HPE6-A84$^{Q\&As}$

Aruba Certified Network Security Expert Written

Pass HP HPE6-A84 Exam with 100% Guarantee

Following Questions and Answers are all new published by HP Official Exam Center

QUESTION 1

Refer to the exhibit.

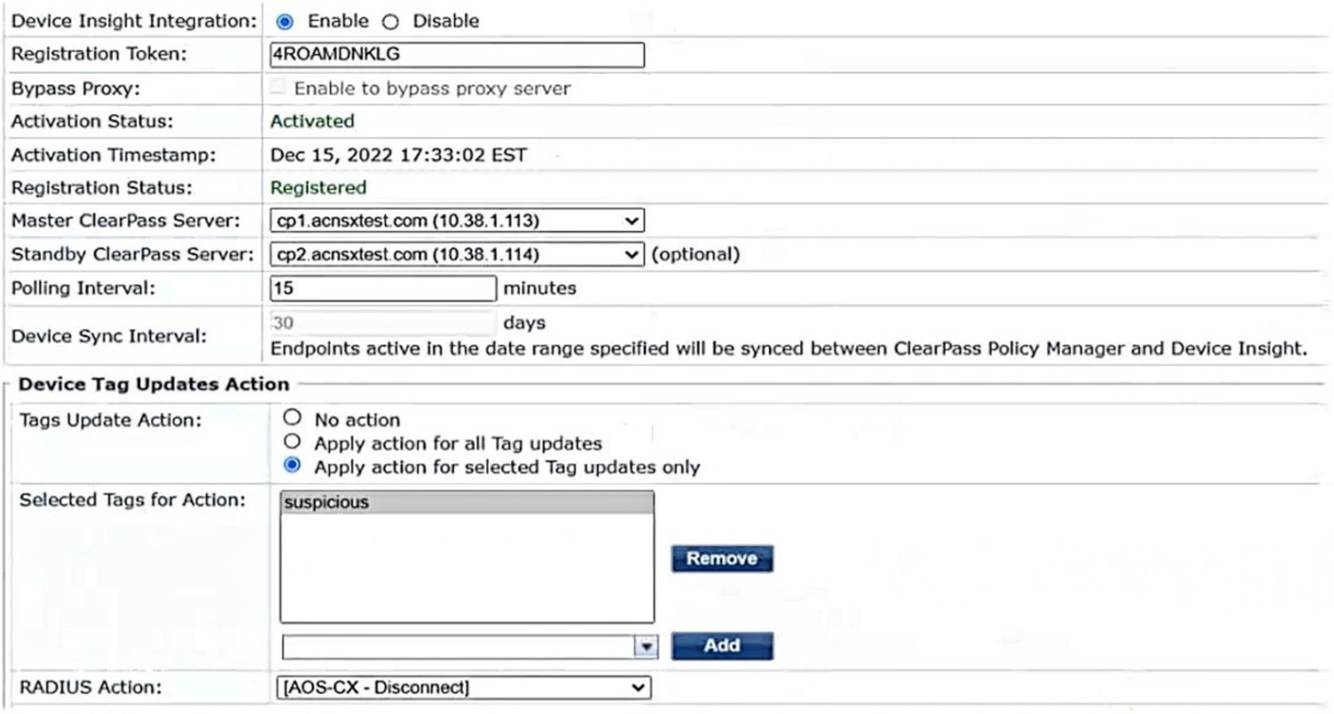

Device Insight Integration:	● Enable ○ Disable
Registration Token:	4ROAMDNKLG
Bypass Proxy:	□ Enable to bypass proxy server
Activation Status:	Activated
Activation Timestamp:	Dec 15, 2022 17:33:02 EST
Registration Status:	Registered
Master ClearPass Server:	cp1.acnsxtest.com (10.38.1.113)
Standby ClearPass Server:	cp2.acnsxtest.com (10.38.1.114) (optional)
Polling Interval:	15 minutes
Device Sync Interval:	30 days Endpoints active in the date range specified will be synced between ClearPass Policy Manager and Device Insight.

Device Tag Updates Action

Tags Update Action:	○ No action ○ Apply action for all Tag updates ● Apply action for selected Tag updates only
Selected Tags for Action:	suspicious
RADIUS Action:	[AOS-CX - Disconnect]

Aruba ClearPass Policy Manager (CPPM) is using the settings shown in the exhibit. You reference the tag shown in the exhibit in enforcement policies related to NASes of several types, including Aruba APs, Aruba gateways, and AOS-CX switches.

What should you do to ensure that clients are reclassified and receive the correct treatment based on the tag?

A. Change the RADIUS action to [Aruba Wireless -Terminate Session] which is supported by all the NASes in question.

B. Change the RADIUS action to [Aruba Wireless - Bounce Switch Port] which is supported by all the NASes in question.

C. Enable profiling in each service using one of these enforcement profiles. Set the profiling action to the correct one for the NASes using that service.

D. Set the Tags Update Action to No Action. Then instead enable the RADIUS CoAs using enforcement profiles in the rules that match clients with the tag shown in the exhibit.

Correct Answer: C

According to the ClearPass Policy Manager User Guide1, the tag shown in the exhibit is a Device Insight tag, which is used to classify and identify devices based on their behavior and characteristics. Device Insight tags can be used as conditions in enforcement policies to apply different actions or roles to devices based on their tags. However, in order to ensure that devices are reclassified and receive the correct treatment based on their tags, profiling must be enabled in each service that uses one of these enforcement profiles. Profiling is a feature that allows ClearPass to dynamically discover and profile devices on the network, and update their attributes and tags accordingly. Profiling also allows ClearPass to send RADIUS Change of Authorization (CoA) messages to the network access servers (NASes) that control the access of the devices, and instruct them to reauthenticate or terminate the sessions of the devices that have changed their tags. The profiling action must be set to the correct one for the NASes using that service, as different NASes may support different types of CoA messages. Therefore, option C is the correct answer.

QUESTION 2

Refer to the scenario.

A customer has an AOS10 architecture that is managed by Aruba Central. Aruba infrastructure devices authenticate clients to an Aruba ClearPass cluster.

In Aruba Central, you are examining network traffic flows on a wireless IoT device that is categorized as "Raspberry Pi" clients. You see SSH traffic. You then check several more wireless IoT clients and see that they are sending SSH also. You want a fast way to find a list of all the IoT clients that have used SSH.

What step can you take?

A. Create and apply a Central client profile tag that selects the SSH application and the clients\\' category.

B. Run a search for SSH traffic and loT client IDs in Aruba ClearPass Policy Manager\\'s (CPPM\\'s) accounting information.

C. Use Central\\'s Live Events monitoring tool to detect which clients meet the desired criteria.

D. Use Central\\'s Gateway IDS/IPS Security Dashboard to search for SSH events and sources.

Correct Answer: C

This is because the Live Events monitoring tool is a feature that allows you to view and filter real-time events and alerts from your network devices and clients on Aruba Central. You can use the Live Events monitoring tool to detect which IoT clients have used SSH by applying the following filters: Category: IoT Application: SSH The Live Events monitoring tool will then display a list of all the IoT clients that have used SSH, along with other information such as their IP address, MAC address, hostname, SSID, AP name, etc. You can also export the list as a CSV file for further analysis or reporting.

A. Create and apply a Central client profile tag that selects the SSH application and the clients\\' category. This is not the fastest way to find a list of all the IoT clients that have used SSH because creating and applying a client profile tag is a

process that involves several steps and might take some time to take effect. A client profile tag is a feature that allows you to group and classify clients based on various criteria, such as device type, OS, category, application, etc. To create

and apply a client profile tag that selects the SSH application and the clients\\' category, you need to do the following:

Navigate to Clients > Client Profile Tags on Aruba Central. Click Add Tag and enter a name and description for the tag. Click Add Rule and select Application as the attribute and SSH as the value. Click Add Rule again and select Category as

the attribute and IoT as the value.

Click Save to create the tag.

Navigate to Clients > Client List on Aruba Central. Select the clients that you want to apply the tag to and click Assign Tag.

Select the tag that you created and click Apply.

After applying the tag, you can then filter the client list by the tag name and see a list of all the IoT clients that have used SSH. However, this method might not be as fast or accurate as using the Live Events monitoring tool, as it depends on

how often the client profile tags are updated and synchronized with Aruba Central.

B. Run a search for SSH traffic and loT client IDs in Aruba ClearPass Policy Manager\\'s (CPPM\\'s) accounting information. This is not the fastest way to find a list of all the IoT clients that have used SSH because running a search in CPPM\\'s

accounting information is a process that involves accessing another system and querying a large amount of data. Accounting information is a feature that allows CPPM to collect and store data about network sessions, such as start time, end

time, duration, bytes sent/received, etc. To run a search for SSH traffic and IoT client IDs in CPPM\\'s accounting information, you need to do the following:

Log in to CPPM and navigate to Monitoring > Live Monitoring > Accounting. Click on Advanced Search and enter SSH as the value for Service Name. Click on Add Filter and enter IoT as the value for Endpoint Category.

Click on Search to run the query.

The query will then return a list of all the network sessions that involved SSH traffic and IoT clients. However, this method might not be as fast or convenient as using the Live Events monitoring tool, as it requires logging in to another system

and searching through a large amount of data that might not be relevant or current. D. Use Central\\'s Gateway IDS/IPS Security Dashboard to search for SSH events and sources. This is not a valid way to find a list of all the IoT clients that

have used SSH because the Gateway IDS/IPS Security Dashboard is a feature that only applies to wired network devices connected to Aruba gateways, not wireless devices connected to Aruba APs. The Gateway IDS/IPS Security

Dashboard is a feature that allows you to monitor and manage security events and alerts from your wired network devices on Aruba Central. You can use the Gateway IDS/IPS Security Dashboard to search for security events related to SSH,

such as brute force attacks or unauthorized access attempts, but not for normal SSH traffic from wireless IoT devices. Therefore, this method will not help you find a list of all the IoT clients that have used SSH.

QUESTION 3

You are designing an Aruba ClearPass Policy Manager (CPPM) solution for a customer. You learn that the customer has a Palo Alto firewall that filters traffic between clients in the campus and the data center.

Which integration can you suggest?

A. Sending Syslogs from the firewall to CPPM to signal CPPM to change the authentication status for misbehaving clients

B. Importing clients\\' MAC addresses to configure known clients for MAC authentication more quickly

C. Establishing a double layer of authentication at both the campus edge and the data center DMZ

D. Importing the firewall\\'s rules to program downloadable user roles for AOS-CX switches more quickly

Correct Answer: A

This option allows CPPM to receive real-time information about the network activity and security posture of the clients from the firewall, and then apply appropriate enforcement actions based on the configured policies 12. For example, if a client is detected to be infected with malware or violating the network usage policy, CPPM can quarantine or disconnect the client from the network 2.

QUESTION 4

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:

1.

Subject: CN=mc41.site94.example.com

2.

No SANs

3.

Issuer: CN=ca41.example.com

4.

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.

Correct Answer: D

A DNS SAN (Subject Alternative Name) is an extension of the X.509 certificate standard that allows specifying additional hostnames or IP addresses that the certificate can be used for. A DNS SAN is useful for validating the identity of the server or client that presents the certificate, especially when the common name (CN) field does not match the hostname or IP address of the server or client. In this case, the certificate has a CN of mc41.site94.example.com, which is the fully qualified domain name (FQDN) of the standalone Aruba Mobility Controller (MC). However, this CN may not match the hostname or IP address that the MC uses for the Web UI or for implementing RadSec with Aruba ClearPass Policy Manager. For example, if the MC uses a different FQDN, such as mc41.example.com, or an IP address, such as 192.168.1.41, for these purposes, then the certificate would not be valid for them. Therefore, the certificate should have a DNS SAN that includes all the possible hostnames or IP addresses that the MC may use for the Web UI and RadSec.

QUESTION 5

Refer to the scenario.

A hospital has an AOS10 architecture that is managed by Aruba Central. The customer has deployed a pair of Aruba 9000 Series gateways with Security licenses at each clinic. The gateways implement IDS/IPS in IDS mode.

The Security Dashboard shows these several recent events with the same signature, as shown below:

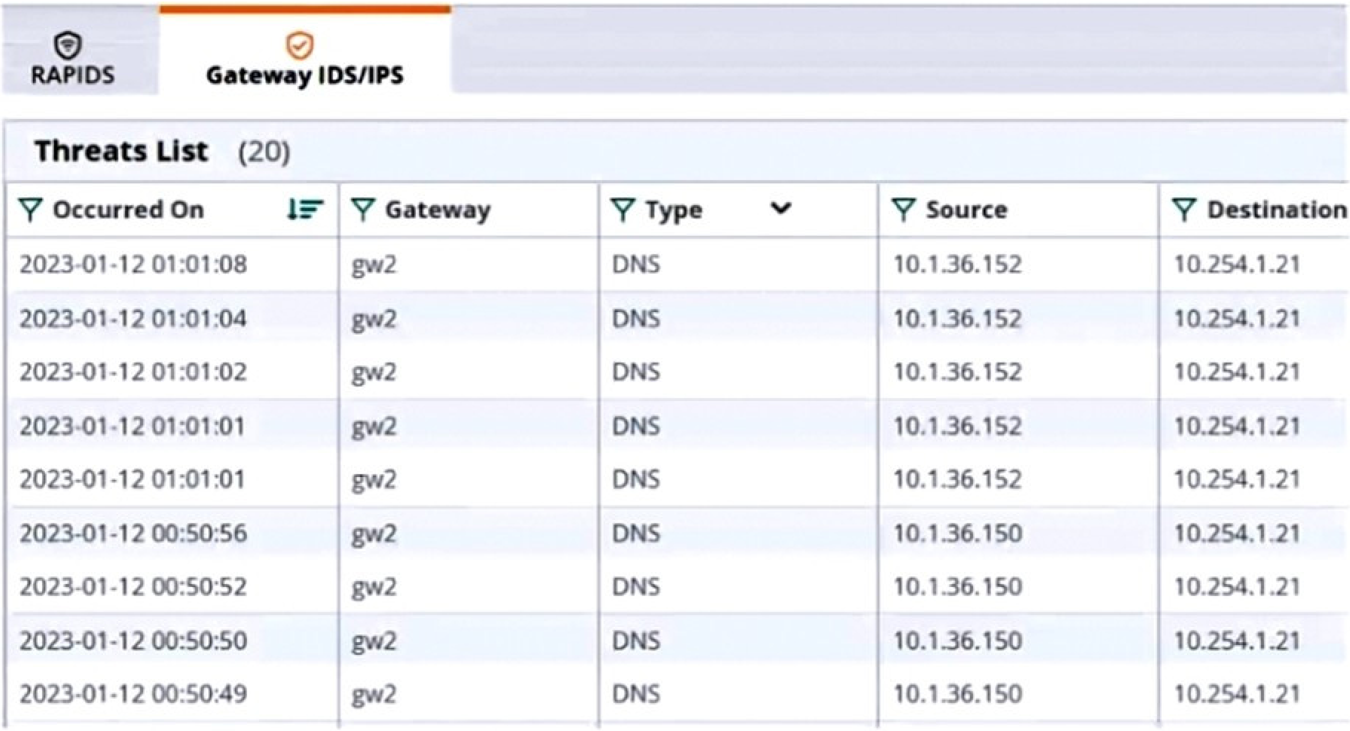

Threats List (20)

Occurred On	Gateway	Type	Source	Destination
2023-01-12 01:01:08	gw2	DNS	10.1.36.152	10.254.1.21
2023-01-12 01:01:04	gw2	DNS	10.1.36.152	10.254.1.21
2023-01-12 01:01:02	gw2	DNS	10.1.36.152	10.254.1.21
2023-01-12 01:01:01	gw2	DNS	10.1.36.152	10.254.1.21
2023-01-12 01:01:01	gw2	DNS	10.1.36.152	10.254.1.21
2023-01-12 00:50:56	gw2	DNS	10.1.36.150	10.254.1.21
2023-01-12 00:50:52	gw2	DNS	10.1.36.150	10.254.1.21
2023-01-12 00:50:50	gw2	DNS	10.1.36.150	10.254.1.21
2023-01-12 00:50:49	gw2	DNS	10.1.36.150	10.254.1.21

Which step could give you valuable context about the incident?

A. View firewall sessions on the APs and record the threat sources\\' type and OS.

B. View the user-table on APs and record the threat sources\\' 802.11 settings.

C. View the RAPIDS Security Dashboard and see if the threat sources are listed as rogues.

D. Find the Central client profile for the threat sources and note their category and family.

Correct Answer: C

The RAPIDS Security Dashboard is a feature of Aruba Central that provides a comprehensive view of the network security status, including IDS/IPS events, rogue APs, and wireless intrusion detection. By viewing the RAPIDS Security Dashboard, you can see if the threat sources are rogue APs that are spoofing legitimate DNS servers or clients. This can give you valuable context about the incident and help you identify the root cause of the attack1 Reference: Aruba Central User Guide

QUESTION 6

Refer to the scenario.

An organization wants the AOS-CX switch to trigger an alert if its RADIUS server (cp.acnsxtest.local) rejects an unusual number of client authentication requests per hour. After some discussions with other Aruba admins, you are still not sure how many rejections are usual or unusual. You expect that the value could be different on each switch. You are helping the developer understand how to develop an NAE script for this use case.

The developer explains that they plan to define the rule with logic like this:

monitor > value

However, the developer asks you what value to include.

What should you recommend?

A. Checking one of the access switches\\' RADIUS statistics and adding 10 to the number listed for rejects

B. Defining a baseline and referring to it for the value

C. Using 10 (per hour) as a good starting point for the value

D. Defining a parameter and referring to it (self ^ramsfname]) for the value

Correct Answer: D

This is because a parameter is a variable that can be defined and modified by the user or the script, and can be used to customize the behavior and output of the NAE script. A parameter can be referred to by using the syntax self ^ramsfname], where ramsfname is the name of the parameter. By defining a parameter for the value, the developer can make the NAE script more flexible and adaptable to different scenarios and switches. The parameter can be set to a default value, such as 10, but it can also be changed by the user or the script based on the network conditions and requirements. For example, the parameter can be adjusted dynamically based on the average or standard deviation of the number of rejects per hour, or based on the feedback from the user or other admins. This way, the NAE script can trigger an alert only when the number of rejects is truly unusual and not just arbitrary. A. Checking one of the access switches\\' RADIUS statistics and adding 10 to the number listed for rejects. This is not a good recommendation because it does not account for the variability and diversity of the network environment and switches. The number of rejects listed for one switch might not be representative or relevant for another switch, as different switches might have different traffic patterns, client types, RADIUS configurations, etc. Moreover, adding 10 to the number of rejects is an arbitrary and fixed value that might not reflect the actual threshold for triggering an alert. B. Defining a baseline and referring to it for the value. This is not a bad recommendation, but it is not as good as defining a parameter. A baseline is a reference point that represents the normal or expected state of a network metric or performance indicator. A baseline can be used to compare and contrast the current network situation and detect any anomalies or deviations. However, a baseline might not be easy or accurate to define, as it might require historical data, statistical analysis, or expert judgment. Moreover, a baseline might not be stable or constant, as it might change over time due to network growth, evolution, or optimization.

C. Using 10 (per hour) as a good starting point for the value. This is not a good recommendation because it is an arbitrary and fixed value that might not reflect the actual threshold for triggering an alert. Using 10 (per hour) as the value might result in false positives or false negatives, depending on the network conditions and switches. For example, if the normal number of rejects per hour is 5, then using 10 as the value might trigger an alert too frequently and unnecessarily. On the other hand, if the normal number of rejects per hour is 15, then using 10 as the value might miss some important alerts and risks.

QUESTION 7

You are configuring gateway IDS/IPS settings in Aruba Central.

For which reason would you set the Fail Strategy to Bypass?

A. To permit traffic if the IPS engine falls to inspect It

B. To enable the gateway to honor the allowlist settings configured in IDS/IPS policies

C. To tell gateways to stop enforcing IDS/IPS policies if they lose connectivity to the Internet

D. To avoid wasting IPS engine resources on filtering traffic for unauthenticated clients

Correct Answer: A

The Fail Strategy is a configuration option for the IPS mode of inspection on Aruba gateways. It defines the action to be taken when the IPS engine crashes and cannot inspect the traffic. There are two possible options for the Fail Strategy: Bypass and Block1 If you set the Fail Strategy to Bypass, you are telling the gateway to allow the traffic to flow without inspection when the IPS engine fails. This option ensures that there is no disruption in the network connectivity, but it also exposes the network to potential threats that are not detected or prevented by the IPS engine1 If you set the Fail Strategy to Block, you are telling the gateway to stop the traffic flow until the IPS engine resumes inspection. This option ensures that there is no compromise in the network security, but it also causes a loss of network connectivity for the duration of the IPS engine failure1

QUESTION 8

Refer to the exhibit.

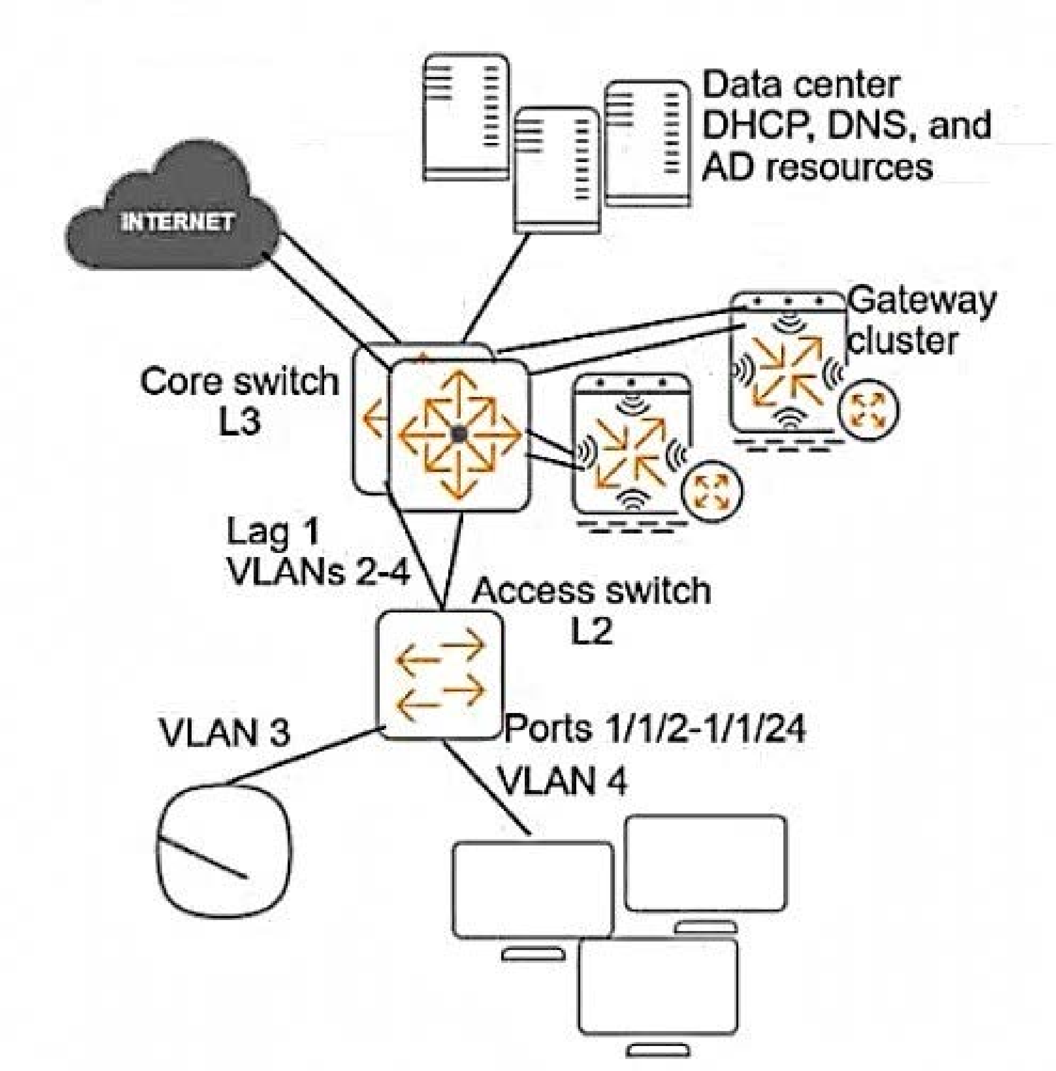

A customer requires protection against ARP poisoning in VLAN 4. Below are listed all settings for VLAN 4 and the VLAN 4 associated physical interfaces on the AOS-CX access layer switch:

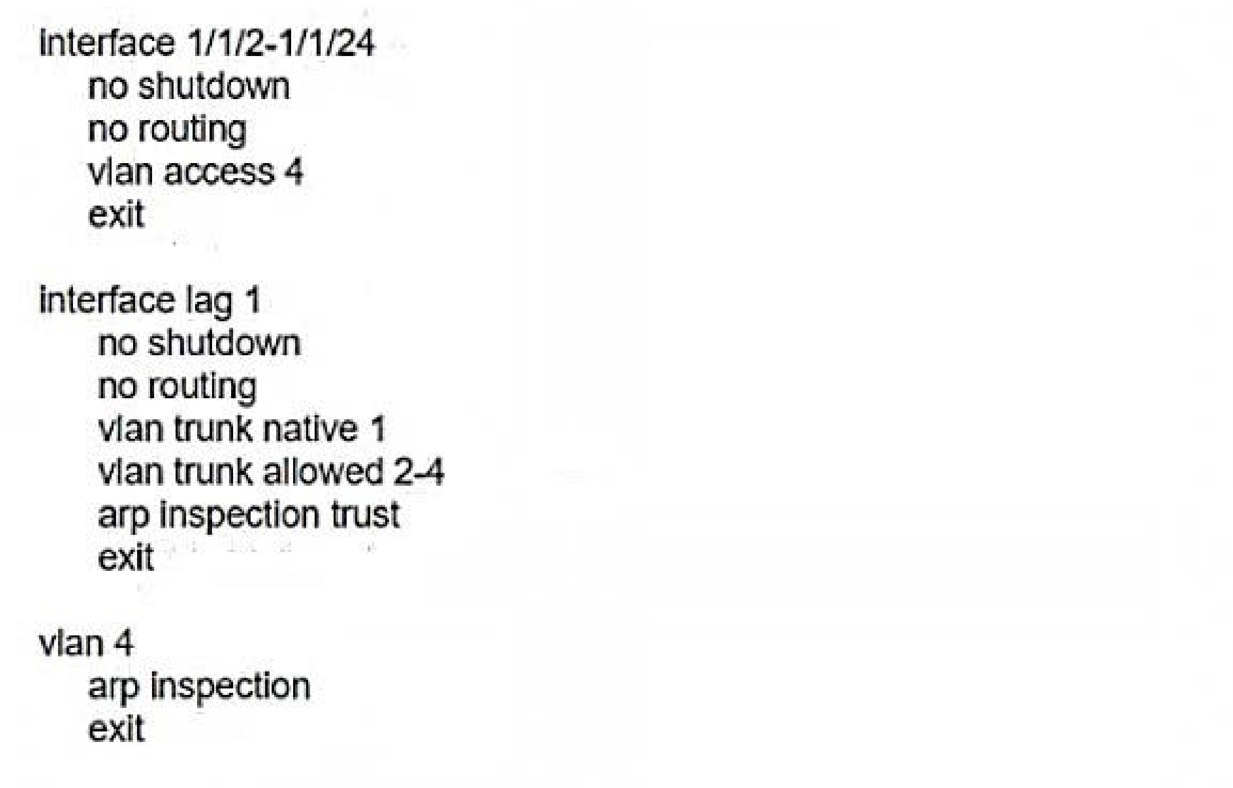

```
interface 1/1/2-1/1/24
    no shutdown
    no routing
    vlan access 4
    exit

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.

Correct Answer: D

This is because ARP inspection is a security feature that validates ARP packets in a network and prevents ARP poisoning attacks12 ARP inspection works by intercepting, logging, and discarding ARP packets with invalid IP-to-MAC address bindings1 To enable ARP inspection, the switch needs to know which ports are trusted and which are untrusted. Trusted ports are those that connect to authorized DHCP servers or other network devices that are not vulnerable to ARP spoofing. Untrusted ports are those that connect to end hosts or devices that might send forged ARP packets13 In the exhibit, LAG 1 is configured as a trusted port for ARP inspection, which is correct because it connects to the core switch. However, the edge ports (1/1/1-1/1/24) are not configured as untrusted ports for ARP inspection, which is incorrect because they connect to end hosts that might be compromised by an attacker. By default, all ports are untrusted for ARP inspection, but this can be changed by using the command ip arp inspection trust on the interface configuration mode3 Therefore, to protect VLAN 4 against ARP poisoning, the edge ports should be configured as untrusted for ARP inspection by using the command no ip arp inspection trust on the interface configuration mode. This way, the switch will validate the ARP packets received on these ports against the DHCP snooping database or an ARP access-list and drop any invalid packets34 A. ARP proxy is not enabled on VLAN 4. This is not an issue because ARP proxy is an optional feature that allows the switch to respond to ARP requests on behalf of hosts in different subnets5 It is not related to ARP poisoning or ARP inspection. B. LAG 1 is configured as trusted for ARP inspection but should be untrusted. This is not an issue because LAG 1 connects to the core switch, which is a trusted device that does not send forged ARP packets.

C. DHCP snooping is not enabled on VLAN 4. This is not an issue because DHCP snooping is a separate feature that prevents rogue DHCP servers from offering IP addresses to clients6 It is not directly related to ARP poisoning or ARP inspection, although it can provide information for ARP inspection validation if enabled

QUESTION 9

A customer\\'s admins have added RF Protect licenses and enabled WIDS for a customer\\'s AOS 8-based solution. The customer wants to use the built-in capabilities of APs without deploying dedicated air monitors (AMs). Admins tested rogue AP detection by connecting an unauthorized wireless AP to a switch. The rogue AP was not detected even after several hours.

What is one point about which you should ask?

A. Whether APs\\' switch ports support all the VLANs that are accessible at the edge

B. Whether admins enabled wireless containment

C. Whether admins set at least one radio on each AP to air monitor mode

D. Whether the customer is using non-standard Wi-Fi channels in the deployment

Correct Answer: C

RF Protect is a feature that enables wireless intrusion detection and prevention system (WIDS/WIPS) capabilities on AOS 8-based solutions. WIDS/WIPS allows detecting and mitigating rogue APs, unauthorized clients, and other wireless threats. RF Protect requires RF Protect licenses to be installed and WIDS to be enabled on the Mobility Master (MM). To use the built-in capabilities of APs for WIDS/WIPS, without deploying dedicated air monitors (AMs), admins need to set at least one radio on each AP to air monitor mode. Air monitor mode allows the AP to scan the wireless spectrum and report any wireless activity or anomalies to the MM. Air monitor mode does not affect the other radio on the AP, which can still serve clients in access mode. By setting at least one radio on each AP to air monitor mode, admins can achieve full coverage and visibility of the wireless environment and detect rogue APs. If admins do not set any radio on the APs to air monitor mode, the APs will not scan the wireless spectrum or report any wireless activity or anomalies to the MM. This means that the APs will not be able to detect rogue APs, even if they are connected to the same network. Therefore, admins should check whether they have set at least one radio on each AP to air monitor mode.

QUESTION 10

When would you implement BPDU protection on an AOS-CX switch port versus BPDU filtering?

A. Use BPDU protection on edge ports to protect against rogue devices when the switch implements MSTP; use BPDU filtering to protect against rogue devices when the switch implements PVSTP+.

B. Use BPDU protection on edge ports to prevent rogue devices from connecting; use BPDU filtering on inter-switch ports for specialized use cases.

C. Use BPDU protection on inter-switch ports to ensure that they are selected as root; use BPDU filtering on edge ports to prevent rogue devices from connecting.

D. Use BPDU protection on edge ports to permanently lock out rogue devices; use BPDU filtering on edge ports to temporarily lock out rogue devices.

Correct Answer: B

BPDU (Bridge Protocol Data Unit) is a message that is exchanged between switches to maintain the spanning tree topology and prevent loops. BPDU protection and BPDU filtering are two features that can be configured on AOS-CX switch ports to enhance security and performance. BPDU protection is a feature that disables a port if it receives a BPDU, indicating that an unauthorized switch or device has been connected to the port. BPDU protection is typically used on edge ports, which are ports that connect to end devices such as PCs or printers, and are not expected to receive BPDUs. BPDU protection prevents rogue devices from connecting to the network and affecting the spanning tree topology. BPDU filtering is a feature that prevents a port from sending or receiving BPDUs, effectively isolating the port from the spanning tree topology. BPDU filtering is typically used on inter-switch ports, which are ports that connect to other switches, for specialized use cases such as creating a separate spanning tree domain or reducing the overhead of BPDUs. BPDU filtering should be used with caution, as it can create loops or inconsistencies in the network. You can find more information about how to configure BPDU protection and BPDU filtering on AOS-CX switch ports in the [Configuring Spanning Tree Protocol - Aruba] page and the [AOS-CX Switching Configuration Guide] page. The other options are not correct because they either use BPDU protection or BPDU filtering on the wrong type of ports or for the wrong purpose. For example, using BPDU protection on inter-switch ports would disable the ports if they receive BPDUs, which are expected in normal operation. Using BPDU filtering on edge ports would allow rogue devices to connect to the network and create loops or affect the spanning tree topology.