

DUMPS ARENA

Aruba Certified Campus Access Switching Expert Written Exam

HP HPE7-A06

Version Demo

Total Demo Questions: 8

Total Premium Questions: 86

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

You want to use OSPF to advertise a only .\16 summary route for the SVIsbelow to a neighbor In the same area (area 0).

Which configuration will achieve this?

- A)
- B)
- C)
- D)
- E.
- A. OptionA
- B. OptionB
- C. OptionC
- D. OptionD
- E. OptionE

F. No configuration can advertise only an internal OSPF summary route to a neighbor in the same area while suppressing the individual SVI routes.


ANSWER: F

Explanation:

No OSPF configuration can advertise only an internal /16 summary while concealing the component SVI prefixes from a neighbor in the same OSPF area. OSPF is a link-state protocol: routers participating in an area synchronize the area's link-state database and calculate routes from the same intra-area topology information. The individual connected SVI networks therefore remain represented by intra-area LSAs and are available to all routers in area 0. An `area range` configuration is applicable on an Area Border Router when it advertises a summary from one area into a different area; it does not summarize or suppress the intra-area topology information exchanged among routers within the source area. Similarly, external-route summarization does not change this intra-area behavior. To advertise a summary instead of the individual SVI routes, the networks must be placed in a non-backbone area and summarized by an ABR when advertising toward another area, such as area 0. This behavior follows the OSPF area architecture and summary-LSA rules defined in [RFC 2328, OSPF Version 2](#). Aruba's OSPF documentation also describes area ranges as ABR route summarization between areas: [AOS-CX area range command](#).

QUESTION NO: 2

Refer to the exhibit.



Acme Corp has VM workload running downstream of ToR-1 and has noticed performance degradation. They suspect ToR-1 uplinks are periodically over utilized. A partner has suggested you migrate your legacy 1U Coie-1 and Cote-2 to the CX 6400 series.

Which aspects of this platform would solve the customer's problem, while focusing on implementing HPE Aruba Networking best practices? (Select two.)

- A. The CX 6400 series supports multiple active forwarding pathways from ToR-1 based on multi-region design.
- B. The proposed new core's VSF capability allows multiple active forwarding pathways from ToR-1 while eliminating the need for STP.
- C. The proposed solutions backplane stacking permits the directly connected ESXI hosts to load balance using active LACP.
- D. MC-LAG permits Core-1 and Core-2 to present to the edge as a common 802.3ad device with a common system ID.

ANSWER: B D

Explanation:

The proposed new core's VSF capability allows multiple active forwarding pathways from ToR-1 while eliminating the need for STP. Virtual Switching Framework combines supported CX switches into one logical switching system. A downstream ToR can therefore form link aggregations to the logical core rather than rely on independent Layer 2 paths, allowing available member links to forward concurrently. This improves utilization of redundant uplinks and provides resiliency if an individual physical link or VSF member fails. Aruba documents VSF architecture and its use of a single logical control and forwarding environment in the [AOS-CX VSF Guide](#).

MC-LAG permits Core-1 and Core-2 to present to the edge as a common 802.3ad device with a common system ID. With a multi-chassis LAG, ToR-1 can use LACP to build one port channel whose physical members terminate on separate core switches. Both member links can carry traffic, so hashing distributes VM traffic flows over the available uplinks rather than leaving a redundant path unavailable. The common LACP system identity is essential because it lets the ToR treat the two peer switches as one LAG partner. This design provides active-active forwarding together with core-switch redundancy, which directly addresses periodic uplink congestion. Aruba describes this operational model in the [AOS-CX VSX and MC-LAG documentation](#).

QUESTION NO: 3

A wired user has successfully authenticated through ClearPass and received a downloadable user role. Later, ClearPass detects a posture change that requires the user to receive a more restrictive role immediately, without waiting for the normal reauthentication interval.

Which mechanism should ClearPass use?

- A. RADIUS Change of Authorization

- B. DHCP lease renewal
- C. RADIUS Access-Accept retransmission
- D. LLDP neighbor update

ANSWER: A

Explanation:

ClearPass should use a RADIUS Change of Authorization message. A CoA allows the policy server to request that the network access device update an existing authenticated session, such as by applying a new user role or triggering session reauthorization.

A new DHCP lease changes addressing information but does not update the authorization state. A RADIUS Access-Accept is part of the initial authentication transaction and is not used to modify an already established session. LLDP can provide endpoint information but does not perform authorization changes.

QUESTION NO: 4 - (SIMULATION)

What is the correct sequence of events that occurs when a user device connects to a network using Dynamic Segmentation?

ANSWER: Check the explanation for the answer

Explanation:

Correct sequence:

ClearPass first authenticates the endpoint and evaluates policy. It returns the appropriate role, which drives VLAN/segment placement and the resulting network-access permissions.

QUESTION NO: 5

Two switches are intended to operate as internal bridges in the same MSTP region. Their spanning-tree configuration names and revision numbers match, but each switch has a different VLAN-to-MST-instance mapping.

Which settings must match for the switches to be recognized as members of the same MSTP region? (Select three.)

- A. MSTP configuration name
- B. MSTP revision number
- C. VLAN-to-MST-instance mapping
- D. Bridge priority
- E. Interface path cost
- F. Root-bridge MAC address

ANSWER: A B C

Explanation:

An MSTP region is identified by its configuration name, revision number, and VLAN-to-instance mapping digest. The digest is calculated from the VLAN mappings, so a mapping difference causes the switches to treat the connection as an MST region boundary even when the configured name and revision number are identical.

Bridge priority is used for root-bridge election and does not determine region membership. Interface cost affects path selection, while the number of configured MST instances is not independently compared; it is the resulting VLAN-to-instance mapping that must match.

QUESTION NO: 6

With the configuration of two CX 8325 switches in the VSX cluster, how would you prepare a link-aggregation for a 7000 gateway for a zero-touch provision to support protocol-based port redundancy?

A)

```
int lag 1
  lacp mode active
  lacp fallback
```

B)

```
int lag 1 multi-chassis
  lacp mode active
  lacp fallback
```

C)

```
int lag 1 multi-chassis
  lacp enable
  lacp fail-over
```

D)

```
int lag 1
  lacp enable
  lacp fail-over
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: B

Explanation:


Option B is correct because it creates the required multi-chassis link aggregation group for a gateway that is dual-attached to the two members of a VSX pair. The `multi-chassis` setting identifies the LAG as a VSX MC-LAG, allowing the Aruba 7000 gateway to see its connections to both physical CX 8325 switches as one logical aggregated interface. This design provides the required device and link resiliency without relying on a single upstream switch.

The configuration also uses `lacp mode active`, which enables standards-based Link Aggregation Control Protocol negotiation. LACP is the protocol-based mechanism that detects compatible member links and maintains the operational aggregation as link state changes. Finally, `lacp fallback` is important for the zero-touch provisioning scenario: it permits connectivity through a member link when the gateway is not yet sending LACP control packets during its initial provisioning state. The gateway can therefore obtain the network services needed for provisioning while retaining LACP-based redundancy once its intended configuration becomes active.

Aruba's VSX documentation describes MC-LAG as the mechanism for a LAG spanning both VSX switches, while the AOS-CX LAG documentation covers LACP operation and fallback behavior. See [Aruba AOS-CX VSX MC-LAG documentation](#) and [Aruba AOS-CX LACP fallback documentation](#).

QUESTION NO: 7

Exhibit.



VSX cluster is already configured. Your task is to validate a correct configuration for the Edge-1 switch that is connected to a CCTV provider that will install its switching infrastructure. The CCTV switches do not support STP.

What needs to be configured on the Edge-1 switch ports connecting to CCTV-SW1 and CCTV-SW2 to prevent loop problems with the existing setup with automatic recovery features?

- A. configure lag with lacp fallback for CCTV switch ports
- B. configure spanning-tree and TCN-guard timeout for CCTV switch ports
- C. configure spanning-tree with bpdu-guard timeout values for CCTV switch ports
- D. configure spanning-tree with udd for CCTV switch ports

ANSWER: C


Explanation:

The correct configuration is **configure spanning-tree with bpduguard timeout values for CCTV switch ports**. Although the CCTV switches do not participate in STP, they can still transparently flood Ethernet frames, including BPDUs received from Edge-1. If the provider's switching infrastructure creates a Layer 2 path between the two Edge-1-facing links, an Edge-1 BPDU transmitted toward one CCTV switch can be flooded through that external topology and received back on the other CCTV-facing port. BPDU Guard identifies this unexpected inbound BPDU condition on an edge-facing port and protects the campus network by placing the affected interface into a disabled state, stopping the loop from continuing to circulate traffic.

Configuring the BPDU Guard timeout provides the required automatic recovery behavior. After the configured timeout expires, AOS-CX can restore the protected port to service, avoiding a permanent manual shutdown while still allowing the switch to react immediately if the loop reappears. This design is appropriate for unmanaged or non-STP external Layer 2 domains connected at the network edge. Aruba documents BPDU Guard behavior and its protection of edge ports in the [AOS-CX Spanning Tree Protocol Guide](#). Additional STP configuration context is available in the [AOS-CX STP configuration documentation](#).

QUESTION NO: 8

Refer to the exhibit.



Acme Corp has VM workload running downstream of ToR-1 and has noticed performance degradation. They suspect ToR-1 uplinks are periodically over utilized. A partner has suggested you migrate your legacy 1U Coie-1 and Cote-2 to the CX 6400 series.

Which aspects of this platform would solve the customer's problem, while focusing on implementing HPE Aruba Networking best practices? (Select two.)

- A. The CX 6400 series supports multiple active forwarding pathways from ToR-1 based on multi-region design.
- B. The proposed new core's VSF capability allows multiple active forwarding pathways from ToR-1 based while eliminating the need for STP.
- C. The proposed solutions backplane stacking permits the directly connected ESXI hosts to load balance using active LACP.
- D. MC-LAG permits Core-1 and Core-2 to present the edge 602.3ad device as a common system ID"

ANSWER: B D

Explanation:

The proposed new core's VSF capability allows multiple active forwarding pathways from ToR-1 based while eliminating the need for STP. VSF virtualizes supported physical switches into one logical switching system. From the downstream switch perspective, this enables a single LACP aggregation toward the VSF system, rather than independent Layer 2 paths toward

separate core devices. The aggregated member links can forward concurrently, increasing usable uplink bandwidth and avoiding reliance on a spanning-tree-blocked redundant path. This directly addresses periodic congestion where only part of the available physical uplink capacity might otherwise carry traffic.

MC-LAG permits Core-1 and Core-2 to present the edge 602.3ad device as a common system ID". In Aruba AOS-CX architectures, multi-chassis link aggregation provides an active-active LACP design across two core switches. Synchronizing the LACP system identity lets the downstream device treat links terminating on both peers as members of one logical LAG. Hash-based distribution can then use both core-facing uplinks while retaining resiliency if a link or one core peer fails. This is an appropriate design for improving forwarding capacity and availability without depending on STP to select a single active uplink. Aruba describes VSF virtualization and VSX multi-chassis LAG operation in the [AOS-CX VSF Guide](#) and [AOS-CX VSX Guide](#).