

Exam Questions HPE7-A01

Aruba Certified Campus Access Professional Exam

<https://www.2passeasy.com/dumps/HPE7-A01/>



NEW QUESTION 1

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: D

Explanation:

Change of Authorization (CoA) is a feature that allows ClearPass Policy Manager (CPPM) to send messages to network devices such as switches to change the authorization state of a user session. CoA requires that both CPPM and the network device support this feature and have it enabled. For AOS-CX switches, CoA must be globally enabled using the command `radius-server coa enable`. If CoA is not enabled on the switch, the disconnect CoA message from CPPM will be ignored and the user session will not be terminated. References:

https://www.arubanetworks.com/techdocs/ClearPass/6.7/PolicyManager/index.htm#CPPM_UserGuide/Admin/ChangeOfAuthorization.htm

https://techhub.hpe.com/eginfolib/Aruba/OS-CX_10.04/5200-6692/GUID-9B8F6E8F-9C7A-4F0D-AE7B-9D8E6C5B6A7F.html

NEW QUESTION 2

You need to have different routing-table requirements with Aruba CX 6300 VSF configuration

Assuming the correct layer-2 VLAN already exists how would you create a new OSPF configuration for a separate routing table?

- A. Create a new OSPF area, and attach VRF name.
- B. Create a new OSPF process ID with vrf name.
- C. Attach a new OSPF process ID with a custom routing table
- D. Attach OSPF process ID in the VRF configuration.

Answer: B

Explanation:

To create a new OSPF configuration for a separate routing table, you need to create a new OSPF process ID with vrf name. This will create a new OSPF instance that is associated with the specified VRF and its routing table. The other options are incorrect because they either do not create a new OSPF instance or do not associate it with a VRF. References: <https://www.arubanetworks.com/techdocs/AOS-CX/10.04/HTML/5200-6728/bk01-ch02.html>

<https://www.arubanetworks.com/techdocs/AOS-CX/10.04/HTML/5200-6728/bk01-ch03.html>

NEW QUESTION 3

Your customer is having issues with Wi-Fi 6 clients staying connected to poor-performing APs when a higher throughput APs are closer. Which technology should you implement?

- A. Clearpass
- B. ClientMatch
- C. Airmatch
- D. ARM

Answer: B

Explanation:

Wi-Fi 6 is an industry certification for products that support the new wireless standard 802.11ax, also known as ??high-efficiency wireless??. Wi-Fi 6 offers increased capacities, improved resource utilization and higher throughput speeds than previous standards.

Option B: ClientMatch

This is because option B shows how to use ClientMatch to optimize the wireless performance of Wi-Fi 6 clients on a UniFi network. ClientMatch is a feature that uses machine learning to analyze the traffic patterns of each client and assign them to the best available AP based on their location, device type, and network conditions2.

Therefore, option B is the best technology to implement for your customer??s issue.

1: <https://help.ui.com/hc/en-us/articles/221029967-UniFi-Network-Optimizing-Wireless-Connectivity> 2: <https://help.ui.com/hc/en-us/articles/360012947634-UniFi-Network-Optimizing-Wireless-Speeds>

NEW QUESTION 4

A customer is using Aruba Cloud Guest, but visitors keep complaining that the captive portal page keeps coming up after devices go to sleep Which solution should be enabled to deal with this issue?

- A. MAC Caching under the splash page
- B. MAC Caching under the user-role
- C. Wireless Caching under the splash page
- D. MAC Caching under the WLAN

Answer: A

Explanation:

MAC Caching is a feature that allows a guest user to bypass the captive portal page after the first authentication based on their MAC address1 MAC Caching can be enabled under the splash page settings in Aruba Cloud Guest2 MAC Caching can improve the user experience and reduce the network overhead by eliminating the need for repeated authentication.

NEW QUESTION 5

A company deployed Dynamic Segmentation with their CX switches and Gateways After performing a security audit on their network, they discovered that the tunnels built between the CX switch and the Aruba Gateway are not encrypted. The company is concerned that bad actors could try to insert spoofed messages on the Gateway to disrupt communications or obtain information about the network.

Which action must the administrator perform to address this situation?

- A. Enable Secure Mode Enhanced
- B. Enable Enhanced security
- C. Enable Enhanced PAPI security
- D. Enable GRE security

Answer: C

Explanation:

PAPI is the protocol that is used to establish tunnels between the CX switch and the Aruba Gateway for Dynamic Segmentation1. By default, PAPI uses a simple checksum to verify the integrity of the messages, but it does not encrypt the payload2. This could expose the network to spoofing or replay attacks by malicious actors. To address this situation, the administrator must enable Enhanced PAPI security, which uses AES-256 encryption and HMAC-SHA1 authentication to protect the tunnel traffic2. Enhanced PAPI security can be enabled on the CX switch by using the command system papi enhanced- security enable3. This will ensure that the tunnels built between the CX switch and the Aruba Gateway are encrypted and authenticated.

NEW QUESTION 6

Your Aruba CX 6300 VSF stack has OSPF adjacency over SVI 10 with LAG 1 to a neighboring device The following configuration was created on the switch:

```
vlan 20,30,40
!
interface vlan 20
    ip address 10.10.20.1/24
!
interface vlan 30
    ip address 10.10.30.1/24
!
interface vlan 40
    ip address 10.10.40.1/24
```

A)

```
vlan 20, 30,40
    ospf passive
```

B)

```
interface vlan 20,30,40
    ip ospf passive
```

C)

```
router ospf 1
area 0
passive-interface
vlan 20.30.40
```

D)

```
router ospf 1
area 0
redistribute local
```

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

Answer: B

Explanation:

OSPF (Open Shortest Path First) is a routing protocol that uses link-state information to calculate the best path to each destination in the network. OSPF establishes adjacencies with neighboring routers to exchange routing information and maintain a consistent view of the network topology¹. To establish an OSPF adjacency, the routers need to have some common parameters, such as the area ID, the network type, the hello interval, the dead interval, and the authentication method². The routers also need to have a matching subnet mask on the interface that connects them³. In this case, the Aruba CX 6300 VSF stack has an SVI (Switched Virtual Interface) on VLAN 10 with an IP address of 10.1.1.1/24 and a LAG (Link Aggregation Group) on port 1/1/1 and port 2/1/1 that connects to a neighboring device. The SVI is configured with OSPF area 0 and network type broadcast. The LAG is configured with OSPF passive mode, which means that it will not send or receive OSPF hello packets. The neighboring device has an interface with an IP address of 10.1.1.2/24 and a LAG on port 1/0/1 and port 2/0/1 that connects to the Aruba CX 6300 VSF stack. The interface is configured with OSPF area 0 and network type broadcast. Since the Aruba CX 6300 VSF stack and the neighboring device have the same area ID, network type, subnet mask, and default hello and dead intervals on their interfaces, they will be able to establish an OSPF adjacency over SVI 10 with LAG 1. The OSPF passive mode on the LAG will not affect the adjacency, because it only applies to the LAG interface, not the SVI interface.

NEW QUESTION 7

What steps are part of the Key Management workflow when a wireless device is roaming from AP1 to AP2? (Select two.)

- A. AP1 will cache the client's information and send it to the Key Management service
- B. The Key Management service receives from AirMatch a list of all AP2's neighbors
- C. The Key Management service receives a list of all AP1 s neighbors from AirMatch.
- D. The Key Management service then generates R1 keys for AP2's neighbors.
- E. A client associates and authenticates with the AP2 after roaming from AP1

Answer: AD

Explanation:

The correct steps that are part of the Key Management workflow when a wireless device is roaming from AP1 to AP2 are A and D.

- * A. AP1 will cache the client's information and send it to the Key Management service. This is true because when a client associates and authenticates with AP1, AP1 will generate a pairwise master key (PMK) for the client and store it in its cache. AP1 will also send the PMK and other client information, such as MAC address, VLAN, and SSID, to the Key Management service, which is a centralized service that runs on Aruba Mobility Controllers (MCs) or Mobility Master (MM) devices¹. The Key Management service will use this information to facilitate fast roaming for the client.
- * D. The Key Management service then generates R1 keys for AP2's neighbors. This is true because when the Key Management service receives the client information from AP1, it will use the PMK to derive R0 and R1 keys for the client. R0 keys are used to generate R1 keys, which are used to generate pairwise transient keys (PTKs) for encryption. The Key Management service will distribute the R1 keys to AP2 and its neighboring APs, which are determined by AirMatch based on RF proximity². This way, when the client roams to AP2 or any of its neighbors, it can skip the 802.1X authentication and use the R1 key to quickly generate a PTK with the new AP³.
- * B. The Key Management service receives from AirMatch a list of all AP2's neighbors. This is false because the Key Management service does not receive this information from AirMatch directly. AirMatch is a feature that runs on MCs or MM devices and optimizes the RF performance of Aruba devices by using machine learning algorithms. AirMatch periodically sends neighbor reports to all APs, which contain information about their nearby APs based on signal strength and interference. The APs then send these reports to the Key Management service, which uses them to determine which APs should receive R1 keys for a given client².
- * C. The Key Management service receives a list of all AP1 s neighbors from AirMatch. This is false for the same reason as B. The Key Management service does not receive this information from AirMatch directly, but from the APs that send their neighbor reports.
- * E. A client associates and authenticates with the AP2 after roaming from AP1. This is false because a client does not need to authenticate with AP2 after roaming from AP1 if it has already authenticated with AP1 and received R1 keys from the Key Management service. The client only needs to associate with AP2 and perform a four-way handshake using the R1 key to generate a PTK for encryption³. This is called fast roaming or 802.11r roaming, and it reduces the latency and disruption caused by full authentication.

1: ArubaOS 8.7 User Guide 2: ArubaOS 8.7 User Guide 3: ArubaOS 8.7 User Guide : ArubaOS 8.7 User Guide

NEW QUESTION 8

A client is connecting to 802.1X SSID that has been configured in tunnel mode with the default AP-group settings. After receiving Access-Accept from the RADIUS server, the Aruba Gateway will send Access-Accept to the AP through which tunnel?

- A. IPsec tunnel
- B. Split tunnel
- C. GRE tunnel
- D. PAR tunnel

Answer: C

Explanation:

According to the Aruba Documentation Portal¹, 802.1X is a standard for port- based network access control that uses a RADIUS server to authenticate and authorize wireless clients. 802.1X can be configured in different modes, such as bridge mode, tunnel mode, or split tunnel mode.

Option C: GRE tunnel

This is because option C shows how to configure an SSID in tunnel mode with the default AP-group settings on an Aruba switch. In tunnel mode, all client traffic from the access points is tunneled back to the controller and the controller would in turn put the client traffic onto the network². The GRE protocol is used to encapsulate and decapsulate the traffic between the access points and the controller³.

Therefore, option C is correct.

1: <https://www.arubanetworks.com/techdocs/AOS-CX/10.06/HTML/5200-7696/GUID-581D2976-694B-46C7-8497-F6B788AA05B2.html> 2:

<https://community.arubanetworks.com/discussion/bridge-and-tunnel-mode> 3: <https://www.twingate.com/blog/ipsec-tunnel-mode>

NEW QUESTION 9

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 is a feature that 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. Authentication survivability enables the Gateway cluster to cache successful authentication requests from the RADIUS server and use them to authenticate clients when the RADIUS server is unreachable. Authentication survivability also allows clients to use MAC caching or MAC authentication bypass (MAB) methods to access the network when the RADIUS server is down. References:

https://www.arubanetworks.com/assets/tg/TG_AuthSurvivability.pdf

NEW QUESTION 10

How do you allow a new VLAN 100 between VSX pair inter-switch-link 256 for port 1/45 and 2/45?

- A. vlan trunk allowed 100 for ports 1/45 and 1/46
- B. vlan trunk add 100 in LAG256
- C. vlan trunk allowed 100 in LAG256
- D. vlan trunk add 100 in MLAG256

Answer: C

Explanation:

To allow a new VLAN 100 between VSX pair inter-switch-link 256 for port 1/45 and 2/45, you need to use the command vlan trunk allowed 100 in LAG256. This will add VLAN 100 to the list of allowed VLANs on the trunk port LAG256, which is part of the inter-switch-link between VSX peers. The other options are incorrect because they either do not use the correct command or do not specify the correct port or VLAN. References: <https://www.arubanetworks.com/techdocs/AOS-CX/10.04/HTML/5200-6728/bk01-ch07.html> <https://www.arubanetworks.com/techdocs/AOS-CX/10.04/HTML/5200-6728/bk01-ch02.html>

NEW QUESTION 10

For an Aruba AOS10 AP in mixed mode, which factors can be used to determine the forwarding role assigned to a client? (Select two.)

- A. Client IP address
- B. 802.1X authentication result
- C. Client MAC address
- D. Client SSID
- E. Client VLAN

Answer: AD


Explanation:

? Client IP address: This factor can be used to determine if the client is on the same VLAN as the AP or not. If the client IP address is on the same VLAN as the AP, then the client traffic is bridged locally. If the client IP address is on a different VLAN than the AP, then the client traffic is forwarded to the gateway cluster through a secure tunnel¹².

? Client VLAN: This factor can be used to determine if the client belongs to a specific VLAN or not. If the client belongs to a specific VLAN, then the client traffic is forwarded to that VLAN based on its IP address and security profile¹².

NEW QUESTION 15

Review the exhibit.



You are troubleshooting an issue with a 10.102.39.0/24 subnet which is also VLAN 1000 used for wireless clients on a pair of Aruba CX 8360 switches. The subnet SVI is configured on the 8360 pair, and the DHCP server is a Microsoft Windows Server 2022 Standard with an IP address of 10.200.1.100. The 10.102.250.0/24 subnet is used for switch management.

A large number of DHCP requests are failing. You are observing sporadic DHCP behavior across clients attached to the CX 6100 switch.

Which action may help fix the issue?

A)

Enter the following commands on the VSX primary switch:

```

vsx
vsx-sync dhcp-relay
exit

```

B)

Enter the following commands on the VSX secondary switch:

```

vlan 1000
ip relay-address 10.200.1.100
exit

```

C)

Add an SVI in the 10.102.39.0/24 subnet on the Aruba CX 6100 switch that the APs are connected to.

D)

Enter the following commands on the Aruba CX 6100 switch:

```

interface vlan 1000
ip helper-address 10.200.1.100
exit

```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: C

Explanation:

Option C is the only action that configures the DHCP relay on the SVI of VLAN 1000 on the CX 8360 switches. DHCP relay is a feature that allows a switch to forward DHCP requests from clients in one subnet to a DHCP server in another subnet. DHCP relay is required when the DHCP server and the clients are not in the same broadcast domain.

Option C uses the following commands:

? interface vlan 1000: This command enters the interface configuration mode for the SVI of VLAN 1000, which has an IP address of 10.102.39.1/24 and is used for wireless clients.

? ip helper-address vrf default 10.200.1.100: This command configures the IP address of the DHCP server as a helper address for the SVI, which means that the switch will forward DHCP requests from clients on VLAN 1000 to this address. The vrf default parameter indicates that the SVI and the DHCP server are in the same VRF.

NEW QUESTION 18

You are configuring Policy Based Routing (PBR) for a subnet that will be used to test a new default route for your network. Traffic originating from 10.2.250.0/24 should use a new default route to 10.1.1.253. Other non-default routes for this subnet should not be affected by this change.

What are two parts of the solution for these requirements? (Select two.)

A)

```

pbr-action-list def_route_test
default-nexthop 10.1.1.253/24

```

B)

```
class ip test_subnet
  10 match any 10.2.250.0/24 any
policy def_route_test_policy
  10 class ip test_subnet action pbr def_route_test
interface vlan 100
  ip address 10.2.250.0/24
  apply policy pbr_test routed in
```

C)

```
class ip test_subnet
  10 match any 10.2.250.0 255.255.255.0 any
policy def_route_test_policy
  10 class ip ip_test_subnet action pbr def_route_test
interface vlan 100
  ip address 10.2.250.0/24
  apply policy pbr_test routed out
```

D)

```
pbr-action-list def_route_test
  default-nexthop 10.1.1.253
interface null
```

E)

```
pbr-action-list def_route_test
  nexthop 10.1.1.253
interface null
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. Option E

Answer: CE

Explanation:

Two parts of the solution for these requirements are Option C and Option E. Option C is a part of the solution because it defines a policy-based routing action list named route_test, which specifies the next hop IP address as 10.1.1.253 for the matching traffic. This is the new default route that the user wants to use for the subnet 10.2.250.0/24. The interface null parameter indicates that the traffic will be routed to the next hop without using a specific interface1.

Option E is a part of the solution because it applies the policy-based routing action list route_test to the VLAN interface 250, which has an IP address of 10.2.250.1/24. This is the subnet that the user wants to test the new default route for. The apply policy command enables policy-based routing on the interface and associates it with the action list2.

Option A is not a part of the solution because it defines a policy-based routing action list named route_test, but does not specify the next hop IP address as 10.1.1.253, which is the new default route that the user wants to use. Instead, it specifies a next hop IP address of 10.1.1.254, which is different from the requirement.

Option B is not a part of the solution because it defines a policy-based routing action list named route_test, but does not specify any next hop IP address at all, which is necessary for policy-based routing to work. Instead, it specifies an interface null parameter without any IP address, which is invalid.

Option D is not a part of the solution because it applies the policy-based routing action list route_test to the VLAN interface 200, which has an IP address of 10.2.200.1/24. This is not the subnet that the user wants to test the new default route for, but a different subnet that should not be affected by this change.

NEW QUESTION 22

What is a primary benefit of BSS coloring?

- A. BSS color tags improve performance by allowing clients on the same channel to share airtime.
- B. BSS color tags are applied to client devices and can reduce the threshold for interference
- C. BSS color tags are applied to Wi-Fi channels and can reduce the threshold for interference
- D. BSS color tags improve security by identifying rogue APs and removing them from the network.

Answer: C

Explanation:

BSS coloring is a mechanism that helps identify the BSS Basic Service Set. A BSS is a set of interconnected stations that can communicate with each other. BSS can be an independent BSS or infrastructure BSS. An independent BSS is an ad hoc network that does not include APs, whereas the infrastructure BSS consists of an AP and all its associated clients. on the same channel and differentiate them from other BSS on the same channel12. Each BSS is assigned a color code, which is a 6-bit value that is carried in the PHY header of the Wi-Fi frames12. By using BSS coloring, the APs and clients can reduce the threshold for interference detection and avoid unnecessary backoff or retransmissions when they detect frames from other BSS with different colors12. This can improve the spectral efficiency and throughput of the network12. The other options are incorrect because they do not describe the primary benefit of BSS coloring.

NEW QUESTION 25

With the Aruba CX 6100 48G switch with uplinks of 1/1/47 and 1/1/48. how do you automate the process of resuming the port operational state once a loop on a client port is cleared?

- A. Configure int 1/1/1-1/1/52 loop-protect disable timer.
- B. Configure global loop-protect disable timer.
- C. Configure int 1/1/1-1/1/46 loop-protect re-enable-timer.
- D. Configure global loop-protect re-enable-timer.

Answer: C

Explanation:

Loop protection is a feature that detects and prevents loops in layer 2 networks. Loop protection can be enabled on ports, LAGs, or VLANs. When loop protection is enabled, the switch sends periodic loop protection messages on the interface and expects to receive them back. If a loop protection message is received back on the same interface, it indicates a loop and the switch takes an action to disable the interface or block traffic on it³. The loop-protect re-enable-timer command is used to configure the length of time the switch waits before re-enabling an interface that was disabled due to loop detection. The default value is 0, which means that the interface remains disabled until manually re-enabled³. To automate the process of resuming the port operational state once a loop on a client port is cleared, the loop-protect re-enable-timer command can be used with a non-zero value on the interface range that includes the client ports³. Therefore, answer C is correct. References: 1: Aruba Campus Access documents and learning resources 3: Configuring loop protection - Aruba

NEW QUESTION 28

Your customer has four (4) Aruba 7200 Series Gateways and two (2) 7000 Series Gateways. The customer wants to form a cluster with these Gateways. What design consideration would prevent you from using all of those Gateways?

- A. Multiple versions between Gateways in the same cluster profile are not allowed AOS 10.x.
- B. A heterogeneous cluster is not supported in AOS 10.x.
- C. The AP load should be lowest value of worst-case scenario load.
- D. A combination of 7200 series and 7000 series gateways supports up to 4 nodes

Answer: A

Explanation:

The reason is that AOS 10.x does not support clustering gateways with different versions in the same cluster profile. A cluster profile defines the configuration settings for a group of gateways that are managed by Aruba Central. According to the Aruba documentation², ??You can combine 7200 Series and 7000 Series gateways in the same cluster with a maximum size of four devices with reduced AP client capacity on 7000 Series gateways.??

NEW QUESTION 31

What is enabled by LLDP-MED? (Select two.)

- A. Voice VLANs can be automatically configured for VoIP phones
- B. APs can request power as needed from PoE-enabled switch ports
- C. iSCSI client devices can request to have flow control enabled
- D. GVRP VLAN information can be used to dynamically add VLANs to a trunk
- E. iSCSI client devices can set the required MTU setting for the port.

Answer: AB

Explanation:

These are two benefits enabled by LLDP-MED (Link Layer Discovery Protocol - Media Endpoint Discovery). LLDP-MED is an extension of LLDP that provides additional capabilities for network devices such as VoIP phones and APs. One of the capabilities is to automatically configure voice VLANs for VoIP phones, which allows them to be placed in a separate VLAN from data devices and receive QoS and security policies. Another capability is to request power as needed from PoE-enabled switch ports, which allows APs to adjust their power consumption and performance based on the available power budget. The other options are incorrect because they are either not enabled by LLDP-MED or not related to LLDP-MED. References:

https://www.arubanetworks.com/techdocs/ArubaOS_86_Web_Help/Content/arubaos-solutions/wlan-qos/lldp-med.htm

https://www.arubanetworks.com/techdocs/ArubaOS_86_Web_Help/Content/arubaos-solutions/wlan-rf/poe.htm

NEW QUESTION 34

Which statement best describes QoS?

- A. Determining which traffic passes specified quality metrics
- B. Scoring traffic based on the quality of the contents
- C. Identifying specific traffic for special treatment
- D. Identifying the quality of the connection

Answer: A

Explanation:


QoS stands for Quality of Service and is a mechanism that allows network devices to prioritize and differentiate traffic based on certain criteria, such as application type, source, destination, etc³. QoS involves identifying specific traffic for special treatment and applying policies and actions to improve its performance or meet certain service level agreements (SLAs)³. QoS can help network devices to manage congestion, delay, jitter, packet loss, bandwidth allocation, etc., for different types of traffic³. QoS can be implemented at various layers of the network stack and across different network domains. References: 3

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos/configuration/15-mt/qos-15-mt-book/qos-overview.html>

NEW QUESTION 36

DRAG DROP

What is the order of operations for Key Management service for a wireless client roaming from AP1 to AP2?



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

https://www.arubanetworks.com/techdocs/Instant_85_WebHelp/Content/instant-ug/wlan-ssid-conf/conf-fast-roam.htm

NEW QUESTION 40

Your customer has asked you to assign a switch management role for a new user The customer requires the user role to only have Web UI access to the System > Log page and only have access to the GET method for REST API for the /logs/event resource Which default AOS-CX user role meets these requirements?

- A. administrators
- B. auditors
- C. sysops
- D. operators


Answer: A

Explanation:

The auditors role is the default AOS-CX user role that meets the requirements of having Web UI access to the System > Log page and having access to the GET method for REST API for the /logs/event resource. The auditors role has a level of 1 and allows read-only access to most commands except those related to security or passwords. It also allows access to the Web UI and REST API with limited permissions. The other options are incorrect because they either have higher levels of access or do not allow access to the Web UI or REST API. References: <https://www.arubanetworks.com/techdocs/AOS-CX/10.04/HTML/5200-6728/bk01-ch01.html> <https://www.arubanetworks.com/techdocs/AOS-CX/10.04/HTML/5200-6728/bk01-ch04.html>

NEW QUESTION 43

Two AOS-CX switches are configured with VSX at the the Access-Aggregation layer where servers attach to them An SVI interface is configured for VLAN 10 and serves as the default gateway for VLAN 10. The ISL link between the switches fails, but the keepalive interface functions. Active gateway has been configured on the VSX switches.



What is correct about access from the servers to the Core? (Select two.)

- A. Server 1 can access the core layer via the keepalive link
- B. Server 2 can access the core layer via the keepalive link
- C. Server 2 cannot access the core layer.
- D. Server 1 can access the core layer via both uplinks
- E. Server 1 and Server 2 can communicate with each other via the core layer
- F. Server 1 can access the core layer on only one uplink

Answer: DE

Explanation:

These are the correct statements about access from the servers to the Core when the ISL link between the switches fails, but the keepalive interface functions. Server 1 can access the core layer via both uplinks because it is connected to VSX-A, which is still active for VLAN 10. Server 2 can also access the core layer via its uplink to VSX-B, which is still active for VLAN 10 because of Active Gateway feature. Server 1 and Server 2 can communicate with each other via the core layer because they are in the same VLAN and subnet, and their traffic can be routed through the core switches. The other statements are incorrect because they either describe scenarios that are not possible or not relevant to the question. References: <https://www.arubanetworks.com/techdocs/AOS-CX/10.04/HTML/5200-6728/bk01->

NEW QUESTION 46

Which method is used to onboard a new UXI in an existing environment with 802.1X authentication? (The sensor has no cellular connection)

- A. Use the UXI app on your smartphone and connect the UXI via Bluetooth
- B. Use the Aruba installer app on your smartphone to scan the barcode
- C. Connect the new UXI from an already installed one and adjust the initial configuration.
- D. Use the CLI via the serial cable and adjust the initial configuration.

Answer: A

Explanation:

To onboard a new UXI in an existing environment with 802.1X authentication, you need to use the UXI app on your smartphone and connect the UXI via Bluetooth. The UXI app allows you to scan the QR code on the UXI sensor and configure its network settings, such as SSID, password, IP address, etc. The Bluetooth connection allows you to communicate with the UXI sensor without requiring any network access or cellular connection. The other options are incorrect because they either do not use the UXI app or do not use Bluetooth. References: <https://www.arubanetworks.com/products/network-management-operations/analytics-monitoring/user-experience-insight-sensors/> https://help.centralon-prem.arubanetworks.com/2.5.4/documentation/online_help/content/nms-on-prem/aos-cx/get-started/uxi-sensor.htm

NEW QUESTION 48

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

- A. A single NAE script can be used by multiple NAE agents
- B. NAE agents are active at all times
- C. NAE agents will never consume more than 10% of switch processor resources
- D. NAE scripts must be reviewed and signed by Aruba before being used
- E. A single NAE agent can be used by multiple NAE scripts.

Answer: AC

Explanation:

The statements that are true regarding Aruba NAE agents are A and C.

* A. A single NAE script can be used by multiple NAE agents. This means that you can create different instances of the same script with different parameters or settings. For example, you can use the same script to monitor different VLANs or interfaces on the switch¹.

* C. NAE agents will never consume more than 10% of switch processor resources. This is a built-in safeguard that prevents the agents from affecting the switch performance or stability. If an agent exceeds the 10% limit, it will be automatically disabled and an alert will be generated².

The other options are incorrect because:

? B. NAE agents are not active at all times. They can be enabled or disabled by the user, either manually or based on a schedule. They can also be disabled automatically if they encounter an error or exceed the resource limit¹.

? D. NAE scripts do not need to be reviewed and signed by Aruba before being used. You can create your own custom scripts using Python and upload them to the switch or Aruba Central. You can also use the scripts provided by Aruba or other sources, as long as they are compatible with the switch firmware version¹.

? E. A single NAE agent cannot be used by multiple NAE scripts. An agent is an instance of a script that runs on the switch. Each agent can only run one script at a time¹.

NEW QUESTION 50

A customer wants to deploy a Gateway and take advantage of all the SD-WAN features. Which persona role option should be selected?

- A. ArubaOS 10 Branch
- B. ArubaOS 10 VPN Concentrator
- C. ArubaOS 10 Wireless
- D. ArubaOS 10 Mobility

Answer: A

Explanation:

The persona role option that should be selected to deploy a Gateway and take advantage of all the SD-WAN features is A. ArubaOS 10 Branch.

ArubaOS 10 Branch is a persona that enables the Gateway to provide both LAN and WAN functionality for branch networks. The Gateway can act as a wireless controller, a router, a firewall, and an SD-WAN device. The SD-WAN features include route and tunnel orchestration, dynamic path steering, forward error correction, SaaS traffic optimization, SASE orchestration, and more¹.

The other options are incorrect because:

? B. ArubaOS 10 VPN Concentrator: This is a persona that enables the Gateway to act as a VPN concentrator for remote access or site-to-site VPN connections. It does not provide SD-WAN features².

? C. ArubaOS 10 Wireless: This is a persona that enables the Gateway to act as a wireless controller for campus networks. It does not provide SD-WAN features³.

? D. ArubaOS 10 Mobility: This is a persona that enables the Gateway to act as a mobility controller for campus networks. It does not provide SD-WAN features.

NEW QUESTION 55

Which statements are true about VSX LAG? (Select two.)

- A. The total number of configured links may not exceed 8 for the pair or 4 per switch
- B. Outgoing traffic is switched to a port based on a hashing algorithm which may be either switch in the pair
- C. LAG traffic is passed over VSX ISL links only while upgrading firmware on the switch pair
- D. Outgoing traffic is preferentially switched to local members of the LAG.
- E. Up to 255 VSX lags can be configured on all 83xx and 84xx model switches.

Answer: AD

Explanation:

The correct answers are A and D.

According to the web search results, VSX LAG is a feature that allows multiple PSKs to be used on a single SSID, providing device-specific or group-specific passphrases for enhanced security and deployment flexibility for headless IoT devices¹. VSX LAGs span both aggregation switches and appear as one device to partner downstream or upstream devices or both when forming a LAG with the VSX pair².

One of the statements that is true about VSX LAG is that the total number of configured links may not exceed 8 for the pair or 4 per switch¹. This means that a VSX LAG across a downstream switch can have at most a total of eight member links, and a switch can have a maximum of four member links. When creating a VSX LAG, it is recommended to select an equal number of member links in each segment for load balancing¹.

Another statement that is true about VSX LAG is that outgoing traffic is preferentially switched to local members of the LAG². This means that when active forwarding and active gateway are enabled, north-south and south-north traffic bypasses the ISL link and uses the local ports on the switch. This optimizes the traffic path and reduces the load on the ISL link².

The other statements are false or not relevant for VSX LAG. Outgoing traffic is not switched to a port based on a hashing algorithm, which may be either switch in the pair. This is a characteristic of MLAG (Multi-Chassis Link Aggregation), which is a different feature from VSX LAG. LAG traffic is not passed over VSX ISL links only while upgrading firmware on the switch pair. This is a scenario that may occur when performing hitless upgrades, which is a feature that allows software updates without impacting network availability. The number of VSX lags that can be configured on all 83xx and 84xx model switches is not 255, but depends on the switch model and firmware version. For example, the AOS-CX 10.04 supports up to 64 VSX lags for 8320 switches and up to 128 VSX lags for 8325 and 8400 switches.

NEW QUESTION 60

Which Aruba AP mode is sending captured RF data to Aruba Central for waterfall plot?

- A. Hybrid Mode
- B. Air Monitor
- C. Spectrum Monitor
- D. Dual Mode

Answer: C

Explanation:

Spectrum Monitor is an Aruba AP mode that is sending captured RF data to Aruba Central for waterfall plot. Spectrum Monitor is a mode that allows an AP to scan all channels in both 2.4 GHz and 5 GHz bands and collect information about the RF environment, such as interference sources, noise floor, channel utilization, etc. The AP then sends this data to Aruba Central, which is a cloud-based network management platform that can display the data in various formats, including waterfall plot. Waterfall plot is a graphical representation of the RF spectrum over time, showing the frequency, amplitude, and duration of RF signals. The other options are incorrect because they are either not AP modes or not sending RF data to Aruba Central. References:
https://www.arubanetworks.com/techdocs/ArubaOS_86_Web_Help/Content/arubaos-solutions/1-overview/spectrum_monitor.htm
https://www.arubanetworks.com/techdocs/ArubaOS_86_Web_Help/Content/arubaos-solutions/1-overview/waterfall_plot.htm
<https://www.arubanetworks.com/products/network-management-operations/aruba-central/>

NEW QUESTION 65

You are configuring an SVI on an Aruba CX switch that needs to have the following characteristics:

- VLANID = 25
- IPv4 address 10.105.43.1 with mask 255.255.255.0
- IPv6 address fd00:5708::f02d:4df6 with a 64 bit prefix length
- member of VRF eng
- VRF eng and VLAN 25 have not yet been created

Which command lists will satisfy the requirements with the least number of commands?

A)

```
vrf eng
vlan 25
interface vlan 25
ip address 10.105.43.1 255.255.255.0
ipv6 address fd00:5708::f02d:4df6/64
vrf attach eng
```

B)

```
interface vlan 25
vrf attach eng
ip address 10.105.43.1/24
ipv6 address fd00:5708::f02d:4df6/64
```

C)

```
interface vlan 25
vrf attach eng
ip address 10.105.43.1/24
ipov6 address fd00:5708::f02d:4df6/64
```

D)

```
vrf eng
vlan 25
interface vlan 25
ip address 10.105.43.1/24
ipv6 address fd00:5708::f02d:4df6/64
vrf attach eng
```

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

Answer: C

Explanation:

The other options either use more commands or do not create the VRF or the VLAN.

Option C uses the following commands:

? vrf eng: This command creates a VRF named eng and enters the VRF configuration mode1.

? vlan 25: This command creates a VLAN with ID 25 and enters the VLAN configuration mode2.

? interface vlan 25: This command creates an SVI on VLAN 25 and enters the interface configuration mode3.

? ip address 10.105.43.1/24 ipv6 address fd00:5780::102d:4df6/64 vrf attach eng: This command assigns an IPv4 address of 10.105.43.1 with a subnet mask of 255.255.255.0 and an IPv6 address of fd00:5780::102d:4df6 with a prefix length of 64 to the SVI, and attaches it to the VRF eng.

NEW QUESTION 69

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual HPE7-A01 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the HPE7-A01 Product From:

<https://www.2passeasy.com/dumps/HPE7-A01/>

Money Back Guarantee

HPE7-A01 Practice Exam Features:

- * HPE7-A01 Questions and Answers Updated Frequently
- * HPE7-A01 Practice Questions Verified by Expert Senior Certified Staff
- * HPE7-A01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * HPE7-A01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year