

DUMPSBOSS.

Aruba Certified Network Security Expert Written Exam

HP HPE6-A84

Version Demo

Total Demo Questions: 9

Total Premium Questions: 98

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Security Architecture	1
Topic 2, Security Design	11
Topic 3, Security Implementation	58
Topic 4, Security Monitoring and Management	22
Topic 5, Security Troubleshooting	6
Total	98

QUESTION NO: 1

You want to use Device Insight tags as conditions within CPPM role mapping or enforcement policy rules.

What guidelines should you follow?

- A.** Create an HTTP authentication source to the Central API that queries for the tags. To use that source as the type for rule conditions, add it as an authorization source for the service in question.
- B.** Use the Application type for the rule conditions; no extra authorization source is required for services that use policies with these rules.
- C.** Use the Endpoints Repository type for the rule conditions; Add Endpoints Repository as a secondary authentication source for services that use policies with these rules.
- D.** Use the Endpoint type for the rule conditions; no extra authorization source is required for services that use policies with these rules.

ANSWER: D

Explanation:

Use the Endpoint type for the rule conditions; no extra authorization source is required for services that use policies with these rules. Device Insight integrates with ClearPass Policy Manager by synchronizing the endpoint intelligence it discovers, including Device Insight tags, into ClearPass endpoint data. Consequently, policy authors can select the Endpoint condition type and the *Device Insight Tags* attribute when building role-mapping rules or enforcement-policy rules. This makes the tag directly available during policy evaluation, enabling actions such as assigning a role, applying a posture-related response, or returning a specific enforcement profile based on the endpoint's discovered classification or risk context.

The important operational point is that these tags are endpoint attributes maintained as part of the Device Insight/ClearPass integration, rather than data that must be retrieved by a separate authorization lookup for each authentication transaction. A service using rules based on this attribute therefore does not need an additional authorization source solely to expose Device Insight tags. This design keeps policy construction straightforward while allowing ClearPass to use continuously updated endpoint context. Aruba documentation describes Device Insight as providing endpoint visibility and classification for policy decisions; see the [Aruba Central Device Insight documentation](#) and the [ClearPass Policy Manager documentation](#).

QUESTION NO: 2

A customer has an AOS 10-based solution, including Aruba APs. The customer wants to use Cloud Auth to authenticate non-802.1X capable IoT devices.

What is a prerequisite for setting up the device role mappings?

- A.** Configuring a NetConductor-based fabric
- B.** Configuring Device Insight (client profile) tags in Central
- C.** Integrating Aruba ClearPass Policy Manager (CPPM) and Device Insight
- D.** Creating global role-to-role firewall policies in Central

ANSWER: B

Explanation:

Configuring Device Insight (client profile) tags in Central is required before device role mappings can be created for Cloud Auth. Cloud Auth can use device profiling to recognize non-802.1X endpoints, such as IoT equipment, based on the characteristics and behavior observed for the client. Device Insight provides this classification and assigns the relevant client-profile tags in Aruba Central.

Those tags supply the matching criteria for a device role mapping. An administrator maps a recognized device profile or tag to the appropriate client role so that, when the endpoint connects, Central can assign the intended role and apply the access controls associated with it. This makes device identification the foundation of the policy decision: the IoT device must first be

categorized consistently before it can receive a role designed for that device type. In practice, administrators should confirm that Device Insight is enabled, that the required client profiles/tags are available, and that the tags reliably identify the target devices before defining the mappings.

Aruba Central documentation describes Cloud Authentication and Policy configuration and its use of client profiling, while the Device Insight documentation covers profile and tag-based identification. See [Aruba Central Cloud Authentication documentation](#) and [Aruba Central Device Insight documentation](#).

QUESTION NO: 3

A customer allows printers to connect through MAC authentication on AOS-CX switches because the printers do not support 802.1X. The customer currently permits any endpoint whose MAC address is found in a static list.

A security review identifies that an attacker could spoof a permitted printer MAC address. The customer wants to retain MAC authentication for the printers but reduce the impact of spoofing.

What is the best recommendation?

- A. Use ClearPass profiling and enforcement to validate expected printer characteristics and restrict exceptions.
- B. Increase the number of permitted printer MAC addresses in the static endpoint list.
- C. Assign all MAC-authenticated printers to a dedicated static VLAN with unrestricted access.
- D. Hide the wireless SSID used by the printers from client scans.

ANSWER: A


Explanation:

Use ClearPass profiling and enforcement to validate that a MAC-authenticated endpoint exhibits expected printer characteristics, and place exceptions in a restricted role. MAC authentication alone is not strong identity proof because a source MAC address can be copied. Profiling adds context from observed endpoint behavior and characteristics, which can identify a mismatch between a claimed printer identity and a device behaving like a computer or another endpoint type.

Profiling does not make MAC authentication equivalent to certificate-based authentication, but it provides a meaningful additional policy signal. A longer MAC-address list, a hidden SSID, or a static VLAN assignment does not address MAC-address spoofing.

QUESTION NO: 4

Refer to the exhibit.



A customer requires protection against ARP poisoning in VLAN 4. Below are listed all settings for VLAN 4 and the VLAN 4 associated physical interfaces on the AOS-CX access layer switch:

Interface 1/1/2-1/1/24

```
no shutdown
no routing
vlan access 4
exit
```

Interface lag 1

```
no shutdown
no routing
vlan trunk native 1
vlan trunk allowed 2-4
arp inspection trust
exit
```

vlan 4

```
arp inspection
exit
```

What is one issue with this configuration?

- A. ARP proxy is not enabled on VLAN 4.
- B. LAG 1 is configured as trusted for ARP inspection but should be untrusted.
- C. DHCP snooping is not enabled on VLAN 4.
- D. Edge ports are not configured as untrusted for ARP inspection.

ANSWER: C

Explanation:

DHCP snooping is not enabled on VLAN 4. Dynamic ARP Inspection (DAI), referred to as ARP inspection in AOS-CX configuration, protects hosts against ARP spoofing by checking ARP sender IP-to-MAC bindings. For dynamically addressed clients, the authoritative source for those bindings is the DHCP snooping binding database. DHCP snooping must therefore be enabled for the protected VLAN so that valid client lease bindings can be learned and subsequently used when ARP inspection evaluates ARP traffic. Without DHCP snooping operating on VLAN 4, ARP inspection does not have the expected dynamic binding information for VLAN 4 clients, so the intended ARP-poisoning protection is incomplete. The access switch should enable DHCP snooping globally as required by the platform and enable it specifically for VLAN 4; DHCP server-facing uplinks must also be identified as trusted so valid DHCP server messages can establish bindings. After the bindings are available, ARP inspection on VLAN 4 can validate ARP packets received from client-facing interfaces against those bindings. Aruba's AOS-CX security documentation describes DHCP snooping and its binding database at [DHCP snooping](#) and documents ARP inspection configuration at [Dynamic ARP Inspection](#).

QUESTION NO: 5

A customer has an AOS 10 architecture, which includes Aruba APs. Admins have recently enabled WIDS at the high level. They also enabled alerts and email notifications for several events, as shown in the exhibit.

USER	ACCESS POINT	SWITCH	GATEWAY	CONNECTIVITY	AUDIT	SITE
By Clicking on + icon, you can quickly generate notifications with default notification policy. You can also define the policy by clicking on the tiles. GOT IT						
New Virtual Controller Detected	+	Virtual Controller Disconnected	+	New AP Detected	+	
AP Disconnected	✓	Rogue AP Detected	✓	Infrastructure Attack Detected	✓	
Client Attack Detected	✓	Uplink Changed	+	Modem Plugged	+	
Modem Unplugged	+	Insufficient Power Supplied	+	AP With Missing Radios	+	
AP CPU Utilization	+	AP Memory Utilization	+	Radio Channel Utilization	+	
Radio Noise Floor	+	Connected Clients Per VC	+	Connected Clients Per AP	+	
Radio Frames Retry Percent	+	AP Tunnel Down	+	All AP Tunnels Down	✓	
Radio Non Wi-Fi Utilization	+	IAP Firmware Upgrade Failed	+			

Admins are complaining that they are getting so many emails that they have to ignore them, so they are going to turn off all notifications.

What is one step you could recommend trying first?

- A. Send the email notifications directly to a specific folder, and only check the folder once a week.
- B. Disable email notifications for Rogue AP, but leave the Infrastructure Attack Detected and Client Attack Detected notifications on.
- C. Change the WIDS level to custom, and enable only the checks most likely to indicate real threats.
- D. Disable just the Rogue AP and Client Attack Detected alerts, as they overlap with the Infrastructure Attack Detected alert.

ANSWER: C

Explanation:

Change the WIDS level to custom, and enable only the checks most likely to indicate real threats. This is the appropriate first step because a high WIDS setting enables a broad set of wireless intrusion-detection checks. While that provides extensive visibility, it can also create a substantial volume of events in environments where benign activity, neighboring wireless networks, or expected client behavior repeatedly triggers detections. Excessive alerting causes alert fatigue and makes it more likely that administrators will overlook a genuinely important security event.

A custom WIDS profile lets the security team tune detection behavior to the organization's actual risk profile and wireless environment. Administrators can retain the checks that provide meaningful indicators of rogue infrastructure, infrastructure attacks, or client attacks, while suppressing checks that have been determined to be low-value or consistently nonactionable. This reduces notification noise at its source rather than merely changing how messages are delivered. The team should review alert history, identify recurring events, validate whether they represent real risk, and then enable the most relevant checks while continuing to monitor and refine the policy.

Aruba describes WIDS configuration and monitoring in its [Aruba Central WIDS documentation](#). Additional AOS 10 wireless security guidance is available in the [AOS 10 documentation library](#).

QUESTION NO: 6

You need to install a certificate on a standalone Aruba Mobility Controller (MC). The MC will need to use the certificate for the Web UI and for implementing RadSec with Aruba ClearPass Policy Manager. You have been given a certificate with these settings:

- Subject: CN=mc41.site94.example.com
- No SANs
- Issuer: CN=ca41.example.com
- EKUs: Server Authentication, Client Authentication

What issue does this certificate have for the purposes for which the certificate is intended?

- A. It has conflicting EKUs.
- B. It is issued by a private CA.
- C. It specifies domain info in the CN field instead of the DC field.
- D. It lacks a DNS SAN.

ANSWER: D

Explanation:

It lacks a DNS SAN. A Subject Alternative Name extension is the appropriate certificate field for binding one or more DNS names to the certificate identity. The Mobility Controller's Web UI is accessed as a TLS server, and a browser validates the name used in the URL against the certificate's DNS identity. For the RadSec connection to ClearPass, the controller also needs a certificate identity that ClearPass can validate when the controller acts as the TLS client. Including the controller FQDN, mc41.site94.example.com, as a DNS SAN provides an explicit, standards-based name binding for both uses; further names should be included if the controller is reached by additional DNS aliases.

The certificate's Server Authentication and Client Authentication extended key usages support its intended dual server/client roles. A private issuing CA can also be appropriate when its root and any intermediate CA certificates are trusted by the relevant browsers, controller, and ClearPass deployment. Aruba's certificate guidance describes managing trusted CAs and server certificates on ArubaOS devices, while RFC 6125 defines service identity matching for TLS and prioritizes subject alternative names for DNS-based identities. See [ArubaOS certificate documentation](#) and [RFC 6125: Service Identity in TLS](#).

QUESTION NO: 7

Refer to the scenario.

A customer has an Aruba ClearPass cluster. The customer has AOS-CX switches that implement 802.1X authentication to ClearPass Policy Manager (CPPM).

Switches are using local port-access policies.

The customer wants to start tunneling wired clients that pass user authentication only to an Aruba gateway cluster. The gateway cluster should assign these clients to the "eth-internet" role. The gateway should also handle assigning clients to their VLAN, which is VLAN 20.

The plan for the enforcement policy and profiles is shown below:

The gateway cluster has two gateways with these IP addresses:

- Gateway 1

- o VLAN 4085 (system IP) = 10.20.4.21

- o VLAN 20 (users) = 10.20.20.1

- o VLAN 4094 (WAN) = 198.51.100.14

- Gateway 2

- o VLAN 4085 (system IP) = 10.20.4.22

- o VLAN 20 (users) = 10.20.20.2

- o VLAN 4094 (WAN) = 198.51.100.12

- VRRP on VLAN 20 = 10.20.20.254

The customer requires high availability for the tunnels between the switches and the gateway cluster. If one gateway falls, the other gateway should take over its tunnels. Also, the switch should be able to discover the gateway cluster regardless of whether one of the gateways is in the cluster.

Assume that you have configured the correct UBT zone and port-access role settings. However, the solution is not working.

What else should you make sure to do?

- A. Assign VLAN 20 as the access VLAN on any edge ports to which tunneled clients might connect.
- B. Create a new VLAN on the AOS-CX switch and configure that VLAN as the UBT client VLAN.
- C. Assign sufficient VIA licenses to the gateways based on the number of wired clients that will connect.
- D. Change the port-access auth-mode mode to client-mode on any edge ports to which tunneled clients might connect.

ANSWER: B

Explanation:

Create a new VLAN on the AOS-CX switch and configure that VLAN as the UBT client VLAN. AOS-CX User-Based Tunneling requires a designated UBT client VLAN in addition to the UBT zone and the port-access role attributes. When ClearPass returns the tunneled access result, the switch places the authenticated wired endpoint into this internally designated UBT VLAN and forwards the endpoint's traffic through the UBT tunnel to the selected gateway. For this deployment, the VLAN must be created on the switch and configured as the UBT client VLAN; where the intended user VLAN is VLAN 20, that configuration must align with the gateway-side VLAN 20 service.

The UBT zone provides the gateway controller addresses, including primary and backup information needed for gateway resiliency. However, that zone alone does not provide the local switching construct that activates UBT handling for authenticated clients. Configuring the UBT client VLAN completes the required switch-side forwarding setup, allowing the gateway to apply the returned *eth-internet* role and provide the user VLAN service. Aruba's AOS-CX security documentation describes UBT configuration and its client-VLAN requirement in the [AOS-CX User-Based Tunneling documentation](#). Aruba's [technical documentation portal](#) also provides version-specific UBT command references.

QUESTION NO: 8

Refer to the scenario.

A customer requires these rights for clients in the “medical-mobile” AOS firewall role on Aruba Mobility Controllers (MCs):

Permitted to receive IP addresses with DHCP

Permitted access to DNS services from 10.8.9.7 and no other server

Permitted access to all subnets in the 10.1.0.0/16 range except denied access to 10.1.12.0/22

Denied access to other 10.0.0.0/8 subnets

Permitted access to the Internet

Denied access to the WLAN for a period of time if they send any SSH traffic

Denied access to the WLAN for a period of time if they send any Telnet traffic

Denied access to all high-risk websites

External devices should not be permitted to initiate sessions with “medical-mobile” clients, only send return traffic.

The exhibits below show the configuration for the role.

There are multiple issues with this configuration. What is one change you must make to meet the scenario requirements? (In the options, rules in a policy are referenced from top to bottom. For example, “medical-mobile” rule 1 is “ipv4 any any svc-dhcp permit,” and rule 8 is “ipv4 any any any permit.”)

- A. In the “medical-mobile” policy, move rules 2 and 3 between rules 7 and 8.
- B. In the “medical-mobile” policy, change the subnet mask in rule 3 to 255.255.248.0.
- C. Move the rule in the “apprf-medical-mobile-sacl” policy between rules 7 and 8 in the “medical-mobile” policy.
- D. In the “medical-mobile” policy, change the source in rule 8 to “user.”

ANSWER: D

Explanation:

In the “medical-mobile” policy, change the source in rule 8 to “user.” This change is required because the final broad permit rule is currently written as *ipv4 any any any permit*. With *any* as the source, the rule can match traffic initiated by an external host toward a client in the medical-mobile role. That behavior conflicts with the requirement that external devices must not be allowed to initiate sessions with these clients.

Using *user* as the source limits the broad Internet-permit rule to traffic originating from the wireless client to which the role is assigned. Aruba’s stateful firewall maintains session state, so return packets for a session initiated by that client are recognized and permitted without creating a general inbound-access allowance. This preserves client-initiated Internet connectivity while enforcing the intended client-protection boundary against unsolicited inbound sessions.

This is also the appropriate change for a catch-all permit rule because Aruba user-role firewall policies are processed in rule order, with the broad rule normally placed after the more specific DHCP, DNS, internal-network, SSH/Telnet, and web-category controls. Aruba documents user-role access rules and their stateful enforcement behavior in the [ArubaOS User Roles documentation](#) and the [ArubaOS Firewall documentation](#).

QUESTION NO: 9

A customer is replacing UDP-based RADIUS between AOS-CX switches and ClearPass Policy Manager (CPPM) with RadSec. The customer requires encryption of the RADIUS exchange and mutual certificate-based authentication between the switches and CPPM.

Which configuration requirement is essential?

- A.** Configure certificates on both the switches and CPPM, and configure each side to trust the peer CA.
- B.** Configure a longer RADIUS shared secret on the switches and CPPM.
- C.** Enable SNMPv3 authentication and privacy for the ClearPass network devices.
- D.** Configure an IPsec tunnel between each switch management interface and the default gateway.

ANSWER: A

Explanation:

Configure certificates on both the switches and CPPM, and configure each side to trust the CA that issued the peer certificate. RadSec transports RADIUS in TLS, so the switch must validate the ClearPass server certificate and ClearPass must validate the switch client certificate when mutual authentication is required. The certificate identities must also match the configured peer names.

Changing only the shared secret does not establish TLS protection or certificate-based peer authentication. SNMPv3 and IPsec can protect other management or network traffic but do not configure the RadSec TLS relationship.