes are controlled and advertised through an L3Out. When a default static route, 0.0.0.0/0, learned or configured through one L3Out must be advertised through another L3Out using OSPF, the destination L3Out requires aggregate-export treatment for that default prefix. Configuring the default prefix with the Aggregate Export Subnet scope enables the fabric to originate and advertise the aggregate 0.0.0.0/0 route toward the OSPF peer on L3Out-2. This is especially important because the default route represents the least-specific aggregate and must be explicitly enabled for aggregate export rather than relying on advertisement of a more-specific external prefix. The route can then be propagated by the OSPF process associated with L3Out-2, allowing external devices on that connection to use ACI as the next hop for destinations not matched by more-specific routes. Cisco documents external EPG subnet scopes, including aggregate export behavior, in its APIC Layer 3 Configuration Guide. Additional ACI configuration documentation is available from the Cisco APIC support documentation.

QUESTION NO: 8

What does the VXLAN source port add to the overlay packet forwarding when it uses the hash of Layer 2, Layer 3, and Layer 4 headers of the inner packet?

A. ECMP

B. TCP optimization

C. disabled fragmentation

D. jumbo frames

ANSWER: A

Explanation:

ECMP is correct because VXLAN uses the outer UDP source-port field as an entropy value. A VTEP derives this value by hashing header fields from the encapsulated, or inner, traffic flow, such as Layer 2 addresses and Layer 3/Layer 4 information. Although packets belonging to a given flow receive a consistent hash-derived source port, separate flows are likely to receive different values. Underlay routers and switches can include the outer IP and UDP header fields in their load-balancing hash, allowing those devices to distinguish VXLAN flows that would otherwise share the same outer VTEP source and destination IP addresses.

This entropy enables Equal-Cost Multipath forwarding across multiple available underlay paths while maintaining predictable, in-order forwarding for packets in the same flow. In practical fabric designs, this behavior is important for distributing overlay traffic efficiently across routed spine-and-leaf links rather than concentrating all encapsulated traffic on a single path. The VXLAN specification explicitly describes the UDP source port as a mechanism for delivering entropy to the network, and Cisco's VXLAN documentation likewise identifies its role in enabling ECMP load sharing. See RFC 7348: Virtual eXtensible Local Area Network and Cisco VXLAN and Ethernet Virtualization.

QUESTION NO: 9 - (DRAG DROP)

DRAG DROP

A leaf receives unicast traffic that is destined to an unknown source, and spine proxy is enabled in the corresponding bridge domain. Drag and drop the Cisco ACI forwarding operations from the left into the order the operation occurs on the right.

Select and Place:

Answer Area

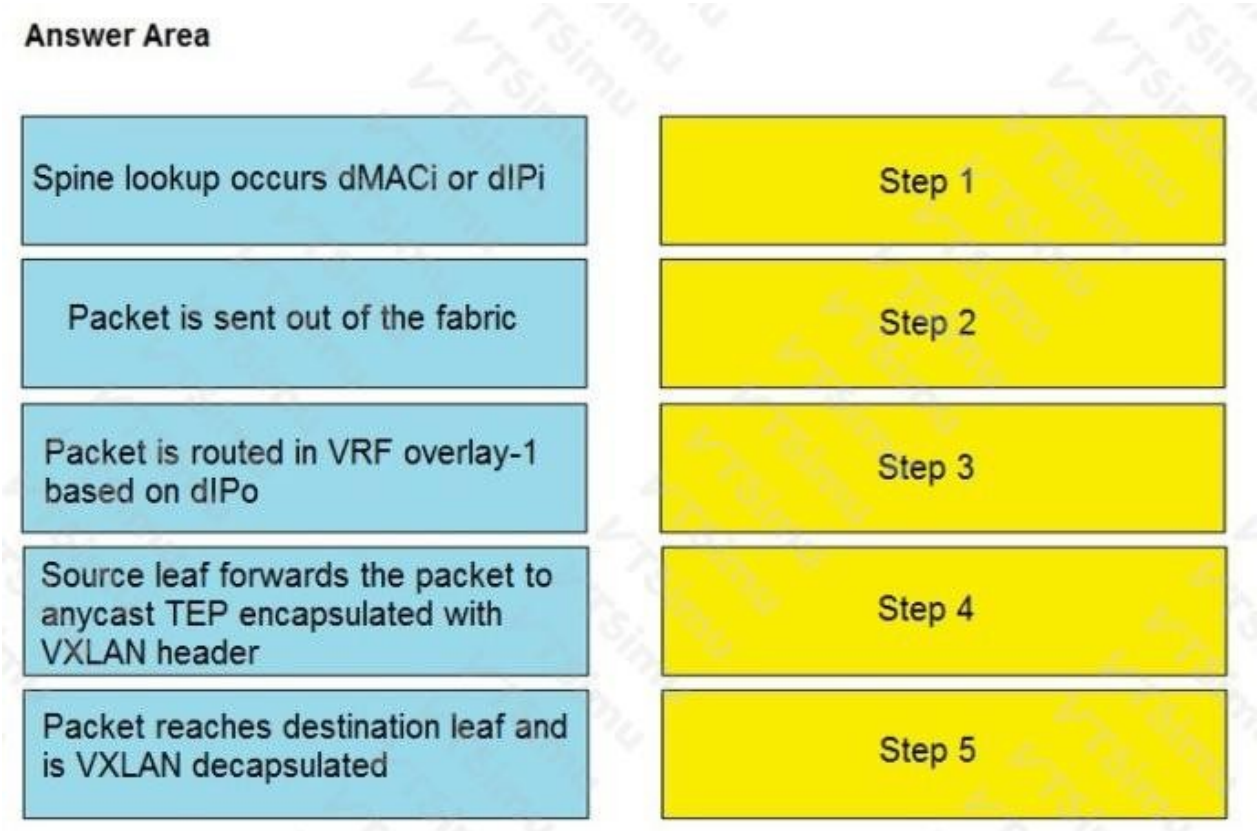

ANSWER:

Answer Area

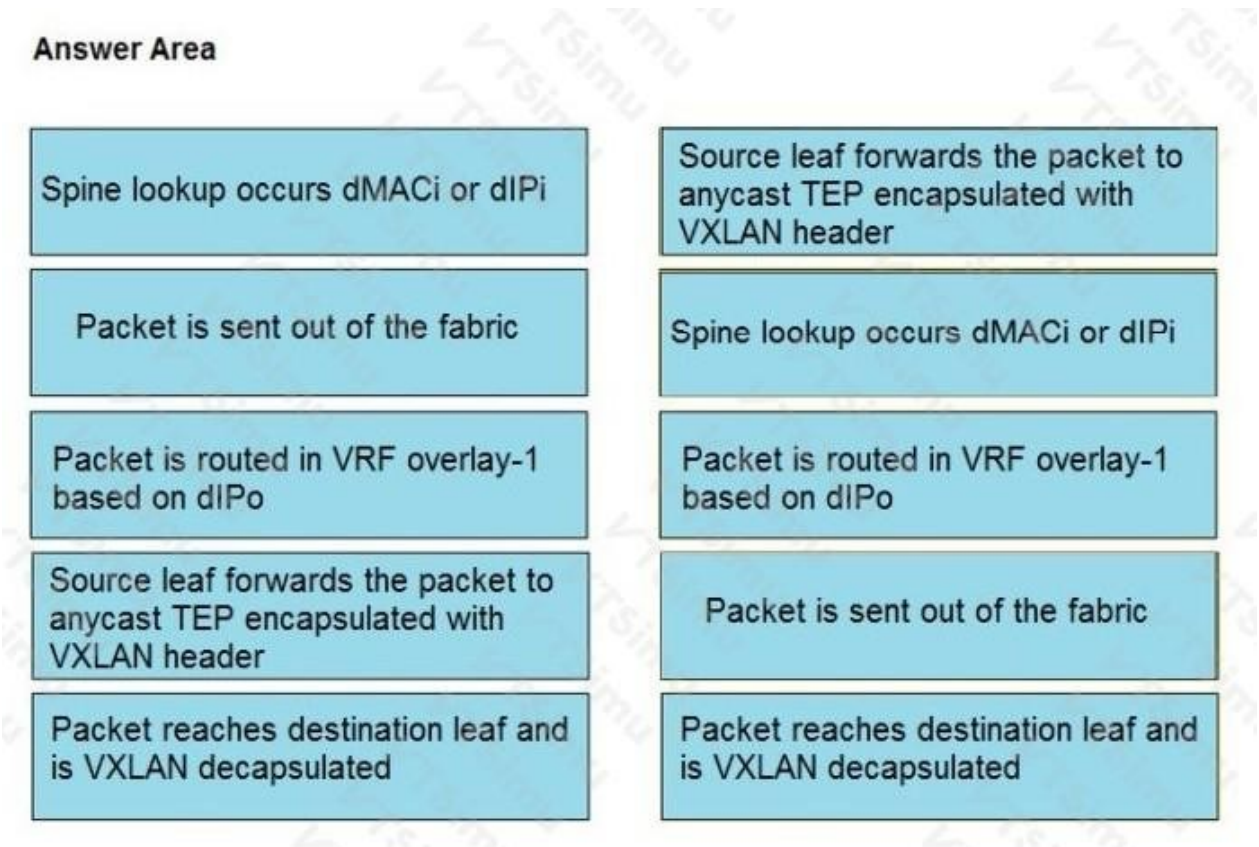

Explanation:

With spine proxy enabled, a leaf that cannot resolve the destination through its local endpoint information sends the traffic toward the spine-proxy anycast TEP. The source leaf applies VXLAN encapsulation so the packet can traverse the ACI IP fabric, and the anycast destination makes the spine layer the next forwarding decision point. This is the purpose of the first

operation: it transfers the unresolved forwarding decision from the source leaf to the spine proxy while preserving tenant forwarding context in the VXLAN header.

At the spine, the proxy performs the destination lookup using the destination MAC or the destination IP information available for the packet. This lookup supplies the endpoint and routing context needed to select the appropriate forwarding treatment. The packet is then routed in the VRF overlay based on the output destination IP. Performing this routing decision before sending traffic onward ensures that the packet is associated with the correct routed overlay path and destination-leaf direction.

After the lookup and VRF-overlay routing decision are complete, the spine forwards the encapsulated traffic out across the fabric to the selected destination leaf. The destination leaf receives the VXLAN packet and decapsulates it, restoring the original packet for the final local forwarding action. This ordering follows the normal ACI separation of responsibilities: the ingress leaf encapsulates toward the proxy, the spine provides proxy lookup and overlay routing intelligence, and the egress leaf terminates VXLAN for local delivery. Cisco's ACI documentation describes the fabric's VXLAN overlay and the role of leaf and spine nodes in endpoint forwarding: Cisco APIC Layer 2 Configuration Guide.

QUESTION NO: 10

An engineer connects a dual-homed server to a pair of Cisco ACI leaf switches by using a vPC. Which two actions are required to configure the leaf-side attachment? (Choose two.)

A. Configure a vPC protection group for the leaf pair

B. Configure a vPC interface policy group for the server links

C. Configure a separate bridge domain for each leaf switch

D. Enable PAgP desirable mode on the leaf interfaces

E. Configure a PC interface policy group on each leaf switch

ANSWER: A B

Explanation:

A **vPC protection group** must be configured for the two leaf switches that form the vPC pair. The protection group identifies the peer leaf switches and establishes the vPC relationship used for dual-homed endpoint connectivity. The engineer must also configure a **vPC interface policy group** and apply it to the appropriate interfaces on both leaf switches.

A regular port interface policy group represents an individual interface and does not create a multi-chassis port-channel. A PC interface policy group is used for a port-channel that terminates on one leaf switch. PAgP is not used by Cisco ACI for vPC negotiation; LACP is the supported protocol when dynamic link aggregation is required. See the Cisco APIC Layer 2 Configuration Guide.