

DUMPSBOSS.

Aruba Certified Campus Access Professional Exam

HP HPE7-A01

Version Demo

Total Demo Questions: 17

Total Premium Questions: 170

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Design an Aruba Campus Access Solution	65
Topic 2, Implement an Aruba Campus Access Solution	90
Topic 3, Troubleshoot an Aruba Campus Access Solution	13
Topic 4, Mix Questions	2
Total	170

QUESTION NO: 1

What is the best practice for handling voice traffic with dynamic segmentation on AOS-CX switches?

- A. Switch authentication and local forwarding of the voice traffic
- B. Switch authentication and user-based tunneling of the voice traffic.
- C. Central authentication and port-based tunneling of the voice traffic.
- D. Controller authentication and port-based tunneling of all traffic

ANSWER: A

Explanation:

Switch authentication and local forwarding of the voice traffic is the recommended design because voice is highly sensitive to delay, jitter, and packet loss. In this model, the AOS-CX switch acts as the enforcement point at the access edge: it authenticates the endpoint, receives the appropriate role or policy information from the policy infrastructure, and applies the resulting access controls directly on the switch. Once authorized, voice packets remain in the local wired forwarding path rather than being carried through an overlay tunnel to another device. This preserves a short, predictable path for real-time media and helps maintain call quality while still allowing identity-based access policy and segmentation to be enforced. It also keeps voice handling efficient at the access layer, where QoS marking, trust behavior, and prioritization can be applied close to the phone. Aruba's AOS-CX security documentation describes the switch-based role and policy enforcement used with dynamic segmentation and User-Based Tunneling capabilities; see the [AOS-CX User-Based Tunneling documentation](#) and the [AOS-CX 802.1X documentation](#).

QUESTION NO: 2

An operations team wants to deploy the same AOS-CX Network Analytics Engine monitoring logic to several interfaces while using different threshold values for each interface.

Which statements regarding NAE agents are true? (Select two.)

- A. Each NAE script can create only one agent on a switch.
- B. A single NAE script can be used to create multiple agents.
- C. Agent parameters can provide different values for separate instances of the same script.
- D. NAE agents require an external telemetry collector to evaluate monitored conditions.
- E. NAE agents can report data but cannot initiate an action when a rule condition is met.

ANSWER: B C

Explanation:

A single NAE script can be used to create multiple agents is correct. An agent is an instance of an NAE script, allowing the same monitoring logic to be applied to different interfaces or other targets.

Agent parameters can provide different values for separate instances of the same script is also correct. Parameters allow each agent instance to use values appropriate to the target it monitors without requiring duplicate scripts. NAE agents run on the switch and do not require an external collector for their monitoring logic. An agent is not limited to only reporting data; rules can trigger configured actions when monitored conditions are met.

QUESTION NO: 3

Which statements are true regarding a VXLAN implementation on Aruba Switches? (Select two.)

- A. MTU size must be increased beyond the default

- B. VNIs encapsulate and decapsulate VXLAN traffic
- C. VTEPs encapsulate and decapsulate VXLAN traffic
- D. They are only available for datacenter switches (CX 8k, 9k, 10k)
- E. All Aruba CX switches support VXLAN.

ANSWER: A C

Explanation:

MTU size must be increased beyond the default is correct because VXLAN carries the original Ethernet frame inside additional outer Ethernet, IP, UDP, and VXLAN headers. This encapsulation adds approximately 50 bytes for an IPv4 underlay, so the underlay path must have sufficient MTU to transport the larger encapsulated packet without fragmentation or drops. Aruba CX VXLAN deployment guidance therefore requires configuring an adequately sized MTU consistently on the relevant underlay interfaces; a 1550-byte MTU is a common minimum when the tenant workload uses standard 1500-byte frames.

VTEPs encapsulate and decapsulate VXLAN traffic is also correct. A Virtual Tunnel Endpoint is the switch function that maps locally received traffic into the appropriate VXLAN Network Identifier and encapsulates it for transport across the IP underlay. At the remote endpoint, the VTEP removes the VXLAN, UDP, and outer IP/Ethernet headers and forwards the restored inner frame within its associated virtual network. This endpoint behavior allows a Layer 2 overlay segment to span a routed Layer 3 fabric while preserving tenant and broadcast-domain separation. Aruba documents VXLAN configuration, VTEP operation, and MTU planning in its [AOS-CX VXLAN Guide](#) and [VXLAN overview documentation](#).

QUESTION NO: 4

With Aruba CX 6300, how do you configure ip address 10 10 10 1 for the interface in default state for interface 1/1/1?

- A. int 1/1/1. switching, ip address 10 10 10 1/24
- B. interface 1/1/1; no routing; ip address 10.10.10.1/24
- C. int 1/1/1. ip address 10.10.10.1/24
- D. int 1/1/1. routing, ip address 10.10.10 1/24

ANSWER: B

Explanation:

To configure an IPv4 address directly on physical interface 1/1/1 of an Aruba CX 6300, the port must operate as a routed interface. ArubaOS-CX interfaces are Layer 2 switchports by default, so entering `no routing` in the interface context converts the port to Layer 3 operation. Once the interface is routed, `ip address 10.10.10.1/24` assigns the requested address and prefix length. The `/24` prefix corresponds to the 255.255.255.0 subnet mask, defining the directly connected IPv4 network for the interface. The commands are entered in sequence: `interface 1/1/1`, followed by `no routing`, then the IP-address command. If the port has been administratively disabled, `no shutdown` can additionally be required to bring the interface operationally up, but it is not the command that changes the port from its default switching behavior to routed behavior. Aruba documents the `no routing` interface command and IPv4 address configuration in its [AOS-CX IP Services command reference](#) and [ip address command documentation](#).

QUESTION NO: 5

Which statements are true regarding a VXLAN Implementation on HPE Aruba Networking switches? (Select two.)

- A. They are only available for datacenter switches (CX 8k, 9k, 10k).
- B. VNIs encapsulate and decapsulate VXLAN traffic.

- C. MTU size must be increased beyond the default.
- D. All AOS-CX switches support VXLAN.
- E. VTEPs encapsulate and decapsulate VXLAN traffic.

ANSWER: C E

Explanation:

MTU size must be increased beyond the default. VXLAN transports an original Layer 2 Ethernet frame inside an outer IP/UDP/VXLAN header. This encapsulation adds approximately 50 bytes when using IPv4 underlay addressing (and more with IPv6), so the physical underlay path must support a sufficiently large MTU to carry the encapsulated packet without fragmentation or drops. In an Aruba AOS-CX VXLAN design, the required MTU must be configured consistently across the underlay interfaces and any intermediate network devices. Aruba documentation commonly uses an MTU of at least 1550 bytes to accommodate VXLAN overhead when endpoints use standard 1500-byte frames.

VTEPs encapsulate and decapsulate VXLAN traffic. A VXLAN Tunnel Endpoint is the device function at the edge of the VXLAN overlay. For traffic entering the overlay, the VTEP maps the local Layer 2 segment to its VXLAN Network Identifier, encapsulates the Ethernet frame in VXLAN-over-UDP, and sends it across the routed IP fabric. At the receiving overlay edge, the destination VTEP removes the outer headers and forwards the original Ethernet frame into the appropriate local segment. This VTEP behavior permits Layer 2 logical networks to extend across a Layer 3 underlay. See Aruba's [AOS-CX VXLAN documentation](#) and the VXLAN format defined in [RFC 7348](#).

QUESTION NO: 6

What are two advantages of splitting a larger OSPF area into a number of smaller areas? (Select two)

- A. It extends the LSDB
- B. It increases stability
- C. it simplifies the configuration.
- D. It reduces processing overhead.
- E. It reduces the total number of LSAs

ANSWER: B D

Explanation:

It increases stability because OSPF uses a separate link-state database (LSDB) and SPF calculation scope for each area. A topology change is flooded within its local area, so routers in unrelated areas do not need to receive the intra-area topology update or rerun SPF because of that local event. This containment improves convergence behavior and limits the operational impact of failures or frequent changes in one part of a campus network.

It reduces processing overhead because each router maintains and calculates shortest paths from the LSDB for its own area rather than from one flat, network-wide topology database. Smaller area databases generally mean fewer link-state records to store, age, flood, and process during SPF recalculation. Area Border Routers communicate reachability between areas using inter-area summary information, allowing the hierarchical design to scale. OSPF's area hierarchy and the role of Area Border Routers are defined in [RFC 2328, OSPF Version 2](#), including the separate LSDB maintained per area and the inter-area routing procedure. Aruba's OSPF configuration guidance also describes the use of areas to organize OSPF routing domains: [AOS-CX OSPF Guide](#).

QUESTION NO: 7

Which feature allows the device to remain operational when a remote link failure occurs between a Gateway cluster and a RADIUS server that is either in the cloud or a datacenter?

- A. MAC caching

- B. MAC Authentication
- C. Authentication survivability
- D. Opportunistic key caching

ANSWER: C

Explanation:

Authentication survivability allows an Aruba Gateway cluster to continue providing network access when connectivity to its remote RADIUS authentication infrastructure is lost, such as when the RADIUS service resides in a cloud environment or a data center that becomes unreachable. During normal operation, the gateway receives successful RADIUS authentication results and securely caches the information needed to validate subsequent client access attempts. If the remote link fails, the gateway can use these cached authentication records to authenticate eligible clients locally rather than requiring a live RADIUS transaction. This preserves service availability for users and devices that have previously authenticated successfully, reducing the operational impact of a WAN, cloud, or data-center connectivity outage. The feature is especially useful for distributed branches that rely on centrally hosted identity services, because access decisions can continue at the edge during a temporary upstream failure. Aruba also documents authentication survivability support for MAC-based client access scenarios, enabling continued access for applicable endpoint deployments while the RADIUS server cannot be contacted. See Aruba's [Authentication Survivability Technical Guide](#) and the [Aruba Central authentication survivability documentation](#).

QUESTION NO: 8

With the CX 6000 48G switch with uplinks of 1/1/47 and 1/1/48, what does the switch do when a client port detects a loop and tx-disable parameter is used?

- A. The ports that confirmed the loop are disabled.
- B. The ports that transmitted and received the loop are disabled.
- C. The port that transmitted the loop is disabled.
- D. The port that received the loop is disabled.

ANSWER: D

Explanation:

The port that received the loop is disabled. ArubaOS-CX loop protection works by transmitting loop-protection packets from enabled interfaces and monitoring for receipt of a packet that the switch itself originated. Receiving its own loop-protection packet on an interface establishes that a Layer 2 forwarding loop exists through that interface. When the configured loop-protect action is `tx-disable`, the switch applies the protective action to the interface on which the loop packet was received. Specifically, it disables transmission from that loop-detecting interface, preventing the switch from continuing to inject traffic into the loop while preserving the intended loop-detection behavior. In practical terms, the affected client-facing port is placed into the tx-disabled condition; its transmit direction is disabled rather than the uplink ports being selected merely because they provide connectivity. This action isolates the forwarding path that exposed the loop and limits broadcast, multicast, and unknown-unicast traffic amplification. The event and affected interface can be verified through loop-protect status and event logging, and the port can be restored after the cabling or downstream switching loop has been corrected. Aruba documents loop protection and its available actions in the ArubaOS-CX Layer 2 Bridging Guide: [ArubaOS-CX Layer 2 Bridging Guide](#). Product documentation for the CX 6000 family is available from [Aruba CX 6000 documentation](#).

QUESTION NO: 9 - (SIMULATION)

Select the Aruba stacking technology matching each option (Options may be used more than once or not at all.)

ANSWER: See the explanation for the answer

Explanation:

Match the options as follows:

- a) Supports up to 10 devices per stack: VSF
- b) Supports two devices per stack: VSX
- c) Individual ISL links up to 400G are supported: VSX
- d) Individual ISL links up to 50G are supported: VSF
- e) Maximum aggregate ISL bandwidth of 200G is supported: VSF

VSF (Virtual Switching Framework) creates a multi-member virtual stack, while VSX is a two-switch virtualization technology designed for high availability and can use higher-speed inter-switch links.

QUESTION NO: 10

A customer has deployed APs at three offices. All APs require the same corporate WLAN, authentication, and radio configuration, but each office must be represented separately for location-based monitoring and floor-plan management in HPE Aruba Networking Central.

How should the devices be organized?

- A. Create one group for each floor plan and use labels for configuration.
- B. Place all APs in one site and create a different group for each office.
- C. Place the APs in a common group and assign each office to a separate site.
- D. Place every AP in its own group and its own site.

ANSWER: C

Explanation:

The APs should be assigned to a common **group** for their shared configuration and assigned to separate **sites** for their physical locations. Groups provide the configuration scope for devices with common policy requirements, while sites provide location organization used for operational visibility and map or floor-plan context.

Creating a separate group for every physical office would duplicate the same WLAN and radio configuration. A label can assist filtering but does not replace the configuration and location functions of groups and sites.

QUESTION NO: 11

The administrator notices that wired guest users that have exceeded their bandwidth limit are not being disconnected. Access Tracker in ClearPass indicates a disconnect CoA message is being sent to the AOS-CX switch.

An administrator has performed the following configuration

```
Access1(config)# ip dns host cppm.arubatraining.com 10.254.1.23 vrf mgmt
Access1(config)# radius-server host cppm.arubatraining.com key plaintext aruba123 vrf mgmt
Access1(config)# aaa group server radius cppm
Access1(config-sg)# server cppm.arubatraining.com vrf mgmt
Access1(config-sg)# exit
Access1(config)# aaa accounting port-access start-stop interim 5 group cppm
Access1(config)# radius dyn-authorization client cppm.arubatraining.com secret-key plaintext aruba123 vrf mgmt
Access1(config)# radius dyn-authorization enable
```

What is the most likely cause of this issue?

- A. Change of Authorization has not been globally enabled on the switch
- B. The SSL certificate for CPPM has not been added as a trust point on the switch
- C. There is a mismatch between the RADIUS secret on the switch and CPPM.
- D. There is a time difference between the switch and the ClearPass Policy Manager

ANSWER: A

Explanation:

Change of Authorization has not been globally enabled on the switch. ClearPass can create and transmit a Disconnect-Request CoA when the guest user reaches the configured bandwidth limit, and Access Tracker reporting that the message was sent confirms that ClearPass has initiated the action. However, an AOS-CX switch must also be configured to accept RADIUS CoA requests. The global `radius-server coa enable` setting enables the switch's CoA listener and processing capability. Without that global setting, the switch does not act on the Disconnect-Request, so the authenticated wired client session remains active despite ClearPass sending the request.

CoA is a RADIUS mechanism used after initial authentication to alter an active session, including terminating it with a Disconnect-Request. For this workflow to operate, ClearPass must address the request to the correct network-access device and the switch must have CoA enabled globally, as well as the appropriate RADIUS/CoA communication settings. Enabling CoA on the AOS-CX switch allows it to receive, validate, and process ClearPass disconnect requests, which causes the guest port session to be removed when the enforcement condition is met. Aruba documents ClearPass CoA behavior in the [ClearPass Change of Authorization documentation](#) and the required AOS-CX RADIUS capabilities in the [AOS-CX RADIUS documentation](#).

QUESTION NO: 12

What are the requirements to ensure that WMM is working effectively ' ? (Select two)

- A. The APs and the controller are Wi-Fi CERTIFIED for WMM which is enabled
- B. All APs need to be from the AP-5xx series and AP-6xx series which are Wi-Fi CERTIFIED 6.
- C. The Client must be Wi-Fi CERTIFIED for WMM and configured for WMM marking.
- D. The Aruba AOS10 APs installed have to be converted to controlled mode
- E. The AP needs to be connected via a tagged VLAN to the wired port

ANSWER: A C

Explanation:

WMM operation depends on QoS capability and configuration at both ends of the wireless connection. **The APs and the controller are Wi-Fi CERTIFIED for WMM which is enabled** is correct because WMM must be enabled in the WLAN/SSID configuration for the infrastructure to advertise and apply the WMM access categories. ArubaOS uses WMM to support wireless QoS handling, mapping traffic to the voice, video, best-effort, and background access categories. Enabling it allows the AP to contend for wireless medium access according to the appropriate category parameters and to preserve priority treatment for applicable traffic.

The Client must be Wi-Fi CERTIFIED for WMM and configured for WMM marking. is also correct because client participation is required for effective end-to-end wireless prioritization. A WMM-capable client recognizes the AP's WMM advertisement and can use the appropriate user-priority markings/access categories for traffic such as voice and video. The AP can then process those markings and schedule frames using the appropriate WMM category. Wi-Fi Alliance certification provides interoperability assurance for the WMM implementation, while correct client marking ensures the traffic enters the intended QoS treatment path. Aruba's WLAN QoS documentation describes WMM configuration and traffic handling, and the Wi-Fi Alliance summarizes the WMM certification program: [ArubaOS WMM documentation](#) and [Wi-Fi CERTIFIED WMM](#).

QUESTION NO: 13

You are doing tests in your lab and with the following equipment specifications:

• AP1 has a radio that generates a 20 dBm signal

• AP2 has a radio that generates a 8 dBm signal

• AP1 has an antenna with a gain of 7 dBi.

â€¢ AP2 has an antenna with a gain of 12 dBi.

â€¢ The antenna cable for AP1 has a 3 dB loss

â€¢ The antenna cable for AP2 has a 3 dB loss.

What would be the calculated Equivalent Isotropic Radiated Power (EIRP) for AP1?

- A. 2 dBm
- B. 8 dBm
- C. 22 dBm
- D. 24 dBm

ANSWER: D

Explanation:

24 dBm is the calculated EIRP for AP1. EIRP expresses the effective radiated power referenced to an ideal isotropic antenna. To calculate it, start with the radio's conducted transmit power, add the antenna gain, and subtract any loss between the radio and antenna. For AP1, the conducted power is 20 dBm, the antenna gain is 7 dBi, and the antenna-cable loss is 3 dB. Therefore, the calculation is $20 \text{ dBm} + 7 \text{ dBi} - 3 \text{ dB}$, which results in 24 dBm EIRP.

This calculation is important in WLAN design and deployment because the radiated power, rather than radio output power alone, is what must be evaluated against regulatory limits and used when estimating RF coverage. Antenna gain focuses energy relative to an isotropic radiator, while cable attenuation reduces the power that reaches the antenna. The units can be combined in this link-budget calculation because gain and loss are logarithmic ratios expressed in decibels, applied to the absolute power level expressed in dBm. See the FCC's [Office of Engineering and Technology bulletins](#) and Aruba's [RF terminology documentation](#) for RF and antenna-power concepts.

QUESTION NO: 14

Your manufacturing client is deploying twenty headless scanners in their warehouse. These new devices do not support 802.1X authentication.

How does the gateway determine the device's role and VLAN derivation-rules when using MPoSK Local?

- A. From the Type-Length-Value based on the Aruba-MPSK-Key-Name.
- B. It pulls the device roles from HPE Aruba Networking Central during deployment.
- C. From the device's Calling-Station-ID in the RADIUS Access-Request.
- D. From the MPoSK roles defined in HPE Aruba Networking Central's security dashboard.

ANSWER: A

Explanation:

With MPoSK Local, the gateway uses the *Aruba-MPSK-Key-Name* Type-Length-Value information associated with the pre-shared key to identify the applicable MPoSK role. An MPoSK deployment can assign a distinct key to a device or device group, which makes it suitable for headless scanners that cannot perform 802.1X authentication. The key name acts as the identifier that links the successfully validated MPoSK to its configured access policy. Once that MPoSK role is identified, the gateway applies the role's configured user-role and VLAN derivation behavior to the client session. This enables consistent segmentation and access control without requiring interactive credentials, a certificate, or an 802.1X supplicant on the scanner. The configuration is local from the gateway's operational perspective: the key-name TLV provides the policy lookup context at connection time, and the resulting MPoSK role determines the access treatment for that session. Aruba's MPoSK documentation describes using MPoSKs to provide differentiated access for non-802.1X clients, while the AOS documentation covers the gateway policy and role framework used to enforce client access. See [Aruba Central MPoSK documentation](#) and [ArubaOS user-role documentation](#).

QUESTION NO: 15 - (SIMULATION)

List the firewall role derivation flow in the correct order

ANSWER: See the explanation for the answer

Explanation:

The firewall role derivation flow, in precedence order, is:

1. **Server-derived role**
2. **User-derived role**
3. **Authentication default role**
4. **Initiation role assigned**

Aruba first uses a role returned by the authentication server. If none is supplied, it evaluates user-role derivation rules, then uses the AAA authentication default role, and finally applies the initiation role as the last fallback.

QUESTION NO: 16

A network operations team must monitor Aruba infrastructure over a shared management network. The security policy requires that management messages cannot be read or modified by an observer.

Which protections are provided when SNMPv3 is configured with the authPriv security level? (Select two.)

- A. Use of an SNMP community string for user authentication.
- B. Authentication and integrity validation of SNMP messages.
- C. Anonymous access to all managed-object values.
- D. Automatic authorization to modify every MIB object.
- E. Encryption of SNMP protocol data.

ANSWER: B E

Explanation:

Authentication and integrity validation of SNMP messages is correct because the authentication component verifies the SNMPv3 user and detects message modification. **Encryption of SNMP protocol data** is also correct because the privacy component protects the confidentiality of the SNMP message contents.

Community strings are associated with SNMPv1 and SNMPv2c, not SNMPv3 USM authentication. authPriv does not make management traffic anonymous and does not replace the need to authorize users for the required MIB access.

QUESTION NO: 17

A high-density wireless deployment has many low-bandwidth clients that contend for airtime on the same Wi-Fi 6 radio channel.

Which statements describe how OFDMA improves this situation? (Select two.)

- A. OFDMA divides a channel into resource units that can be assigned to different clients.
- B. OFDMA requires a 160 MHz channel before it can be used.
- C. OFDMA eliminates CSMA/CA contention for all wireless traffic.
- D. OFDMA supports scheduled uplink and downlink multiuser transmissions.

E. OFDMA separates clients by assigning each client a dedicated spatial stream.

ANSWER: A D

Explanation:

OFDMA divides a channel into resource units that can be assigned to different clients is correct. Rather than allocating the entire channel to one client transmission, the AP can schedule smaller groups of subcarriers for multiple clients.

OFDMA supports scheduled uplink and downlink multiuser transmissions is also correct. This scheduling reduces contention overhead and is particularly useful for many clients carrying short or low-rate frames. OFDMA does not require a 160 MHz channel, does not eliminate all CSMA/CA behavior, and should not be confused with MU-MIMO, which separates simultaneous clients through spatial streams.