HPE7-A01Q&As

Aruba Certified Campus Access Professional

QUESTION 1

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

Correct 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

QUESTION 2

You need to drop excessive broadcast traffic on an ingress port or an ArubaOS-CX switch. What is the best feature to use for this task?

A. DWRR queuing

B. Strict queuing

C. Rate limiting

D. QoS shaping

Correct Answer: C

Explanation: According to the Aruba Documentation Portal1, the ArubaOS-CX switch supports various features to control the ingress traffic on specific ports, such as rate limiting, QoS shaping, and access control. These features can help reduce the impact of excessive broadcast traffic on the network performance and availability. This is because rate limiting is a feature that allows you to limit the inbound or outbound traffic on a port based on a percentage of the port capacity or a fixed amount of bytes per second. Rate limiting can help prevent broadcast storms by reducing the amount of broadcast packets that enter or leave a port https://www.arubanetworks.com/techdocs/central/latest/content/nms/aos-cx/cfg/conf-cx- access-control.htm 2: https://community.arubanetworks.com/blogs/esupport1/2021/02/08/broadcast-storm- containment-in-aruba-pvos-switches 3: https://techhub.hpe.com/eginfolib/networking/docs/switches/K-KA-KB/15-18/5998- 8160_ssw_mcg/content/ch05.html

QUESTION 3

Match the solution components of NetConductor (Options may be used more than once or not at all.)

Select and Place:

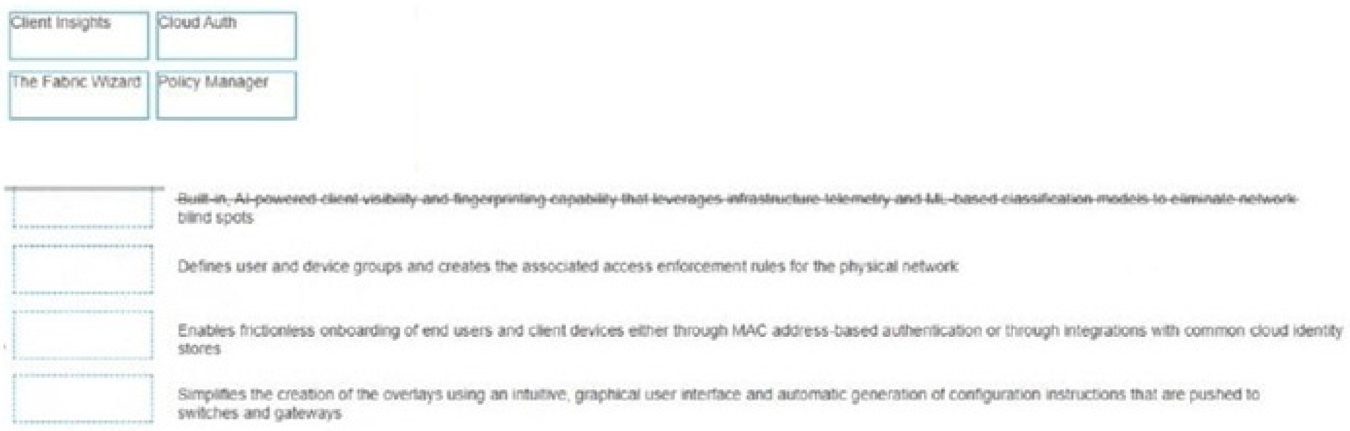

Client Insights	Cloud Auth
The Fabric Wizard	Policy Manager

	~~Built-in, AI-powered client visibility and fingerprinting capability that leverages infrastructure telemetry and ML-based classification models to eliminate network~~ blind spots
	Defines user and device groups and creates the associated access enforcement rules for the physical network
	Enables frictionless onboarding of end users and client devices either through MAC address-based authentication or through integrations with common cloud identity stores
	Simplifies the creation of the overlays using an intuitive, graphical user interface and automatic generation of configuration instructions that are pushed to switches and gateways

Correct Answer:

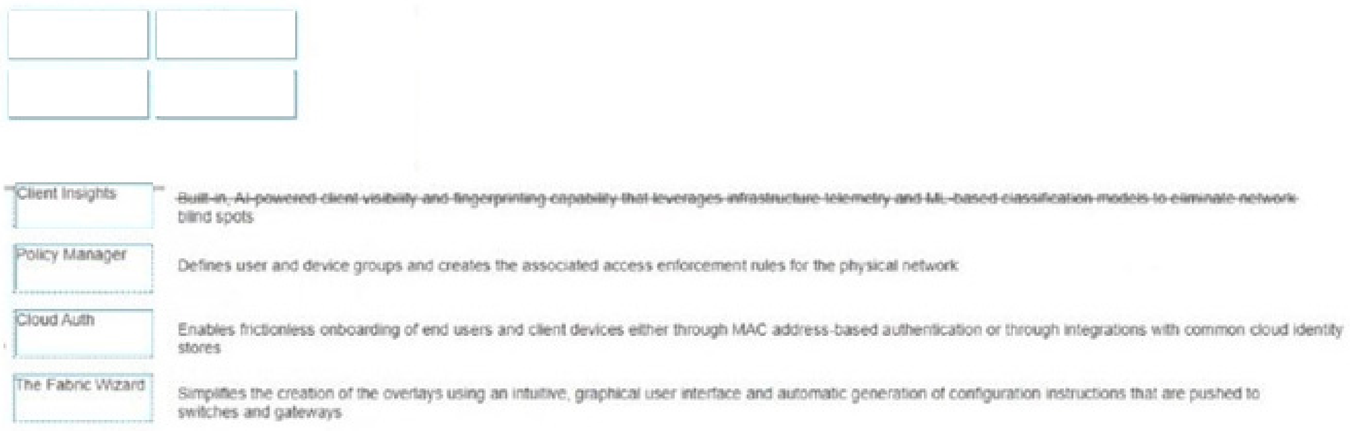

Client Insights	~~Built-in, AI-powered client visibility and fingerprinting capability that leverages infrastructure telemetry and ML-based classification models to eliminate network~~ blind spots
Policy Manager	Defines user and device groups and creates the associated access enforcement rules for the physical network
Cloud Auth	Enables frictionless onboarding of end users and client devices either through MAC address-based authentication or through integrations with common cloud identity stores
The Fabric Wizard	Simplifies the creation of the overlays using an intuitive, graphical user interface and automatic generation of configuration instructions that are pushed to switches and gateways

Client Insights matches with Built in , AI powered client visibility and fingerprinting capability that leverages infrastructure telemetry and ML based classification models to eliminate network bling spots

Client Insights is a solution component of NetConductor that provides built-in, AI-powered client visibility and fingerprinting capability that leverages infrastructure telemetry and MLbased classification models to eliminate network blind spots.

Client Insights uses machine learning to automatically detect, identify, and classify devices on the network, such as IoT devices, BYOD devices, or rogue devices. Client Insights also provides behavioral analytics and anomaly detection to

monitor device performance and security posture. Client Insights helps network administrators gain visibility into the device landscape, enforce granular access policies, and troubleshoot issues faster.

References:

https://www.arubanetworks.com/products/network-managementoperations/central/netconductor/

https://www.arubanetworks.com/assets/wp/WP_NetConductor.pdf

QUESTION 4

Using Aruba best practices what should be enabled for visitor networks where encryption is needed but authentication is not required?

A. Wi-Fi Protected Access 3 Enterprise

B. Opportunistic Wireless Encryption

C. Wired Equivalent Privacy

D. Open Network Access

Correct Answer: B

Explanation: Opportunistic Wireless Encryption (OWE) is a feature that provides encryption for open wireless networks without requiring authentication. OWE uses an enhanced version of the 4-way handshake to establish a pairwise key between the client and the AP, which is then used to encrypt the wireless traffic using WPA2 or WPA3 protocols. OWE can be used for visitor networks where encryption is needed but authentication is not required. References: https://www.arubanetworks.com/assets/tg/TG_OWE.pdf

QUESTION 5

A customer just upgraded aggregation layer switches and noticed traffic dropping for 120 seconds after the aggregation layer came online again. What is the best way to avoid having this traffic dropped given the topology below?

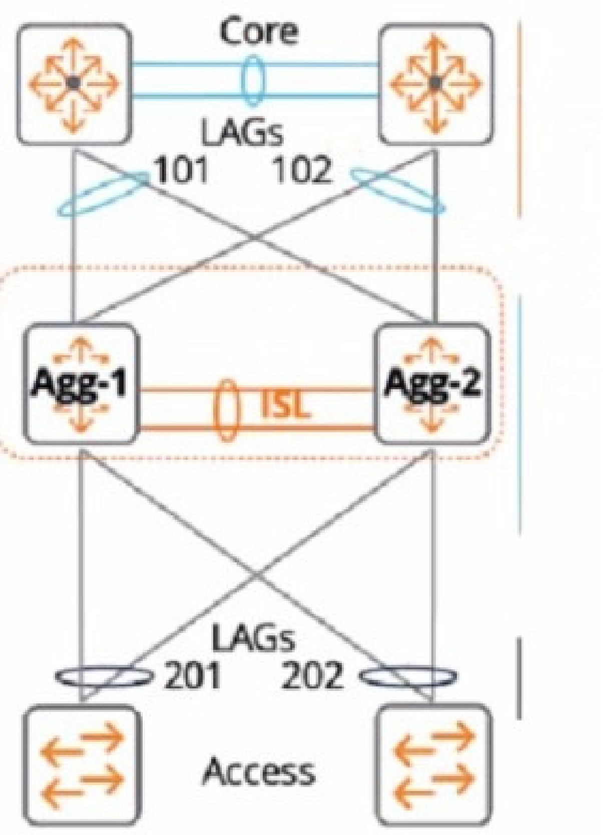

A. Configure the linkup delay timer to 240 seconds to double the amount of lime for the initial phase to sync

B. Configure the linkup delay timer to exclude LAGS 101 and 102, which will allow time for routing adjacencies to form and to learn upstream routes

C. Configure the linkup delay timer to include LAGs 101 and 102, which will allow time for routing adjacencies lo form and to learn upstream routes

D. Configure the linkup delay timer to 120 seconds, which will allow the right amount of time for the initial phase to sync

Correct Answer: C

Explanation: The reason is that the linkup delay timer is a feature that delays bringing downstream VSX links up, following a VSX device reboot or an ISL flap. The linkup delay timer has two phases: initial synchronization phase and link-up delay phase. The initial synchronization phase is the download phase where the rebooted node learns all the LACP+MAC+ARP+STP database entries from its VSX peer through ISLP. The initial synchronization timer, which is not configurable, is the required time to download the database information from the peer. The link-up delay phase is the duration for installing the downloaded entries to the ASIC, establishing router adjacencies with core nodes and learning upstream routes. The link-up delay timer default value is 180 seconds. Depending on the network size, ARP/routing tables size, you might be required to set the timer to a higher value (maximum 600 seconds). When both VSX devices reboot, the link-up delay timer is not used. Therefore, by configuring the linkup delay timer to include LAGs 101 and 102, which are part of the same VSX device as LAG 201, you can ensure that both devices have enough time to synchronize their databases and form routing adjacencies before bringing down their downstream links.