Boost up Your Certification Score

HP

HPE7-A06

HPE Campus Access Switching Expert Written Exam

For More Information – Visit link below:

https://www.examsboost.com/

Product Version

✓ Up to Date products, reliable and verified.
✓ Questions and Answers in PDF Format.

Visit us at: https://www.examsboost.com/test/hpe7-a06

Latest Version: 6.0

Question: 1

The customer is experiencing periodic uplink congestion between campus-1's AGG-1 and core. This has boon negatively affecting voice communications. The VOIP phones edge mark their packets with DSCP EF. The uplink from AGG-1 to core is LAG1.
The customer has already configured the following class and policy on AGG-1:

```
class ip voip_class
 10 match udp any any dscp 46
 exit
policy voip_policy
 10 class ip voip_class action local-priority 6
 exit
apply policy voip_policy in
```

Based on this policy, which scrip), when deployed on AGG-1. will improve the reliable forwarding of voice traffic between AGG-1 and its uplink to the core?
A)

```
qos schedule-profile 8qDwrrStrict
 dwrr queue 0 weight 1
 dwrr queue 1 weight 1
 dwrr queue 2 weight 1
 dwrr queue 3 weight 1
 dwrr queue 4 weight 1
 dwrr queue 5 weight 1
 strict queue 6  max-bandwidth 300000
 strict queue 7
 exit
apply qos schedule-profile 8qDwrrStrict
```

B)

```
qos schedule-profile 8qDwrrStrict
 dwrr queue 0 weight 1
 dwrr queue 1 weight 1
 dwrr queue 2 weight 1
 dwrr queue 3 weight 1
 dwrr queue 4 weight 1
 dwrr queue 5 weight 1
interface lag 1
 apply qos schedule-profile 8qDwrrStrict
 exit
policy voip_policy
 10 class ip voip_class action local-priority 7
 exit
```

C)

```
qos schedule-profile 8qDwrrStrict
 dwrr queue 0 weight 1
 dwrr queue 1 weight 1
 dwrr queue 2 weight 1
 dwrr queue 3 weight 1
 dwrr queue 4 weight 1
 dwrr queue 5 weight 1
 strict queue 6 max-bandwidth 400000
 dwrr queue 7 weight 5
interface lag 1
 apply qos schedule-profile 8qDwrrStrict
 exit
policy voip_policy
 10 class ip voip_class action local-priority 7
 exit
```

D)

```
qos schedule-profile 8qDwrrStrict
 dwrr queue 0 weight 1
 dwrr queue 1 weight 1
 dwrr queue 2 weight 1
 dwrr queue 3 weight 1
 dwrr queue 4 weight 1
 dwrr queue 5 weight 1
 strict queue 6 max-bandwidth 500000
 dwrr queue 7 weight 1
apply qos schedule-profile 8qDwrrStrict
```

A. Option A
B. Option B

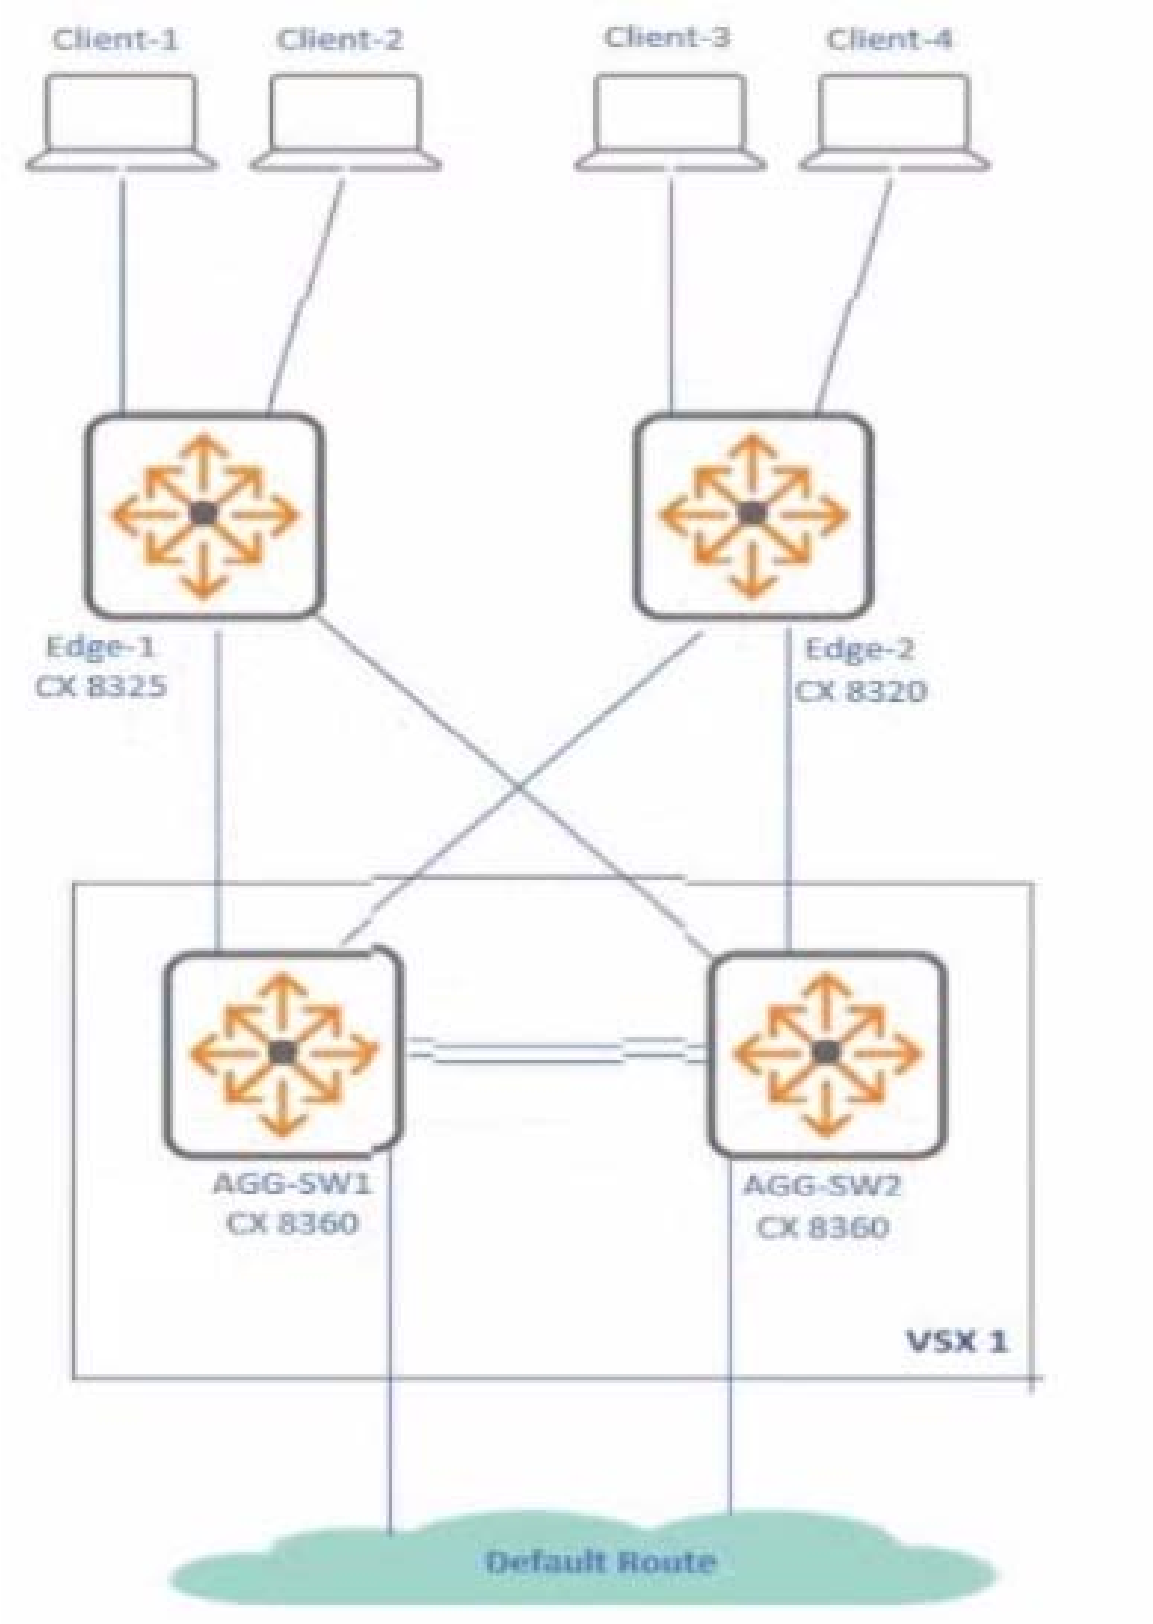
C. Option C
D. Option D

Answer: B

Explanation:
The problem describes uplink congestion affecting VoIP traffic (marked with DSCP EF, value 46) on AGG-1's LAG1 uplink. The existing configuration classifies this traffic into voip_class and applies voip_policy inbound, setting local-priority 6. To improve reliable forwarding during congestion, VoIP traffic needs strict priority queuing on the egress interface (LAG1).
Analysis of Options:
Option A applies a QoS schedule profile globally but doesn't modify the policy's local-priority or apply the schedule profile specifically to the congested LAG.
Option B modifies voip_policy to set local-priority 7 (mapping DSCP 46 traffic to queue 7) and applies the 8qDwrStrict schedule profile to the egress interface lag 1. In the 8qDwrStrict profile, queue 7 is configured for strict priority, ensuring voice traffic gets precedence over lower-priority traffic during congestion. This aligns with best practices for QoS for VoIP.
Option C also sets local-priority 7 and applies the schedule profile to lag 1, but the profile itself configures queue 7 with DWRR (Deficit Weighted Round Robin) instead of strict priority, which is less suitable for delay-sensitive voice traffic.
Option D applies a schedule profile globally and uses DWRR for queue 7.
Conclusion: Option B is the correct solution because it maps the DSCP EF traffic to the highest local priority (7) and applies a QoS schedule profile to the specific congested uplink (lag 1) that treats queue 7 with strict priority. This ensures voice traffic is prioritized reliably.

Question: 2

Exhibit.



Visit us at: https://www.examsboost.com/test/hpe7-a06

A conference venue has a requirement to secure independent network users from each other in their network. The following configurations are created on Edge-1:

A. change the VLAN 152 private-vlan community
B. change the VLAN 151 private-vlan community
C. change the VLAN 152 type. primary-vlan 152
D. change the VLAN 151 primary-vlan 151

Answer: B

Explanation:
The requirement is to secure independent network users from each other in a conference venue using Edge-1. This scenario typically calls for Private VLANs, specifically using the 'isolated' type to prevent communication between hosts within the same secondary VLAN.
Analysis of Options:
Private VLANs consist of a primary VLAN and one or more secondary VLANs (isolated or community). Isolated ports cannot communicate with other isolated ports in the same VLAN; they can only communicate with promiscuous ports (usually the router uplink). Community ports can communicate with each other and promiscuous ports.
Option A: Configures VLAN 152 as private-vlan community.
Option B: Configures VLAN 151 as private-vlan community.
Option C: Defines VLAN 152 as a primary-vlan associated with itself, which isn't standard syntax/logic.
Option D: Defines VLAN 151 as a primary-vlan associated with itself.
The goal is isolation. None of the options directly configure an isolated VLAN. Options A and B configure community VLANs, which allow communication between users within that VLAN, contradicting the requirement. Options C and D attempt to define primary VLANs in a potentially incorrect way.
Caveat: There seems to be an issue with the provided options. Standard configuration to make VLAN 151 isolated would involve defining a primary VLAN (e.g., vlan 152 private-vlan primary) and then defining VLAN 151 as isolated (vlan 151 private-vlan isolated). Since none of the options correctly configure an isolated VLAN, and the requirement is isolation, the question or options are likely flawed. However, if forced to interpret intent, questions sometimes test understanding of the types of private VLANs. Changing a VLAN to community type (Option B for VLAN 151) is a distinct action, even if it doesn't meet the stated isolation goal. Without correct options for 'isolated', selecting the 'best' flawed option is difficult. Assuming the question intends to configure VLAN 151 as some type of private secondary VLAN, Option B modifies VLAN 151's private VLAN characteristic.
Conclusion: Based on the requirement for isolation, none of the provided options are correct. However, if assuming a potential error in the question or options and needing to select the closest modification related to private VLAN types for VLAN 151, Option B is chosen tentatively, despite configuring 'community' instead of the required 'isolated'.

Question: 3

Which tables arc synchronized between a pair of CX 8325 switches in a VSX cluster? (Select two.)

A. BGP Neighbors
B. MAC address
C. Spanning-Tree Protocol (STP)
D. IP Routing
E. Link Layer Discovery Protocol (LLDP)

Answer: B, D

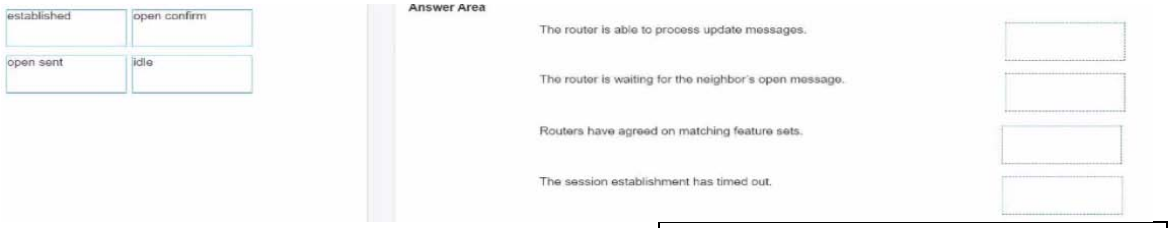
Explanation:
The question asks which tables are synchronized between a pair of CX 8325 switches in a Virtual Switching Extension (VSX) cluster. VSX is a high-availability solution that synchronizes specific tables to ensure consistent operation across both switches.
Analysis of Options:
A . BGP Neighbors: BGP neighbor tables are not synchronized in VSX; each switch maintains its own BGP sessions.
B . MAC address: Correct. VSX synchronizes the MAC address table to ensure consistent Layer 2 forwarding across both switches.
C . Spanning-Tree Protocol (STP): STP states are not synchronized; each switch runs its own STP instance, though they coordinate to avoid loops.
D . IP Routing: Correct. VSX synchronizes the IP routing table to ensure consistent Layer 3 forwarding.
E . Link Layer Discovery Protocol (LLDP): LLDP information is not synchronized; each switch maintains its own neighbor information.
Why B and D are Correct:

In a VSX cluster, the MAC address table and IP routing table are synchronized to ensure seamless Layer 2 and Layer 3 operations. This synchronization allows both switches to share a common view of the network, enabling features like active-active forwarding and hitless failover. The vsx-sync feature in AOS-CX ensures these tables are kept consistent across the VSX pair.
Relevance to Certification Objectives:
Network Resiliency and Virtualization (8%): Involves designing and troubleshooting VSX for resiliency and redundancy.
Switching (19%): Includes implementing and troubleshooting Layer 2 technologies like MAC address tables.
Routing (16%): Covers IP routing table synchronization in VSX environments.
HPE Aruba Networking AOS-CX Configuration Guide: VSX Configuration, detailing table synchronization.
HPE7-A06 Study Guide: Covers VSX architecture and synchronization mechanisms.
HPE Aruba Networking Technical Documentation: VSX Overview, explaining MAC and routing table synchronization.
VSX (Virtual Switching Extension) synchronizes state information between the two switches in a cluster to enable active-active forwarding and provide a single logical view to downstream devices.
Analysis of Options:

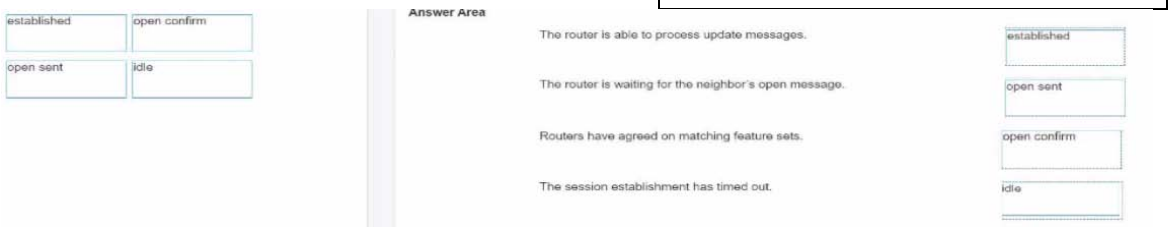
A . BGP Neighbors: BGP sessions are typically established independently by each VSX member. While configurations can be synced, the dynamic state/neighbor table itself is not a core VSX synchronization item.
B . MAC address: The MAC address table is synchronized between VSX members. This is crucial for Layer 2 forwarding consistency and allowing either switch to forward traffic destined for a known MAC address learned via the VSX pair.
C . Spanning-Tree Protocol (STP): STP runs independently on each physical switch. VSX uses technologies like MC-LAG to provide loop-free active-active paths downstream, reducing reliance on STP blocking, but the STP state itself isn't synchronized via the ISL.
D . IP Routing: While the full IP routing table (RIB) is built independently on each switch via routing protocols, VSX Active Gateway synchronizes necessary Layer 3 information (like virtual gateway IP and MAC, and potentially ARP entries) to ensure consistent first-hop routing and failover. Some sources might broadly categorize ARP synchronization under L3/IP routing context in VSX. Given that the ARP table (essential for L3 forwarding consistency) is synchronized, and it's not listed separately, "IP Routing" might encompass this synchronization aspect.
E . Link Layer Discovery Protocol (LLDP): LLDP information relates to physically connected neighbors of each switch and is not synchronized across the VSX ISL.
Conclusion: The MAC address table (B) is definitively synchronized. The ARP table is also synchronized, which is fundamental for Layer 3 forwarding consistency provided by Active Gateway. As ARP is not explicitly listed, and "IP Routing" (D) is, D is the most likely second answer intended to cover the necessary L3 state synchronization (primarily ARP and Active Gateway state) performed by VSX.

Question: 4

DRAGDROP
Match the BGP connection slates to the conditions that could have caused that state.



Answer:



Question: 5

Exhibit.

agg-sw1:

```
show bgp ipv4 unicast neighbors 10.250.17.25 advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.250.255.11

    Network                                 Nexthop                                 Metric     LocPrf
*>e 10.216.12.0/24                          10.250.17.26                            0          0
*>e 10.250.10.0/24                          10.250.17.26                            0          0
*>e 10.250.128.0/22                         10.250.17.26                            0          0
*>e 10.250.255.11/32                        10.250.17.26                            0          0
*>e 192.168.65.0/24                         10.250.17.26                            0          0
*>e 192.168.68.0/24                         10.250.17.26                            0          0
*>e 192.168.69.0/24                         10.250.17.26                            0          0
*>e 192.168.72.0/24                         10.250.17.26                            0          0
Total number of entries 8
```

agg-sw2:

```
show bgp ipv4 unicast neighbors 10.250.17.29 advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.250.255.12

    Network                                 Nexthop                                 Metric     LocPrf
*>e 10.216.12.0/24                          10.250.17.30                            0          0
*>e 10.250.10.0/24                          10.250.17.30                            0          0
*>e 10.250.128.0/22                         10.250.17.30                            0          0
*>e 10.250.255.11/32                        10.250.17.30                            0          0
*>e 192.168.65.0/24                         10.250.17.30                            0          0
*>e 192.168.68.0/24                         10.250.17.30                            0          0
*>e 192.168.69.0/24                         10.250.17.30                            0          0
```

AGG-SW1 and AGG-SW2 are configured with iBGP and eBGP to AS65000. Both agg-sw1 and agg-sw2 use route-map BGP-EXPORT and ip-prefix list local-export in the bgp configuration.
What must be done on agg-swl for the adjacent router to prefer to route all exported routes by agg-sw2?

A. Add set local-preference 200 to the route-map BGP-EXPORT
B. Add set as-path prepend 65345 65345 65345 65345 to the route-map BGP-EXPORT Match with local-export Ip prefix-list.
C. Add set metric 200 to the route-map BGP-EXPORT.
D. Add set as-path 65345 65345 65345 65345 to the route-map BGP-EXPORT Match with local-export ip prefix-list.

Answer: B

Explanation:
The goal is to make the adjacent router prefer routes exported by AGG-SW2 over AGG-SW1 for iBGP and eBGP routes to AS65000. Both switches use a route-map BGP-EXPORT with an ip-prefix list local-export. BGP path selection uses attributes like local preference, AS path length, and metric to determine the preferred route.
Analysis of Options:
Option A: Setting local-preference 200 affects iBGP route selection within the same AS but has no impact on eBGP peers (external AS65000), as local preference is not advertised externally.
Option B: Prepending the AS path with 65345 65345 65345 65345 increases the AS path length for routes exported by AGG-SW1, making them less preferred by the adjacent router (both iBGP and eBGP peers) compared to AGG-SW2's routes, which have a shorter AS path.

Option C: Setting metric 200 affects the MED (Multi-Exit Discriminator), which is used for eBGP route selection within the same AS but is less influential than AS path length and not applicable for iBGP.
Option D: Incorrect syntax (set as-path without prepend) and does not achieve the desired effect.
Why Option B is Correct:

BGP route selection prioritizes the shortest AS path for both iBGP and eBGP. By prepending AS 65345 multiple times to AGG-SW1's exported routes, AGG-SW1's routes appear less attractive due to a longer AS path, causing the adjacent router to prefer AGG-SW2's routes. This is a standard BGP traffic engineering technique.
Relevance to Certification Objectives:
Routing (16%): Involves designing and troubleshooting BGP routing topologies, including manipulating path attributes like AS path.
Troubleshooting (10%): Includes remediating BGP routing issues by adjusting route-maps.
HPE Aruba Networking AOS-CX Configuration Guide: BGP Configuration, covering route-maps and AS path prepending.
HPE7-A06 Study Guide: Details BGP path selection and traffic engineering.
HPE Aruba Networking Technical Documentation: BGP Route Manipulation, explaining AS path prepending for route preference.

Thank You for Trying Our Product

For More Information – Visit link below:

https://www.examsboost.com/

15 USD Discount Coupon Code:

G74JA8UF

FEATURES

- ✓ 90 Days Free Updates
- ✓ Money Back Pass Guarantee
- ✓ Instant Download or Email Attachment
- ✓ 24/7 Live Chat Support
- ✓ PDF file could be used at any Platform
- ✓ 50,000 Happy Customer

