PassQueen
a.k.a dumpsqueen.com

Implementing Cisco Application Centric Infrastructure - Advanced (DCACIA)

Cisco 300-630

Version Demo

Total Demo Questions: 16

Total Premium Questions: 164

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Advanced ACI Layer 3	46
Topic 2, Advanced ACI Layer 2	59
Topic 3, Advanced ACI Data Center Integrations	31
Topic 4, Advanced ACI Services Integration	18
Topic 5, Advanced ACI Management and Operations	10
Total	164

QUESTION NO: 1

An engineer must restrict the prefixes that are imported from an L3Out into a Cisco ACI VRF. Which two actions are required? (Choose two.)

A. Enable Import Route Control Enforcement for the VRF.

B. Associate an import route-control profile with the L3Out.

C. Enable ARP flooding in every bridge domain.

D. Configure a static route on each compute leaf switch.

E. Set the bridge domain to Unicast Routing disabled.

ANSWER: A B

Explanation:

Import Route Control Enforcement must be enabled for the VRF so that routes learned from an L3Out are denied unless they are explicitly permitted by route-control policy. Without import route control enforcement, learned external routes are imported according to the default behavior and a route-control policy cannot provide the required default-deny restriction.

A route-control profile must also be associated with the L3Out and include a match rule that permits the required prefixes. The route-control profile provides the prefix-based policy that determines which received routes can be installed in the tenant VRF. Together, VRF enforcement and the L3Out route-control profile provide controlled route import. See the Cisco ACI Layer 3 Configuration Guide.

QUESTION NO: 2

Which action must be taken to connect two nonsilient servers in active/standby mode to two ACI leaves?

A. Implement a virtual PortChannel on ACI leaves

B. Enable ARP flooding for the bridge domain if the MAC address of the active interface changes after NIC failover

C. Configure ARP flooding for the bridge domain if the MAC address of the active interface stays identicalafter a NIC failover

D. Disable dataplane learning to avoid endpoint flapping

ANSWER: B

Explanation:

Enable ARP flooding for the bridge domain if the MAC address of the active interface changes after NIC failover. With active/standby NIC teaming, each server NIC is independently attached to a different leaf; this is not a port-channel design. After a failover in which the newly active NIC uses a different MAC address, the server IP address can remain associated in the fabric endpoint database with the former MAC address and attachment location. Enabling ARP flooding causes ARP requests in that bridge domain to be flooded rather than answered exclusively through the ACI proxy-ARP endpoint cache. The newly active interface can therefore receive the ARP request and reply with its current MAC address. That reply enables the fabric to learn the updated IP-to-MAC relationship and endpoint location, restoring reachability after the failover. This behavior is particularly important for endpoint mobility scenarios where an IP address persists while the active interface identity changes. Cisco documents ARP flooding as a bridge-domain forwarding setting and describes its role alongside endpoint and ARP learning behavior in the ACI Layer 2 configuration documentation. See the Cisco APIC Layer 2 Configuration Guide and the Cisco APIC support documentation.

QUESTION NO: 3

Refer to the exhibit.

A company decided to decrease its routing footprint and remove RT-2 and RT-3 devices from its data center. Because of that, the exit point must be created from all the tenants by using the common tenant. Which two configuration tasks must be completed to meet these requirements? {Choose two.)

A. Move subnets from all the bridge domains to the EPG level and mark them with flag Shared between VRFs.

B. Change contract Ctr-3 scope to Global, consume it by all EPGs. and flag all subnets with flag Shared between VRFs.

C. Update the L3Qut ExtEPG subnet in the common tenant with flag Shared Route Control Subnet and Aggregate Shared Routes.

D. Mark all subnets with flag Shared between VRFs and attach contract Ctr-3 as a provider to all the EPGs.

E. Export contract Clr-2 into (he tenant TN-1 and attach it as a consumer to all the EPGs in the tenant TN-1.

ANSWER: C D

Explanation:

Using a shared exit point in the common tenant requires both route sharing and policy-based communication between the tenant VRFs and the external network. “Update the L3Qut ExtEPG subnet in the common tenant with flag Shared Route Control Subnet and Aggregate Shared Routes.” enables the common-tenant L3Out external EPG prefix to participate in shared route control. The Shared Route Control Subnet setting makes the external prefix eligible for sharing across VRFs, while Aggregate Shared Routes permits APIC to advertise the appropriate aggregate shared routes through that L3Out. This establishes the common tenant’s L3Out as the consolidated north-south routing path after the dedicated routers are removed.

“Mark all subnets with flag Shared between VRFs and attach contract Ctr-3 as a provider to all the EPGs.” completes the tenant-side requirements. Marking the relevant internal subnets as shared between VRFs allows those prefixes to be exported for inter-VRF route sharing. Providing the contract from the application EPGs supplies the required ACI policy relationship for traffic using the shared external connectivity. Together, route sharing and contract policy allow workloads in multiple tenants to use the common-tenant exit point while preserving ACI’s VRF and contract-based controls. Cisco documents these shared-service and route-control concepts in the Cisco APIC Layer 3 Networking Configuration Guide.

QUESTION NO: 4

An engineer deploys a bare-metal server by using an EPG static path binding on a Cisco ACI leaf access interface. Which two parameters must be configured in the static path binding? (Choose two.)

A. leaf interface or interface policy group path

B. VLAN encapsulation

C. bridge domain multicast address

D. COOP group address

E. BGP EVPN router ID

ANSWER: A B

Explanation:

The static path binding must identify the **leaf interface or interface policy group path** where the endpoint is attached. This path associates the EPG with the physical access connectivity in the fabric.

The binding must also specify a **VLAN encapsulation**. The VLAN encapsulation maps traffic received on the server-facing interface into the EPG. The encapsulation must be available through the VLAN pool associated with the physical domain used by the EPG. Cisco documents EPG static path bindings in the Cisco APIC Access Configuration Guide.

QUESTION NO: 5

An engineer implements the Cisco ACI solution and needs to confirm that the leaf switch correctly learns the remote endpoint MAC. Which information must be included in the VXLAN packet that is received from the Cisco ACI spine for the Cisco ACI leaf to learn MAC as a remote endpoint?

A. VRF

B. application profile

C. bridge domain

D. EPG

ANSWER: C

Explanation:

The correct information is **bridge domain**. Cisco ACI uses VXLAN encapsulation in its fabric data plane, and the VXLAN Network Identifier identifies the Layer 2 forwarding domain associated with the traffic. In ACI, this forwarding context is derived from the bridge domain, allowing an egress leaf to place the received source MAC address into the appropriate bridge-domain endpoint table as a remote endpoint.

When traffic arrives through the fabric from a spine, the receiving leaf uses the encapsulation context together with the packet source information to learn where that endpoint is reachable. The bridge-domain context is essential because MAC addresses are learned and resolved within a Layer 2 domain; the same MAC address could potentially be relevant in separate tenant or bridge-domain contexts. This enables correct endpoint lookup and subsequent VXLAN forwarding toward the remote leaf that owns the endpoint.

Cisco describes the ACI fabric as using VXLAN encapsulation for forwarding between leaf switches and explains that bridge domains provide the Layer 2 forwarding scope. See the Cisco ACI VXLAN white paper and the Cisco ACI Design Guide.

QUESTION NO: 6

Which approach does Cisco ACI use to achieve multidestination packet forwarding between leaf switches in the same fabric?

A. Map VXLAN VTEP to the multicast group

B. Map VXLAN to PIM-SM protocol

C. Map VXLAN VNI to the multicast group

D. Map VXLAN to PIM-DM protocol

ANSWER: C

Explanation:

Cisco ACI achieves multidestination forwarding for broadcast, unknown-unicast, and multicast traffic by associating a VXLAN Network Identifier (VNI) with an underlay multicast group. In ACI, the VNI identifies the logical forwarding domain carried in the VXLAN overlay, commonly corresponding to a bridge domain. When a leaf switch receives multidestination traffic for that VNI, it VXLAN-encapsulates the packet and sends it to the multicast group mapped to that VNI. The routed ACI underlay then replicates and delivers the packet only to leaf switches that have endpoints or interfaces participating in the associated logical domain. This approach provides efficient fabric-wide BUM replication without requiring the ingress leaf to create individual unicast copies for every remote leaf. APIC manages the multicast addressing and the VNI-to-group association as part of fabric policy and forwarding-state programming. Cisco documents that ACI uses multicast in the IP fabric for VXLAN BUM traffic and maps the overlay forwarding domain/VNID to multicast forwarding behavior. See the Cisco ACI Fundamentals Guide and the Cisco APIC Layer 2 Networking Configuration Guide.

QUESTION NO: 7

What are two characteristics of Cisco ACI interaction with MSTP? (Choose two.)

A. A Cisco ACI leaf flushes all local endpoints in the received EPG when an MSTP TCN frame is received.

B. A static path binding is required for native VLAN to carry the MST BPDUs in existing EPGs.

C. Mis-cabling protocol is used to snoop the MSTP TCN packets.

D. A dedicated EPG must be created for the native VLAN ID with static bindings toward the MSTP-enabled switches.

E. The PVST simulation must be configured on external network switches to support MSTP.

ANSWER: C D

Explanation:

Cisco ACI does not participate in MSTP as an MST bridge; instead, it transparently handles MST BPDUs across the fabric while providing endpoint-table handling for topology changes. **Mis-cabling protocol is used to snoop the MSTP TCN packets.** MCP is the ACI mechanism that monitors the relevant topology-change notifications, allowing leaf switches to react appropriately to an external MSTP topology change and maintain accurate endpoint reachability information.

A dedicated EPG must be created for the native VLAN ID with static bindings toward the MSTP-enabled switches. MSTP BPDUs are sent untagged on Ethernet trunks and therefore use the native VLAN. ACI requires an EPG representing that native VLAN and static path bindings on the interfaces connected to the MSTP domain so the fabric can carry and process the MSTP control traffic correctly. This design ensures that the native VLAN is explicitly represented in policy and consistently deployed on all relevant leaf-facing paths. Cisco documents ACI Layer 2 and spanning-tree interoperability behavior in its Cisco APIC configuration guides; additional product documentation is available from the Cisco APIC support documentation page.

QUESTION NO: 8

An engineer configures an OSPF-based Cisco ACI L3Out connection to an external router. Which two objects are required to configure the OSPF adjacency on the border leaf switches? (Choose two.)

A. logical node profile

B. logical interface profile

C. VMM domain

D. bridge domain flood domain

E. route reflector policy

ANSWER: A B

Explanation:

A **logical node profile** is required to identify the border leaf nodes that participate in the L3Out. The node profile contains the logical router information, including the router ID and routing configuration applicable to those nodes.

A **logical interface profile** is also required to define the routed interfaces toward the external router. The interface profile contains the interface-specific configuration, including the OSPF interface policy used to establish the OSPF adjacency. Cisco documents L3Out logical node and interface profiles in the Cisco APIC Layer 3 Configuration Guide.

QUESTION NO: 9

An engineer configures a Cisco ACI Multi-Pod for disaster recovery. Which action should be taken for the new nodes to be discoverable by the existing Cisco APICs?

A. Enable subinterfaces with dot1q tagging on all links between the IPN routers.

B. Configure IGMPv3 on the interfaces of IPN routers that face the Cisco ACI spine.

C. Enable DHCP relay on all links that are connected to Cisco ACI spines on IPN devices.

D. Configure BGP as the underlay protocol in IPN.

ANSWER: C

Explanation:

Configure DHCP relay on all links that are connected to Cisco ACI spines on IPN devices. In a Multi-Pod design, APIC controllers are typically deployed in a single pod, while fabric nodes in remote pods still need to perform initial discovery and receive their required bootstrap information from those APICs. During discovery, a newly connected leaf or spine uses DHCP-based communication over the infrastructure network. Because the IPN is a Layer 3 boundary between pods, it does not forward the node's DHCP broadcast traffic by default.

DHCP relay on the IPN interfaces facing the ACI spines relays the discovery-related DHCP requests toward the APIC DHCP service and returns the appropriate responses to the requesting node. This enables the APIC cluster to identify, register, and provision nodes located in pods that do not host APICs. The relay configuration must be present for the spine-facing connectivity used by each relevant pod so that discovery remains available across the IPN during onboarding and recovery scenarios. Cisco documents DHCP relay as an IPN requirement for discovering fabric nodes in remote Multi-Pod sites. See the Cisco ACI Multi-Pod Deployment Guide and the Cisco ACI Multi-Pod Configuration Guide.

QUESTION NO: 10

A Cisco ACI fabric Is configured with these settings:

• A single VRF exists that contains 100 EPGs.

• EPGs 1 to 50 must be permitted to communicate.

• EPGs 51 to 100 must NOT be allowed to communicate with EPGs 1 to 50.

• EPG 51 must be allowed to communicate with EPG 10.

Which two actions must be taken to accomplish these requirements? (Choose two.)

A. Create a vzAny contract.

B. Create a standard contract between EPG 51 and EPG 10.

C. Select the Intra-EPG Isolation checkbox.

D. Set policy control enforcement to Unenforced.

E. Enable the Preferred Group option on EPGs 1 to 50

ANSWER: B E

Explanation:

Enabling the Preferred Group option on EPGs 1 to 50 makes those EPGs preferred-group members. Within a VRF configured to use preferred groups, member EPGs can communicate with one another without requiring individual contracts between every pair of EPGs. This efficiently satisfies the requirement for unrestricted communication among the first 50 EPGs while retaining the default ACI policy model for nonmembers: inter-EPG communication is denied unless an applicable contract permits it.

Creating a standard contract between EPG 51 and EPG 10 provides the required narrowly scoped exception. The contract is applied with EPG 51 and EPG 10 in the appropriate consumer/provider roles and includes filters for the traffic that must be allowed. Because EPG 51 is not a preferred-group member, it does not obtain broad access to the preferred-group EPGs merely through the preferred-group configuration. Its communication is allowed only through the explicitly deployed contract, preserving separation between EPGs 51 to 100 and EPGs 1 to 50 except for the required EPG 51-to-EPG 10 policy relationship. Cisco documents preferred groups and contracts in its APIC Virtual Networking Configuration Guide and ACI policy-model overview.

QUESTION NO: 11

An engineer must connect bare-metal servers to Cisco ACI leaf access ports by using a physical domain. Which two policy associations are required before the physical domain can be selected in an EPG static path binding? (Choose two.)

A. Associate the physical domain with a VLAN pool

B. Associate the physical domain with an attachable access entity profile

C. Associate the physical domain with a route reflector policy

D. Associate the physical domain with a bridge domain subnet

E. Associate the physical domain with an L3Out

ANSWER: A B

Explanation:

A physical domain must be associated with a **VLAN pool** so that valid VLAN encapsulations are available for EPG static path bindings. The VLAN pool defines the VLAN ranges that can be used by the domain.

The physical domain must also be associated with an **attachable access entity profile** (AAEP). The AAEP connects the domain policy to the access interface policy groups assigned to leaf interfaces. This relationship permits endpoints on those interfaces to use the physical domain and its EPG VLAN encapsulations. Cisco describes physical domains, VLAN pools, and AAEP associations in the Cisco APIC Access Configuration Guide.

QUESTION NO: 12

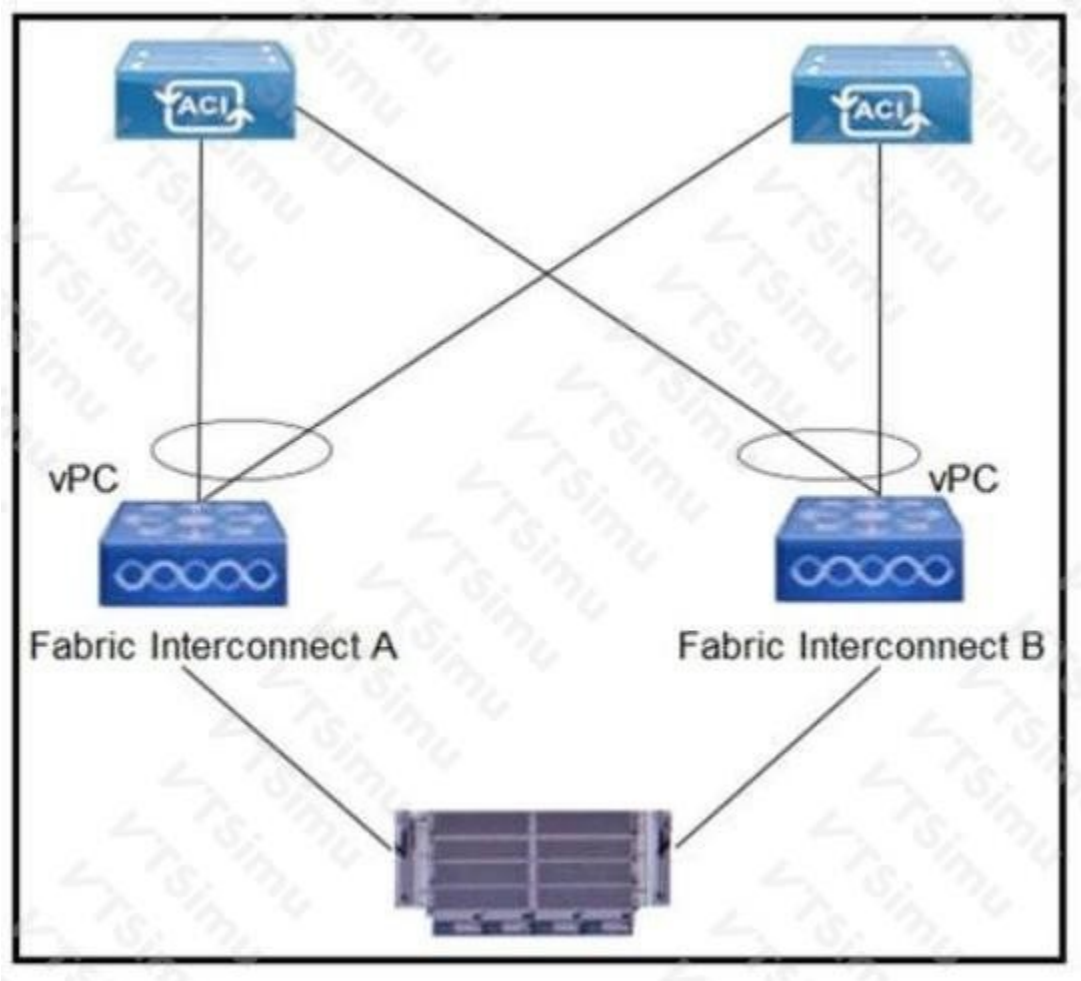

Refer to the exhibit. Which configuration mode must be selected for the VMM vSwitch Port Channel policy to avoid MAC flapping on the Cisco UCS FIs and Cisco ACI leaf switches?

A. LACP Passive

B. MAC Pinning

C. LACP Active

D. Static Channel – Mode On

ANSWER: B

Explanation:

MAC Pinning is the required VMM vSwitch Port Channel policy mode in this Cisco UCS Fabric Interconnect and Cisco ACI design. MAC pinning assigns a virtual machine or host MAC address to one specific physical uplink for forwarding. That consistent forwarding path is important when the ESXi host connects through Cisco UCS Fabric Interconnects, because the upstream Cisco ACI leaf switch must learn a given source MAC address on a stable interface. By ensuring that a MAC is not alternated between available host uplinks during normal traffic forwarding, MAC pinning prevents the leaf switches and Fabric Interconnects from continually relearning the same MAC on different ports, which is observed as MAC flapping. The policy is intended for vSwitch connectivity where a static, per-MAC uplink association is needed instead of link-aggregation negotiation. Cisco's APIC virtualization documentation describes MAC pinning as a vSwitch port-channel policy mode used to pin MAC addresses to physical NICs, and Cisco UCS integration guidance emphasizes selecting the appropriate host-uplink behavior when connecting UCS workloads to the ACI fabric. See the Cisco APIC Virtualization Configuration Guide and Cisco's Application Centric Infrastructure documentation resources.

QUESTION NO: 13

Cisco ACI fabric is migrated from network-centric mode to application-centric mode. What is possible in an application-centric design?

A. remapping of VLAN pools

B. migration of bare metal servers to virtual machines

C. creation of additional EPGs under one bridge domain

D. Cisco ACI VMM integration with hypervisors

ANSWER: C

Explanation:

Creation of additional EPGs under one bridge domain is possible because an Application Centric Infrastructure bridge domain defines the Layer 2 forwarding and flooding scope, while endpoint groups provide the logical segmentation used to represent application tiers, functions, or security groups. Multiple EPGs can therefore share the same bridge domain and its subnet while remaining distinct policy entities. This lets an administrator evolve a network-centric design, where endpoint grouping may initially mirror VLAN and subnet constructs, into an application-centric model that separates endpoints according to application requirements. Contracts can then be applied between the EPGs to explicitly define permitted communications, even though those EPGs use the same underlying bridge domain. This model supports incremental migration: the existing common Layer 2/Layer 3 addressing scope can be retained while policy granularity is increased by introducing additional EPGs and contracts. Cisco describes EPGs as logical groups of endpoints that are commonly placed within a bridge domain, and notes that policy is applied between EPGs through contracts. See the Cisco APIC Basic Configuration Guide and the Cisco ACI overview.

QUESTION NO: 14

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 multidestination traffic.

ANSWER: D

Explanation:

FTAG trees in Cisco ACI are used to load balance multidestination traffic. An FTAG is a forwarding identifier assigned to a multicast distribution tree in the ACI fabric. The APIC computes multiple loop-free trees across the spine-and-leaf topology and assigns an FTAG to each tree. For bridge-domain traffic that must be delivered to multiple destinations—such as broadcast, unknown unicast, and multicast traffic—the ingress leaf selects an FTAG/tree. The selected FTAG is carried in the VXLAN encapsulation so that each spine and leaf can forward the packet along the intended multidestination tree without creating loops or requiring every device to independently calculate a flood path. Using multiple FTAG trees distributes this traffic across available fabric links, improving path utilization while retaining deterministic replication and forwarding behavior. The specific tree choice is typically based on a hash of packet fields, allowing different flows to use different trees. Cisco describes FTAGs as part of the fabric's multicast forwarding mechanism and explains that they provide load balancing for multidestination traffic. See the Cisco APIC Layer 2 Configuration Guide and the Cisco ACI technical documentation.

QUESTION NO: 15

An engineer created a Cisco ACI environment in which multiple tenants reuse the same contract. The requirement is to prevent inter-tenant communication. Which action meets this requirement?

A. Create the contract in the user tenant with the scope set to VRF and exported to other tenants

B. Create the contract in the common tenant with the scope set to Tenant

C. Create the contract in the user tenant with the scope set to Global and exported to other tenants

D. Create the contract in the common tenant with the scope set to Global

ANSWER: B

Explanation:

Create the contract in the common tenant with the scope set to Tenant. The common tenant is intended to hold shared policy objects, including contracts, so that multiple user tenants can reference one centrally managed contract definition instead of maintaining duplicate contracts. Setting the contract scope to *Tenant* preserves tenant isolation: the contract can be used by endpoint groups within an individual tenant, but it does not authorize a provider endpoint group in one tenant to communicate with a consumer endpoint group in another tenant.

This design meets both requirements simultaneously. The contract is reusable because it is defined in the common tenant and can be referenced by other tenants, while the tenant-scoped enforcement boundary ensures that policy matching and communication remain limited to objects belonging to the same tenant. Consequently, each tenant can apply the shared contract to its own application policy without creating inter-tenant connectivity. Cisco describes contracts as policy objects that define permitted communications between endpoint groups, and contract scope determines the boundary across which that policy can apply. See the Cisco APIC REST API Configuration Guide and the Cisco APIC Fundamentals Configuration Guide.

QUESTION NO: 16 - (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.

Answer Area

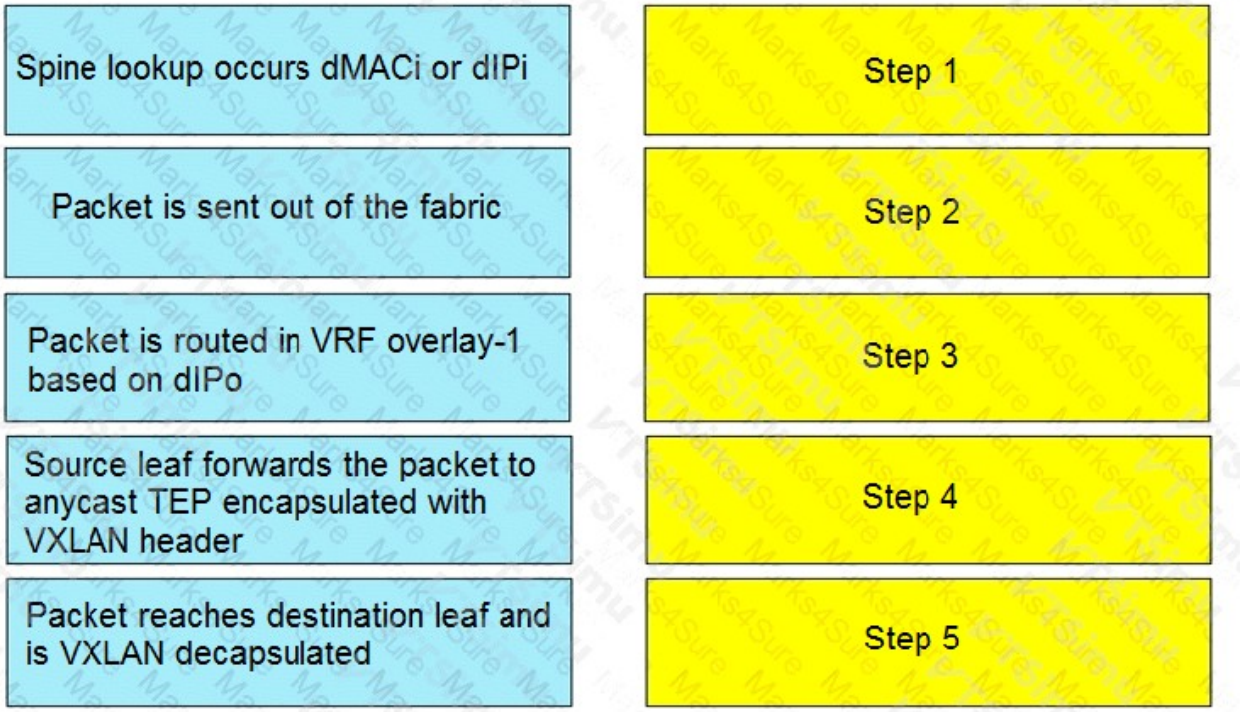

ANSWER:

Explanation:

With spine proxy enabled, an ingress leaf that does not have a local endpoint-table result for the destination does not immediately know which remote leaf should receive the frame. It first forwards the traffic through the VXLAN fabric to the anycast TEP used by the spine-proxy function. Therefore, the forwarding process begins when the source leaf forwards the packet to the anycast TEP encapsulated with a VXLAN header. The original endpoint-facing packet information remains available as inner-header information for the proxy lookup.

At the spine-proxy function, the fabric performs the destination lookup using the inner destination MAC or inner destination IP. This lookup identifies the correct destination location and allows the traffic to be directed to the appropriate leaf rather than treated as broadly flooded unknown traffic. Cisco ACI uses its endpoint-directory and proxy mechanisms to provide this location resolution while retaining distributed forwarding at the leaves. Cisco provides an overview of the fabric's endpoint and forwarding architecture in the Cisco Application Centric Infrastructure documentation.

Once the lookup identifies the destination leaf, the packet is sent through the fabric to that leaf. On arrival, the destination leaf removes the VXLAN encapsulation, restoring the packet for local forwarding treatment. The leaf then performs the applicable VRF overlay forwarding decision based on the outer destination IP information. Thus, the correct ordered mapping is: source leaf forwards the packet to the anycast TEP encapsulated with VXLAN header; spine lookup occurs using dMACi or dIPi; packet is sent out of the fabric; packet reaches the destination leaf and is VXLAN decapsulated; packet is routed in VRF overlay-1 based on dIPo.