

Application Centric Infrastructure Support Representative

Cisco 500-601

Version Demo

Total Demo Questions: 8

Total Premium Questions: 83

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Cisco ACI Overview	4
Topic 2, Cisco ACI Fabric Infrastructure	14
Topic 3, Cisco ACI Fabric Initialization	2
Topic 4, Cisco ACI Fabric Configuration	48
Topic 5, Cisco ACI Troubleshooting	15
Total	83

QUESTION NO: 1

On the Cisco Nexus 93128 chassis, which port range is capable of operating at 40 GB?

- A. 1/97-1/100
- B. 1/97-1/104
- C. 1/97-1/108
- D. 1/100-1/108

ANSWER: B

Explanation:

1/97-1/104 is correct for the Cisco Nexus 93128TX platform. This switch provides 96 fixed 1/10-Gigabit Ethernet downlink interfaces, numbered Ethernet 1/1 through Ethernet 1/96, followed by eight uplink interfaces numbered Ethernet 1/97 through Ethernet 1/104. The uplink ports use QSFP+ optics or cables and support operation at 40 Gbps. This port layout is important during initial provisioning and troubleshooting because it distinguishes the high-speed uplink block from the 10-Gbps host-facing port block. In an ACI deployment, these interfaces can be selected for connections requiring 40-Gbps bandwidth, subject to the configured interface policy, transceiver or cable compatibility, and the capabilities of the connected peer. Cisco documents the 93128TX physical interface arrangement and uplink capabilities in its hardware documentation, and the Nexus 9000 Series data sheet identifies the platform's 40-Gbps QSFP+ uplink support. See the [Cisco Nexus 9000 Series Switches Data Sheet](#) and the [Cisco Nexus 93128TX Hardware Installation Guide](#).

QUESTION NO: 2

Which two settings are required when creating a VMware VMM domain in Cisco ACI? (Choose two.)

- A. vCenter controller connection
- B. dynamic VLAN pool
- C. external EPG
- D. L3Out
- E. PIM-enabled bridge domain

ANSWER: A B

Explanation:

A **vCenter controller connection** and a **dynamic VLAN pool** are required when configuring a VMware VMM domain. The vCenter controller connection supplies the address, credentials, and data-center information that APIC uses to integrate with VMware vCenter. APIC uses this connection to discover virtualization inventory and manage the associated distributed virtual switch configuration.

The dynamic VLAN pool supplies VLAN IDs that APIC allocates to EPG port groups deployed through the VMM domain. APIC assigns these VLAN IDs dynamically as EPGs are associated to the VMM domain. An external EPG and an L3Out are used for routed external connectivity and are not required to establish the VMM domain. See the [Cisco APIC Virtual Networking Configuration Guide](#) for VMM domain configuration requirements.

QUESTION NO: 3

Which two statements about troubleshooting common ARP forwarding issues ma Cisco ACI Fabric are true? (Choose two.)

- A. If IP Manager on the spine rejects a tenant VRF or route configuration, use `show ip route vrf` to verify the tenant VRF routing configuration. Verify the BD SVI IP address with `vsh_lc -c "show platform internal ns forwarding router-ip"`.

B. The `tcpdump` command can be used on all Cisco ACI switches to capture and verify tenant EPGdata plan traffic.

C. The VLAN membership for an interface on the Broadcom ASIC can be verified with the `show hardware vlans programmed` command. VLAN membership of an interface can be verified with `vsh_lc -c "show system internal eltmc info interface ethernet x/x"`.

ANSWER: A C

Explanation:

When IP Manager on a spine rejects a tenant VRF or route configuration, `show ip route vrf <vrf-name>` provides a direct view of the route entries associated with that VRF. This validates that the VRF routing information has been accepted and installed where expected in the fabric control-plane context. The BD subnet gateway address is also a key ARP-forwarding dependency: `vsh_lc -c "show platform internal ns forwarding router-ip"` checks the router IP information programmed in the leaf forwarding namespace, helping confirm that the bridge-domain SVI gateway is present for local endpoint routing and ARP processing.

VLAN programming must also be verified at the hardware and internal software layers when investigating forwarding behavior. `show hardware vlans programmed` confirms that VLAN membership has been installed in the Broadcom ASIC. The ELTMC internal interface command provides complementary visibility into the interface's VLAN membership and associated internal programming. Together, these checks establish that the access interface, VLAN encapsulation, and hardware forwarding entries are aligned, which is essential for frames—including ARP requests and replies—to reach the correct bridge-domain forwarding context. Cisco's ACI troubleshooting and command references provide the relevant platform-level validation workflow: [Cisco ACI Troubleshooting Guide](#) and [Cisco APIC Command Reference](#).

QUESTION NO: 4

Which two interface types can be used to configure interface profiles for the border leaf ports of a Layer 3 external routed network? (Choose two)

- A. Layer 3 subinterfaces
- B. Layer 2 ports with STP-enabled interfaces
- C. Layer 4 - Layer 7 route peering interfaces
- D. Layer 3 physical interfaces
- E. Layer 3 extended GRE tunnel interfaces

ANSWER: A D

Explanation:

Layer 3 subinterfaces and Layer 3 physical interfaces are supported interface types for connecting an ACI border leaf to a Layer 3 external routed network (an L3Out). A Layer 3 physical interface is used when the routed connection uses the complete physical Ethernet port as a routed link. The interface is associated with the L3Out logical interface profile and is configured with the appropriate addressing and routing settings, such as BGP, OSPF, or static routing as required by the external connectivity design.

Layer 3 subinterfaces are used when the routed external connection must be carried over an IEEE 802.1Q VLAN tag on a physical port. This permits separate logical routed links to share a physical interface while retaining independent encapsulation and Layer 3 configuration. In APIC, both routed physical interfaces and routed subinterfaces are configured through the L3Out logical interface profile for the selected border leaf interface profile. Cisco documents routed interfaces and subinterfaces as L3Out interface deployment choices; the subinterface configuration includes VLAN encapsulation in addition to its Layer 3 parameters. See Cisco's [APIC Layer 3 Configuration Guide](#) and the [APIC Basic Configuration Guide](#).

QUESTION NO: 5

By default, how is traffic treated between two endpoints in the same EPG?

- A. Traffic is permitted.
- B. Traffic is denied until a contract is configured.
- C. Traffic is sent automatically to the default VRF.
- D. Traffic is permitted only when both endpoints use the same leaf switch.

ANSWER: A

Explanation:

Traffic is permitted. Cisco ACI permits intra-EPG communication by default. Endpoints assigned to the same EPG share the same security classification, and traffic between them does not require a contract unless intra-EPG isolation is explicitly enabled.

Intra-EPG isolation can be configured when endpoints in the same EPG must be prevented from communicating directly. Cisco documents EPG communication and contract behavior in the [Cisco ACI Fundamentals Guide](#).

QUESTION NO: 6

Where is IP multicast routing enabled for a tenant routing domain in Cisco ACI?

- A. VRF
- B. application profile
- C. contract subject
- D. physical domain

ANSWER: A

Explanation:

VRF is correct. IP multicast routing is enabled on the tenant VRF that is used by the multicast-enabled bridge domains and routed interfaces. The VRF provides the Layer 3 routing context in which multicast routes, rendezvous point information, and multicast control-plane operation are maintained.

Additional multicast configuration can be required on bridge domains, L3Out interfaces, and external multicast routers according to the multicast design. However, enabling multicast routing for the tenant routing domain is a VRF-level setting. Cisco documents the multicast configuration workflow in the [Cisco APIC Layer 3 Configuration Guide](#).

QUESTION NO: 7

Which two statements about troubleshooting common ARP forwarding issues on Cisco ACI Fabric are true? (Choose two.)

- A. The BD SVI IP address can be verified with the vsh_lc show platform internal ns forwarding router-ip command.
- B. The tcpdump command can be used on all Cisco ACI switches to capture and verify tenant EPG data plan traffic.
- C. The VLAN membership for an interface on the Broadcom ASIC can be verified with the show hardware vlans programmed command.
- D. VLAN membership of an interface can be verified with the vsh_lc show system internal eltmc info interface x/x command.

ANSWER: A D

Explanation:

The BD SVI IP address can be verified with the `vsh_lc show platform internal ns forwarding router-ip` command. In ACI, a bridge domain may have an L3 gateway IP address, and ARP forwarding depends on the relevant router-IP information being programmed on the leaf that provides the anycast gateway function. Viewing the platform namespace forwarding router-IP entries helps confirm that this critical L3 forwarding information is present in the local leaf programming.

VLAN membership of an interface can be verified with the `vsh_lc show system internal eltmc info interface x/x` command. This internal ELTMC view exposes interface-level VLAN membership details, allowing an engineer to validate that the VLAN associated with an encapsulation has been applied to the expected physical interface. Confirming the programmed interface membership is a useful ARP troubleshooting step because ARP requests and replies must be carried over the correct local VLAN before endpoint learning and forwarding can succeed. Cisco documents the relevant NX-OS and ACI internal command environment in its [Cisco APIC and NX-OS Command Reference](#) and provides ARP and endpoint troubleshooting context in the [Cisco APIC Troubleshooting Guide](#).

QUESTION NO: 8

Which action is needed on a Cisco UCS to make ACI dynamic discovery work?

- A. Enable VLAN grouping on the fabric interconnect uplinks.
- B. Enable the Cisco Discovery Policy on the Cisco UCSM vNICs
- C. Forward the infrastructure VLAN toward the ESX.
- D. Use MTU 9000.

ANSWER: C

Explanation:

Forward the infrastructure VLAN toward the ESX. is required because ACI dynamic VMM discovery depends on infrastructure control-plane traffic reaching the ESXi host through the Cisco UCS Fabric Interconnect and its vNICs. The ACI infrastructure VLAN carries the discovery communication that enables the fabric to associate a hypervisor and its virtual-switch information with the physical leaf-facing connectivity. Therefore, the UCS configuration must permit this VLAN on the relevant path toward ESXi rather than limiting the path to only tenant or management VLANs. This allows the necessary discovery traffic to traverse the Fabric Interconnect and reach the host, so APIC can dynamically learn the VMM attachment information and maintain accurate endpoint and virtual-network visibility. The requirement applies independently of jumbo-frame settings; the essential configuration is providing end-to-end forwarding for the ACI infrastructure VLAN on the UCS-to-ESXi connection. Cisco documents the infrastructure VLAN as the VLAN used for ACI infrastructure communication and describes VMM integration and host discovery requirements in the [Cisco APIC Virtualization Guide](#). Additional UCS networking and vNIC configuration context is available in the [Cisco UCS Manager Configuration Guide](#).