Free Questions for HPE7-A01 by ebraindumps

Shared by Newton on 24-05-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

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
```

Options:

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.

Question 2

Question Type: **MultipleChoice**

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

Options:

A- It extends the LSDB

B- It increases stability

C- it simplifies the configuration.

D- It reduces processing overhead.

E- It reduces the total number of LSAs

Answer:

B, D

Explanation:

Splitting a larger OSPF area into a number of smaller areas has several advantages for network scalability and performance. Some of these advantages are:

It increases stability by limiting the impact of topology changes within an area. When a link or router fails in an area, only routers within that area need to run the SPF algorithm and update their routing tables. Routers in other areas are not affected by the change and do not need to recalculate their routes.

It reduces processing overhead by reducing the size and frequency of link-state advertisements (LSAs). LSAs are packets that contain information about the network topology and are flooded within an area. By dividing a network into smaller areas, each area has fewer LSAs to generate, store, and process, which saves CPU and memory resources on routers.

It reduces bandwidth consumption by reducing the amount of routing information exchanged between areas. Routers that connect different areas, called area border routers (ABRs), summarize the routing information from one area into a single LSA and advertise it to another area. This reduces the number of LSAs that need to be transmitted across area boundaries and saves network bandwidth.

Question 3

Question Type: **MultipleChoice**

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

Options:

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:

A, D

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

Question 4

Question Type: MultipleChoice

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

Options:

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:

A, D

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.

Question 5

Question Type: **MultipleChoice**

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

Options:

A- Wi-Fi Protected Access 3 Enterprise

B- Opportunistic Wireless Encryption

C- Wired Equivalent Privacy

D- Open Network Access

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. Reference: https://www.arubanetworks.com/assets/tg/TG_OWE.pdf

Question 6

Question Type: MultipleChoice

A network administrator is troubleshooting some issues guest users are having when connecting and authenticating to the network The access switches are AOS-CX switches.

What command should the administrator use to examine information on which role the guest user has been assigned?

Options:

A- show aaa authentication port-access interface all client-status

B- show port-access captiveportal profile

C- show port-access role

D- diag-dump captiveportal client verbose

Answer:

A

Explanation:

The show aaa authentication port-access interface all client-status command displays the status of all clients authenticated by port-based access control on all interfaces. The output includes the MAC address, user role, VLAN ID, and session timeout for each client. This command can be used to examine information on which role the guest user has been assigned by the AOS-CX switch. Reference: https://techhub.hpe.com/eginfolib/Aruba/OS-CX_10.04/5200-6692/GUID-9B8F6E8F-9C7A-4F0D-AE7B-9D8E6C5B6A7F.html

Question 7

Question Type: MultipleChoice

Describe the difference between Class of Service (CoS) and Differentiated Services Code Point (DSCP).

Options:

A- CoS has much finer granularity than DSCP

B- CoS is only contained in VLAN Tag fields DSCP is in the IP Header and preserved throughout the IP packet flow

C- They are similar and can be used interchangeably.

D- CoS is only used to determine CLASS of traffic DSCP is only used to differentiate between different Classes.

Answer:

B

Explanation:

CoS and DSCP are both methods of marking packets for quality of service (QoS) purposes. QoS is a mechanism that allows network devices to prioritize and differentiate traffic based on certain criteria, such as application type, source, destination, etc. CoS stands for Class of Service and is a 3-bit field in the 802.1Q VLAN tag header. CoS can only be used on Ethernet frames that have a VLAN tag, and it can only be preserved within a single VLAN domain. DSCP stands for Differentiated Services Code Point and is a 6-bit field in the IP header. DSCP can be used on any IP packet, regardless of the underlying layer 2 technology, and it can be preserved throughout the IP packet flow, unless it is modified by intermediate devices. Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos/configuration/15-mt/qos-15-mt-book/qos-overview.html https://www.cisco.com/c/en/us/support/docs/lan-switching/8021q/17056-741-4.html https://www.cisco.com/c/en/us/support/docs/quality-of-service-qos/qos-packet-marking/10103-dscpvalues.html

Question 8

Question Type: MultipleChoice

What is an OSPF transit network?

Options:

A- a network that uses tunnels to connect two areas

B- a special network that connects two different areas

C- a network on which a router discovers at least one neighbor

D- a network that connects to a different routing protocol

Answer:

A

Explanation:

An OSPF transit network is a network that has at least two routers that are connected by a multi-access link and can forward traffic for other networks1. A transit network is different from a stub network, which has only one router connected to it and does not forward traffic for other networks2. A transit network is also different from a virtual link, which is a logical connection between two areas that are not physically adjacent2. A transit network is not necessarily connected to a different routing protocol, although it can be if the router performs redistribution2. Therefore, the correct answer is C. A network on which a router discovers at least one neighbor.

Question 9

Question Type: MultipleChoice

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

Options:

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:

A, C

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.

Question 10

Question Type: MultipleChoice

You are troubleshooting an issue with a pair of Aruba CX 8360 switches configured with VSX Each switch has multiple VRFs. You need to find the IP address of a particular client device with a known MAC address You run the "show arp" command on the primary switch in the pair but do not find a matching entry for the client MAC address.

The client device is connected to an Aruba CX 6100 switch by VSX LAG.

Which action can be used to find the IP address successfully?

A)

Run the following command on the CX 6100 switch:

```
show mac-address-table
```

B)

Run the following command on the VSX primary switch:

```
show arp all-vrfs
```

C)

Run the following command on the VSX primary switch:

```
show mac-address-table
```

D)

Run the following command on the CX 6100 switch:

```
show arp all-vrfs
```

Options:

A- Option A

B- Option B

C- Option C

D- Option D

Answer:

B

Explanation:

The show arp command displays the ARP table for a specific VRF or all VRFs on the switch. The ARP table contains the IP address to MAC address mappings for hosts that are directly connected to the switch or reachable through a gateway. If the client device is connected to another switch by VSX LAG, the ARP entry for the client device will not be present on the primary switch unless it has communicated with it recently. Therefore, to find the IP address of the client device, the administrator should run the show arp command on the secondary switch in the VSX pair, specifying the VRF name that contains the client device's subnet. Reference: https://techhub.hpe.com/eginfolib/Aruba/OS-CX_10.04/5200-6692/GUID-9B8F6E8F-9C7A-4F0D-AE7B-9D8E6C5B6A7F.html

Question 11

Question Type: MultipleChoice

A customer is looking Tor a wireless authentication solution for all of their IoT devices that meet the following requirements

- The wireless traffic between the IoT devices and the Access Points must be encrypted

- Unique passphrase per device

- Use fingerprint information to perform role-based access

Which solutions will address the customer's requirements? (Select two.)

Options:

A- MPSK and an internal RADIUS server

B- MPSK Local with MAC Authentication

C- ClearPass Policy Manager

D- MPSK Local with EAP-TLS

E- Local User Derivation Rules

Answer:

C, D

Explanation:

The correct answers are C and D.

MPSK (Multi Pre-Shared Key) 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. MPSK requires MAC authentication against a ClearPass Policy Manager server, which returns the encrypted passphrase for the device in a RADIUS VSA2. ClearPass Policy Manager is a platform that provides role- and device-based network access control for any user across any wired, wireless and VPN infrastructure3. ClearPass Policy Manager can also use device profiling and posture assessment to assign roles based on device fingerprint information4.

MPSK Local is a variant of MPSK that allows the user to configure up to 24 PSKs per SSID locally on the device, without requiring ClearPass Policy Manager5. MPSK Local can be combined with EAP-TLS (Extensible Authentication Protocol-Transport Layer Security), which is a secure authentication method that uses certificates to encrypt the wireless traffic between the IoT devices and the access points6. EAP-TLS can also use device certificates to perform role-based access control6.

Therefore, both ClearPass Policy Manager and MPSK Local with EAP-TLS can meet the customer's requirements for wireless authentication, encryption, unique passphrase, and role-based access for their IoT devices.

MPSK and an internal RADIUS server is not a valid solution, because MPSK does not support internal RADIUS servers and requires ClearPass Policy Manager789. MPSK Local with MAC Authentication is not a valid solution, because MAC Authentication does not encrypt the wireless traffic or use fingerprint information for role-based access2. Local User Derivation Rules are not a valid solution, because they do not provide unique passphrase per device or use fingerprint information for role-based access101112.

Question 12

Question Type: **OrderList**

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

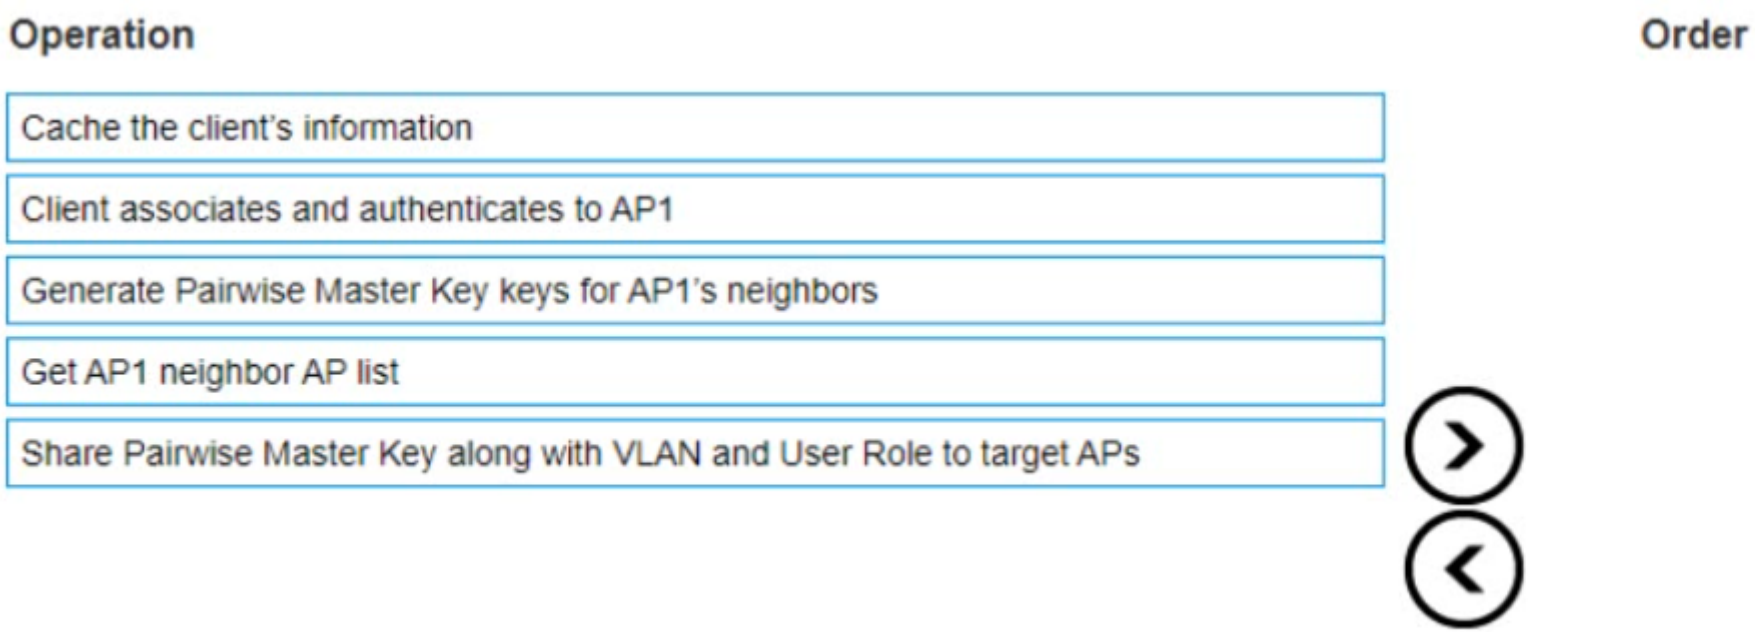

Answer: