

Aruba Certified Campus Access Mobility Expert Written Exam

HP HPE7-A07

Version Demo

Total Demo Questions: 16

Total Premium Questions: 166

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

QUESTION NO: 1

A pair of AOS-CX aggregation switches operates as a VSX pair. An access switch must use both uplinks simultaneously while maintaining connectivity if either aggregation switch fails. The access switch supports LACP but does not support VSX.

Which configuration actions are required to implement this design? (Select two.)

- A. Configure both access-switch uplinks as members of one LACP port channel.
- B. Configure the corresponding aggregation-switch LAGs as the same VSX multi-chassis LAG.
- C. Create two independent LACP port channels on the access switch.
- D. Configure an OSPF adjacency across each access-switch uplink.
- E. Configure VRRP between the access switch and each aggregation switch.

ANSWER: A B

Explanation:

The access switch must place its two physical uplinks in one LACP port channel. This lets the downstream switch treat the two physical links as members of one logical aggregated connection.

The corresponding port-channel interfaces on the VSX peers must be configured as the same multi-chassis LAG. This allows the independent VSX control planes to present the two aggregation-switch links as one logical LAG to the access switch. Creating separate LAGs on the access switch would not provide the intended active-active MC-LAG connection. A routed OSPF adjacency or VRRP configuration is not required for Layer 2 MC-LAG operation.

QUESTION NO: 2

A customer is experiencing authentication failures when clients connect to a new EAP-TLS SSID.

Based on the logs and packet capture above, what is the cause of the failure?

- A. The client cannot validate the RADIUS server's certificate
- B. The MTU in the path between the AP and HPE Aruba Networking ClearPass is too small
- C. HPE Aruba Networking ClearPass cannot validate the user's certificate
- D. The access point doesn't have the correct root CA certificate installed

ANSWER: B

Explanation:

The MTU in the path between the AP and HPE Aruba Networking ClearPass is too small. EAP-TLS transports certificate chains and related TLS messages within EAP, which can make the RADIUS/EAP exchange significantly larger than a typical password-based 802.1X transaction. The reported ClearPass timeout and the AP's server-timeout indication show that an expected EAP/RADIUS message was not received in time, rather than that certificate validation completed and failed. When a packet carrying a large EAP-TLS fragment exceeds an effective Layer-3 MTU, it can be fragmented; if a network device, firewall, VPN tunnel, or security policy drops IP fragments, the RADIUS conversation cannot continue. The NAS and ClearPass then each eventually report a timeout or an incomplete EAP transaction. The observed Framed-MTU value is especially relevant because it indicates that MTU handling is part of the exchange, but the complete AP-to-ClearPass path must still support the packet sizes involved. Validate the effective MTU across every hop, including tunnels and firewalls, and ensure RADIUS UDP fragments are not filtered. Aruba's ClearPass documentation is available through the [ClearPass Technical Documentation portal](#); ArubaOS authentication guidance is available in the [ArubaOS 802.1X documentation](#).

QUESTION NO: 3

An engineer is extending selected user VLANs between buildings through an AOS-CX EVPN-VXLAN fabric. The physical network between the buildings is routed, and the design must continue forwarding if one routed underlay path fails while avoiding flood-based learning of all endpoint MAC addresses.

Which two functions are required? (Select two.)

- A. Provide routed IP reachability between VTEP loopback addresses.
- B. Use BGP EVPN to distribute endpoint reachability information.
- C. Extend every user VLAN as an untagged VLAN across the routed underlay links.
- D. Configure VRRP on every VTEP loopback interface.
- E. Enable DHCP relay on each underlay transit interface.

ANSWER: A B

Explanation:

The routed underlay must provide IP reachability between the VTEP loopback addresses. VXLAN packets are encapsulated with the remote VTEP address as the destination, so resilient routed reachability between VTEPs is required for overlay traffic to traverse the fabric.

BGP EVPN distributes MAC/IP endpoint reachability information as a control plane. This allows VTEPs to learn remote endpoint locations through EVPN route advertisements instead of relying solely on data-plane flooding for host learning. A single shared Layer 2 VLAN across all fabric links is not required and would defeat the routed-underlay design. VRRP and DHCP relay can be used for other services but do not provide the required EVPN control plane.

QUESTION NO: 4

A customer is evaluating device profiles on a CX 6300 switch. The test device has the following attributes:

- MAC address = 81:cd:93:13:ab:31
- LLDP sys-desc = iotcontroller

The test device is being assigned to the "iot-dev" role. However, the customer requires the "iot-prod" role be applied.

Given the configuration, what is causing the "iot-dev" role to be applied to the device'?

- A. The test device does not support CDP.
- B. The device-profile precedence order is not configured.
- C. An external RADIUS server is unreachable.
- D. The LLDP system description matches the lldp-group configuration.

ANSWER: D

Explanation:

The LLDP system description matches the lldp-group configuration. AOS-CX device profiles classify connected endpoints by evaluating configured match groups against attributes learned on the access port, including the device MAC address and LLDP information. In this case, the endpoint advertises an LLDP system description of `iotcontroller`. When that value satisfies the LLDP system-description match defined in the `lldp-group`, the associated device profile is selected and its configured `iot-dev` role is applied. The result is driven by the successful identity match from the LLDP advertisement, rather than by an assumption about the endpoint type. To have the production role assigned, the relevant device-profile matching logic must be changed so that this device's attributes select a profile associated with `iot-prod`, while accounting

for the configured profile evaluation behavior. Aruba documents device-profile classification and the use of LLDP-based matching in the AOS-CX Security Guide: [AOS-CX Device Profiles documentation](#). Additional CX configuration guidance is available in the [Device Profile Configuration documentation](#).

QUESTION NO: 5

A customer uses ClearPass to send Change of Authorization requests to an AOS 10 mobility gateway cluster when an employee role changes. Initial 802.1X authentication succeeds, but ClearPass CoA requests are rejected by the gateway cluster.

Which configuration items must be verified for ClearPass-initiated CoA processing? (Select two.)

- A. Enable Dynamic Authorization CoA in the gateway cluster configuration.
- B. Configure ClearPass as an authorized Dynamic Authorization client with the correct shared secret.
- C. Assign the client to a downloadable user role.
- D. Configure the employee SSID for WPA3-Enterprise.
- E. Place ClearPass in the same client VLAN as the wireless user.

ANSWER: A B

Explanation:

Dynamic Authorization CoA must be enabled in the gateway cluster configuration so that the cluster accepts RADIUS Dynamic Authorization requests. ClearPass must also be configured as an authorized Dynamic Authorization client with the correct source address and shared secret. The gateway validates the sender and the RADIUS authenticator before accepting a CoA request.

The client VLAN does not determine whether the gateway accepts CoA, and a downloadable user role is not required for Dynamic Authorization. WPA3-Enterprise controls wireless authentication and encryption but does not enable RADIUS CoA processing.

QUESTION NO: 6



What is the expected behavior for ARP traffic sent from H1?

- A. A2 will send the ARP traffic out of ports 1/1/1 and 1/1/3.
- B. A2 will send the ARP traffic out of ports 1/1/1–1/1/4.
- C. A2 will drop the ARP traffic.
- D. A2 will flood the ARP traffic out of all interfaces.

ANSWER: B

Explanation:

A2 will send the ARP traffic out of ports 1/1/1–1/1/4. ARP resolution starts with an Ethernet broadcast because H1 does not yet know the destination MAC address associated with the requested IPv4 address. On receipt of that broadcast, A2 performs normal VLAN flooding: it transmits the frame through the other local forwarding ports that belong to VLAN 100, while not sending it back through the ingress interface. In the illustrated topology, those eligible egress interfaces are ports 1/1/1 through 1/1/4.

Group-Based Policy role-to-role rules govern classified IP communications between source and destination roles. An ARP request is not an IPv4 or IPv6 data packet and does not contain the destination IP traffic flow needed for a role-to-role policy decision. ARP must be able to reach VLAN members so that endpoints can resolve MAC addresses before subsequent IP communication is evaluated and enforced. Accordingly, the ARP broadcast is flooded within the VLAN; the later IP traffic is the traffic to which the Marketing-to-Sales, Marketing-to-Finance, and Marketing-to-HR policy outcomes apply.

For ARP's defined broadcast request behavior, see [RFC 826](#). Aruba's AOS-CX documentation portal, including security and policy configuration guidance, is available at [Aruba AOS-CX Technical Documentation](#).

QUESTION NO: 7

A campus topology uses VSX with a collapsed core topology. The customer added redundant SFP+ transceivers and reconfigured their mobility gateways from a single link to an aggregate Link. You are asked to verify the CLI output for the link aggregation configuration for one of the mobility gateway cluster members below.

What is a valid configuration?

- A)
- B)
- C)
- D)
- A. Option A
- B. Option B
- C. Option C
- D. Option D

ANSWER: A

Explanation:

The valid configuration is the one identified as *Option A*. For a mobility gateway connected redundantly to two members of an Aruba CX VSX pair, the switch-facing LAG must be configured as a multi-chassis LAG so that both physical member links belong to one logical aggregate across the VSX pair. LACP provides the required dynamic negotiation and monitoring of the aggregate, and active LACP operation with a fast LACPDU rate provides rapid member-link detection. The aggregate must also operate as a VLAN trunk when it carries gateway traffic for more than one VLAN. Allowing the required VLANs on the trunk and setting native VLAN 100 means untagged traffic is classified consistently while tagged client, management, and service VLANs can traverse the aggregate. This produces a resilient active-active uplink: loss of one transceiver, cable, or VSX peer link member does not require the gateway to move all traffic to a separate standby interface. Aruba AOS-CX documents that the `multi-chassis` LAG setting is used for VSX-connected downstream devices, while LACP manages aggregation membership and forwarding behavior. See the [Aruba AOS-CX Link Aggregation Guide](#) and the [Aruba AOS-CX VSX Guide](#).

QUESTION NO: 8

An OSPF router has learned a path to an external network by both an E1 and an E2 advertisement. Both routes have the same path cost. Which path will the router prefer?

- A. The router will prefer the E1 path.
- B. The router will use both paths equally utilizing ECMP.
- C. The router will prefer the E2 path.
- D. Both routes will be suppressed until the path conflict has been resolved.

ANSWER: A

Explanation:

The router will prefer the E1 path. OSPF defines Type 1 external routes (E1) and Type 2 external routes (E2) with an explicit route-selection preference: Type 1 external paths are considered preferable to Type 2 external paths. This preference applies before comparing the routes as equivalent candidates for forwarding, so matching path costs do not cause the router to load-balance between the E1 and E2 advertisements.

An E1 metric represents the complete cost to the destination: the external metric advertised by the ASBR plus the internal OSPF cost required to reach that ASBR. This makes the E1 route metric meaningful across the OSPF domain, particularly when multiple ASBRs can advertise reachability to the same external prefix. The OSPF specification describes Type 1 external metrics as comparable to OSPF internal metrics and states that Type 1 external paths are always shorter than Type 2 external paths. ArubaOS routing behavior follows standard OSPF route types and preference rules. See [RFC 2328, Section 16.4](#) and [RFC 2328, AS-external-LSA metric types](#).

QUESTION NO: 9

What directly affects the MCS used by wireless stations? (Select two.)

- A. SNR
- B. retry rate
- C. channel utilization
- D. number of connected clients
- E. frequency band

ANSWER: A B

Explanation:

SNR directly influences the modulation and coding scheme a wireless client can use because it represents the usable signal level relative to the noise floor at the receiver. Each MCS requires a minimum signal quality for frames to be decoded reliably. As SNR improves, rate adaptation can select more efficient modulation and coding combinations; when SNR falls, it selects a more robust MCS to maintain reliable communication.

retry rate is also a direct rate-adaptation input. Retries provide immediate evidence of whether frames sent at the current PHY rate are being delivered successfully. A rising retry rate indicates that the current transmission conditions are not reliably sustaining the selected MCS, so the client or access point rate-control algorithm can reduce the MCS. Conversely, consistently successful transmissions allow the algorithm to probe or sustain higher MCS values. This continuous feedback loop lets Wi-Fi devices balance PHY efficiency with frame-delivery reliability. Aruba's Adaptive Radio Management documentation describes its use of radio and client performance information for RF optimization, while the Wi-Fi Alliance's Wi-Fi 6 overview describes the PHY enhancements built around modulation and coding: [ArubaOS ARM overview](#) and [Wi-Fi Alliance: Wi-Fi CERTIFIED 6](#).

QUESTION NO: 10

You are troubleshooting a WLAN deployment with APs and gateways set up with an 802.1X tunneled SSID. End-users are complaining that they can't connect to the enterprise SSID. Which possible AP tunnel states could be the cause of the issue? (Select two.)

- A. SM_STATE_CONNECTING
- B. SM_STATE_SURVIVED
- C. SM_STATE_SURVIVING
- D. SM_STATE_CONNECTED
- E. SM_STATE_REKEYING

ANSWER: A C

Explanation:

SM_STATE_CONNECTING is a possible cause because it denotes that the AP tunnel is still being established rather than being in a usable operational state. For an SSID whose client traffic is tunneled to a gateway, the AP requires an operational tunnel before it can carry the client's authentication exchanges and subsequent user data to the gateway. A tunnel that remains in this transitional state can therefore prevent enterprise clients from completing connectivity.

SM_STATE_SURVIVING is also a possible cause because it indicates that the AP is operating in survival handling while attempting to maintain or recover tunnel service after a tunnel-related disruption. This is not the normal steady-state condition for a tunnel carrying centralized WLAN traffic. If the AP cannot restore a usable tunnel while in this state, 802.1X authentication and client forwarding for the tunneled enterprise SSID can be unavailable.

In a tunneled WLAN design, gateway tunnel availability is a prerequisite for the AP to deliver centrally processed client traffic. Aruba's documentation describes the role of AP-to-gateway tunneling and IPsec protection in forwarding WLAN traffic through the gateway: [ArubaOS IPsec documentation](#) and [Aruba Central AP tunnel documentation](#).

QUESTION NO: 11

You have been tasked to ensure that audit logs on mobility gateways contain accurate timestamps, keeping security in mind. Which configuration change would best secure the time clock against attacks?

- A. Modify the ACL AllowList to deny NTP
- B. Turn on Use NTP authentication toggle and set the parameters
- C. Use an ACL in the communication path
- D. Modify the audit log timezone to match the mobility gateways

ANSWER: B

Explanation:

Turn on Use NTP authentication toggle and set the parameters is the appropriate configuration because it protects the gateway's time synchronization process with cryptographic authentication. Accurate time is fundamental to the usefulness of audit logs: timestamps must reliably establish the sequence of authentication events, configuration changes, security alerts, and administrative actions. Without a trusted time source, an attacker able to spoof or alter NTP responses could shift the device clock, weakening incident investigation and correlation with logs from other systems.

NTP authentication requires the mobility gateway and its approved NTP server to use a shared authentication key. The gateway is configured with the key details, identifies the key as trusted, and associates it with the NTP server. It can then validate that received time updates originate from an authorized source before accepting them. This adds source authenticity to NTP synchronization and directly addresses spoofing and time-manipulation attacks while preserving centralized, consistent time for gateway logging.

Aruba's NTP guidance documents configuring authentication keys and trusted keys when securing NTP synchronization. See the [ArubaOS NTP authentication documentation](#) and the [Network Time Protocol Version 4 specification](#) for the NTP authentication model.

QUESTION NO: 12

A network administrator wants to configure an 802.1X supplicant for a wireless network that includes the following:

AES encryption

EAP-MSCHAPv2-based user and machine authentication

Validation of server certificate in Microsoft Windows 10

The network administrator creates a WLAN profile and selects the **Change connection settings** option. Then the network administrator changes the security type to **Microsoft: Protected EAP (PEAP)** and enables **user and machine authentication** under Additional Settings.

What must the network administrator do next to accomplish the task? (Select two)

- A. Enable server certificate validation
- B. Select Secured password (EAP-MSCHAP v2) as the authentication method
- C. EAP-TLS-based user and machine authentication
- D. Change default RC4 encryption for AES

ANSWER: A B

Explanation:

After selecting Microsoft: Protected EAP (PEAP), the administrator must configure the inner EAP method used within the protected TLS tunnel. Selecting **Secured password (EAP-MSCHAP v2)** as the authentication method fulfills the requirement for PEAP with EAP-MSCHAPv2 authentication. PEAP establishes an encrypted tunnel using the authentication server certificate, and EAP-MSCHAPv2 then carries the machine or user credential exchange inside that tunnel. The previously selected user-and-machine mode controls whether the Windows supplicant authenticates in the computer context before logon and in the user context after logon.

Enable server certificate validation is also required. This causes the Windows supplicant to validate the certificate presented by the RADIUS server during the PEAP TLS negotiation, including trust in the issuing certification authority and, when configured, the expected server identity. Certificate validation is an essential part of securely deploying PEAP because it lets the client establish that it is communicating with a trusted authentication server before supplying credentials. AES is supplied by the WPA2/WPA3-Enterprise wireless security configuration, while these PEAP settings configure the 802.1X authentication behavior. Microsoft documents the available PEAP profile elements at [PEAPProfile Schema](#), and Aruba documents PEAP operation at [EAP-PEAP Authentication](#).

QUESTION NO: 13

Based on best practices, if an SSID is configured for a primary and secondary gateway cluster with cluster preemption enabled, which will decide if the APs move to the secondary gateway cluster if all of the nodes in the primary gateway cluster are down?

- A. Tunnel orchestrator for LAN tunnel service in HPE Aruba Networking Central
- B. Cluster leader in the primary gateway cluster
- C. Every AP individually
- D. Cluster leader in the secondary gateway cluster

ANSWER: C

Explanation:

Every AP individually is correct. An AP maintains its own tunnel connectivity to the gateway cluster selected by the SSID configuration. It uses the configured primary cluster as its preferred destination and monitors tunnel availability through its normal control-plane and keepalive behavior. If no gateway in that primary cluster is reachable, the AP independently determines that its preferred cluster is unavailable and establishes its tunnel to an available gateway in the configured secondary cluster. This behavior avoids dependence on a single cluster-leader decision during a complete primary-cluster outage and lets APs recover according to their own reachability and timing.

Cluster preemption affects the preferred-cluster behavior after connectivity is restored. With preemption enabled, APs using the secondary cluster can be directed by the configured preference logic to return to the primary cluster when that primary service is again available. It does not make a primary or secondary cluster leader responsible for detecting an individual AP's tunnel loss and initiating the initial failover. Aruba Central gateway-cluster configuration and AOS tunnel-management guidance are available in the [HPE Aruba Networking Central documentation](#) and the [HPE Aruba Networking technical documentation portal](#).

QUESTION NO: 14 - (SIMULATION)

Match each Group Based Policy (GBP) role description to its respective role ID.

ANSWER: Check the explanation for the answer

Explanation:

Correct role-ID matching:

Default GBP role → 0


Infrastructure GBP role → 2

User-defined GBP role → 100-8191

Role ID 0 is reserved for the default role, and role ID 2 is reserved for infrastructure. Custom/user-defined Group Based Policy roles must use IDs from 100 through 8191.

QUESTION NO: 15 - (SIMULATION)

An administrator is creating a fabric with NetConductor in HPE Aruba Networking Central. Considering an EVPN VXLAN fabric, click on the most appropriate layer to be configured as a Rome-Reflector Persona.



ANSWER: Check the explanation for the answer

Explanation:

Click the Core layer (the two centrally located switches labeled *Core*).

In an EVPN-VXLAN campus fabric, the Route Reflector (RR) persona is best placed on redundant core devices. The core has resilient, high-bandwidth connectivity to the aggregation/access fabric and is an appropriate centralized location for reflecting EVPN BGP routes. The Services Aggregation layer is intended to connect shared services, such as gateway/security and service infrastructure, rather than to provide the fabric's RR control-plane role.

QUESTION NO: 16

Which statements accurately describe OSPF Graceful Restart (when the restarting router is able to Keep its forwarding tables across the restart)? (Select two.)

- A. The GR helper role is supported on AOS-CX 6100 switches.
- B. VSF Failover and Graceful-Restart require a VSF secondary member in the VSF stack
- C. Bidirectional Forwarding Detection for OSPF and GR are mutually exclusive features.
- D. OSPF Routers listen for Grace-LSAs on each network segment where there is an OSPF adjacency.
- E. You must ensure your VSF stack has a secondary member when acting as a GR helper

ANSWER: A D

Explanation:

The GR helper role is supported on AOS-CX 6100 switches. In the OSPF graceful-restart model, a router that is not restarting can act as a helper for an adjacent restarting router. The helper recognizes the restart procedure and maintains the relevant adjacency and topology information for the negotiated grace period, allowing the restarting router's preserved forwarding state to continue carrying traffic while its OSPF control plane recovers. This helper capability is available on the AOS-CX 6100 platform; "AOX-CX" in the statement is evidently a product-name typo for AOS-CX.

OSPF Routers listen for Grace-LSAs on each network segment where there is an OSPF adjacency. Grace-LSAs are the protocol mechanism used to signal a planned graceful restart to neighboring OSPF routers. A neighbor receives and processes the Grace-LSA on the interface associated with the adjacency, which enables it to provide helper behavior rather than immediately treating the restarting neighbor as unavailable. This per-interface signaling is fundamental to OSPF graceful restart and is defined by the OSPF graceful-restart standard. See Aruba's [AOS-CX OSPF Graceful Restart documentation](#) and [RFC 3623](#).