DUMPSBOSS.

HCIA-Datacom V2.0

Huawei H12-811_V2.0

Version Demo

Total Demo Questions: 7

Total Premium Questions: 70

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Data Communication Basics	5
Topic 2, Routing and Switching Basics	35
Topic 3, WLAN Basics	10
Topic 4, Security Basics	6
Topic 5, Network Services and Applications Basics	6
Topic 6, IPv6 Basics	5
Topic 7, SDN and Automation Basics	3
Total	70

QUESTION NO: 1

On the network shown in the figure, GE1/0/1 and GE1/0/2 of SW1 are access interfaces, and their PVIDs are VLAN 2 and VLAN 3 respectively. GE1/0/1 of SW2 is also an access interface, and its PVID is VLAN 4. Which of the following configurations on SW1 and SW2 can ensure that data packets sent from PC1 and PC2 can reach PC3?

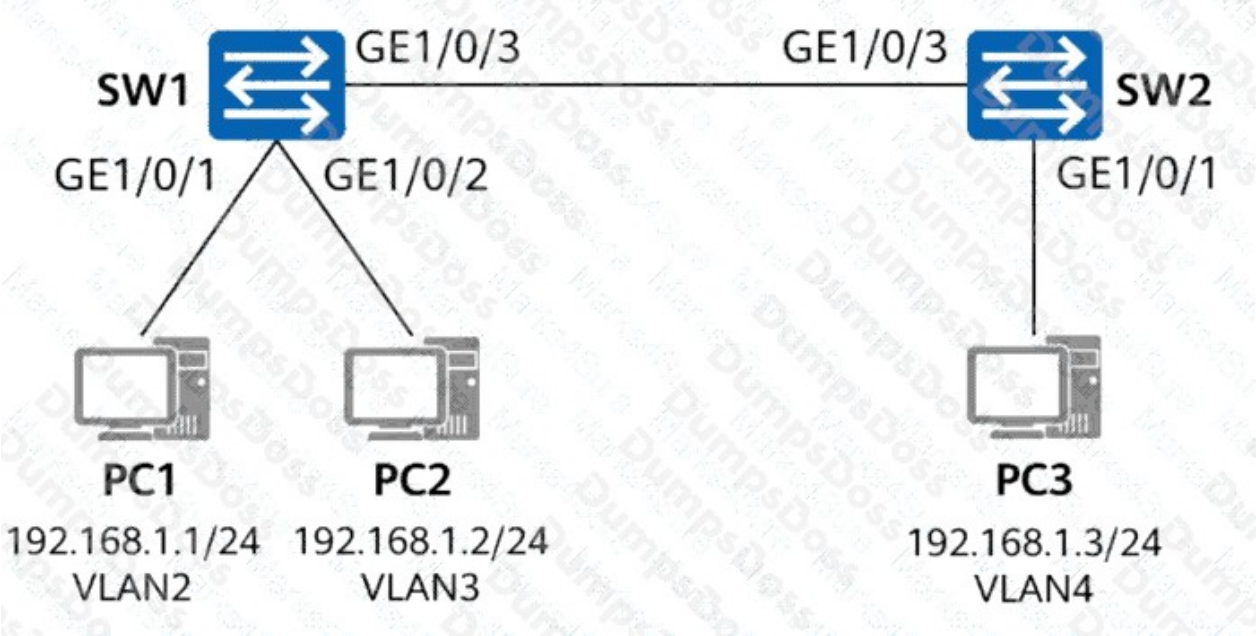

A. Configure GE1/0/3 of SW2 as a hybrid interface, add it to VLAN 4 in untagged mode, and retain the default PVID.

B. Configure GE1/0/3 of SW1 as a hybrid interface, add it to VLAN 2 and VLAN 3 in untagged mode, and set its PVID to VLAN 4.

C. Configure GE1/0/3 of SW2 as a trunk interface, configure it to allow packets from VLAN 4 to pass through, and set its PVID to VLAN 4.

D. Configure GE1/0/3 of SW1 as a trunk interface, configure it to allow packets from VLAN 2 and VLAN 3 to pass through, and set its PVID to VLAN 4.

ANSWER: B C

Explanation:

The required result is achieved by combining “Configure GE1/0/3 of SW1 as a hybrid interface, add it to VLAN 2 and VLAN 3 in untagged mode, and set its PVID to VLAN 4.” with “Configure GE1/0/3 of SW2 as a trunk interface, configure it to allow packets from VLAN 4 to pass through, and set its PVID to VLAN 4.” PC1 and PC2 send traffic into VLAN 2 and VLAN 3 respectively through their access ports. On SW1, the hybrid inter-switch interface is permitted to transmit frames for both of these VLANs without an 802.1Q VLAN tag. Thus, although the traffic originated in two distinct VLANs, it is sent across the physical link in untagged form. On SW2, the trunk interface’s PVID of VLAN 4 classifies every received untagged frame as VLAN 4. The frame can then be switched within VLAN 4 and delivered through SW2’s access interface toward PC3, whose port also belongs to VLAN 4. This configuration intentionally performs VLAN mapping by removing the source VLAN tags at SW1 and assigning received untagged traffic to VLAN 4 at SW2. Huawei’s VLAN interface behavior, including PVID processing and tagged versus untagged VLAN membership, is documented in the Huawei VRP configuration documentation.

QUESTION NO: 2

If the neighbor relationship between R1 and R2 has reached the Full state in OSPF, which of the following statements are true? (Select all that apply)

A. R1 and R2 have different router IDs.

B. R1 and R2 belong to the same area.

C. R1 and R2 use the same process ID.

D. R1 and R2 have the same LSDB in the same area.

ANSWER: A B D

Explanation:

The statements “R1 and R2 have different router IDs,” “R1 and R2 belong to the same area,” and “R1 and R2 have the same LSDB in the same area” are true. OSPF uses the router ID as a unique identifier for each router within an OSPF routing domain. A duplicate router ID prevents a valid adjacency from being maintained, so routers that have successfully reached Full must have unique IDs. OSPF Hello packets also carry the area ID, and neighbors can establish an adjacency only through interfaces assigned to the same area. After the neighbor state progresses through database exchange and loading, Full indicates that the routers have completed link-state database synchronization for that area. This synchronized database enables each router to run the SPF calculation against a consistent topology view and derive routes reliably. The OSPF process ID is not an adjacency parameter exchanged between devices; it is a local configuration identifier, so it has no bearing on whether the completed Full state is achieved. These behaviors are defined in the OSPF neighbor-state and database synchronization procedures in RFC 2328; Huawei’s OSPF configuration documentation likewise describes OSPF area-based operation and neighbor establishment at Huawei Enterprise Support.

QUESTION NO: 3

The administrator configures an Eth-Trunk in LACP mode between two switches, and sets the maximum number of active links in the Eth-Trunk to 3 and the number of remaining standby links to 1. If one of the active links fails, the two switches automatically adjust the number of active links to 2 through negotiation, and the standby link remains in the standby state.

A. TRUE

B. FALSE

ANSWER: B

Explanation:

FALSE is correct because LACP is designed to maintain the configured number of active Eth-Trunk member links when an eligible standby member is available. With a maximum of three active links and one standby link, the bundle initially has three forwarding members and one backup member. If an active member link fails, LACP detects the loss of the member and renegotiates the aggregation state with the peer. The available standby member can then be selected as an active member, restoring the Eth-Trunk to three active forwarding links rather than leaving it at two.

This active/standby selection is a key benefit of LACP-based link aggregation: it provides redundancy while limiting the number of links that actively forward traffic. The standby member is not merely an unused link; it is a candidate selected to replace an unavailable active member, provided that its physical state and LACP negotiation are normal. Consequently, the stated outcome—two active links while the healthy standby link remains standby—does not describe the intended LACP recovery behavior. Huawei’s Eth-Trunk documentation describes LACP’s active and backup link selection and its adjustment following member-link status changes; see Huawei product documentation and Huawei Enterprise Knowledge Base.

QUESTION NO: 4

On the network shown in the figure, the administrator creates Eth-Trunk 1 in manual mode on SW1 and SW2, and adds GE1/0/1 and GE1/0/2 on both switches to the Eth-Trunk. After the administrator runs the display interface brief command on SW1, the administrator finds that GE1/0/1 is Up and GE1/0/2 is Down. Which of the following statements is true about this scenario?

A. Eth-Trunk 1 is Up because it is Up as long as one member interface is Up.

B. Eth-Trunk 1 is Down because the states of member interfaces must be the same in manual mode.

C. Eth-Trunk 1 is Down because it has only two member interfaces and the minimum number of active links is not reached after one member interface is Down.

D. Eth-Trunk 1 is Down because all member interfaces of the Eth-Trunk must be Up.

ANSWER: A

Explanation:

“Eth-Trunk 1 is Up because it is Up as long as one member interface is Up.” is correct. A manually configured Eth-Trunk aggregates multiple physical Ethernet interfaces into one logical interface, providing both increased available bandwidth and link-level redundancy. The operational state of the logical Eth-Trunk does not require every configured member interface to be physically operational. When GE1/0/1 remains Up and is an active member of the bundle, it supplies a usable forwarding path for Eth-Trunk 1. Consequently, the Eth-Trunk interface can maintain its Up state and traffic assigned to the trunk can continue to be forwarded through the remaining available member link. The loss of GE1/0/2 reduces the aggregate capacity available to the bundle, but it does not by itself remove all connectivity from the logical link. This behavior is fundamental to link aggregation: member-link failures can be tolerated without necessarily interrupting traffic carried by the Eth-Trunk. Huawei’s Ethernet switching documentation describes Eth-Trunk as a mechanism that bundles physical links into a logical link and improves reliability through backup among member links. See Huawei’s Ethernet Switching documentation and the Huawei configuration documentation portal for Eth-Trunk configuration and operational guidance.

QUESTION NO: 5

Which of the following hosts are reachable to the host at 192.168.1.200/27 at Layer 2? (Select all that apply)

A. 192.168.1.221/27

B. 192.168.1.192/27

C. 192.168.1.193/27

D. 192.168.1.222/27

ANSWER: A C D

Explanation:

The prefix length `/27` corresponds to the subnet mask `255.255.255.224`. This leaves five host bits in the final octet, so each subnet contains 32 addresses. Subnet boundaries therefore occur at final-octet values of 0, 32, 64, 96, 128, 160, 192, and 224. The address `192.168.1.200/27` falls within the subnet beginning at `192.168.1.192`.

Within this subnet, `192.168.1.192` identifies the network and `192.168.1.223` is the directed broadcast address. The usable unicast host range is consequently `192.168.1.193` through `192.168.1.222`, inclusive. Thus, `192.168.1.221/27`, `192.168.1.193/27`, and `192.168.1.222/27` are valid host addresses in the same IPv4 subnet as `192.168.1.200/27`. Assuming the hosts share the same VLAN/broadcast domain and no Layer 2 filtering prevents communication, the source host can resolve their MAC addresses using ARP and send Ethernet frames directly, without forwarding traffic to a router. Cisco’s IPv4 subnetting overview describes how masks determine network and host portions of an address: Cisco IPv4 subnetting reference. ARP’s role in mapping IPv4 addresses to link-layer addresses is specified in RFC 826.

QUESTION NO: 6

On an STP network, an administrator wants SW2 to be elected as the root bridge. Which of the following methods can be used to achieve this goal?

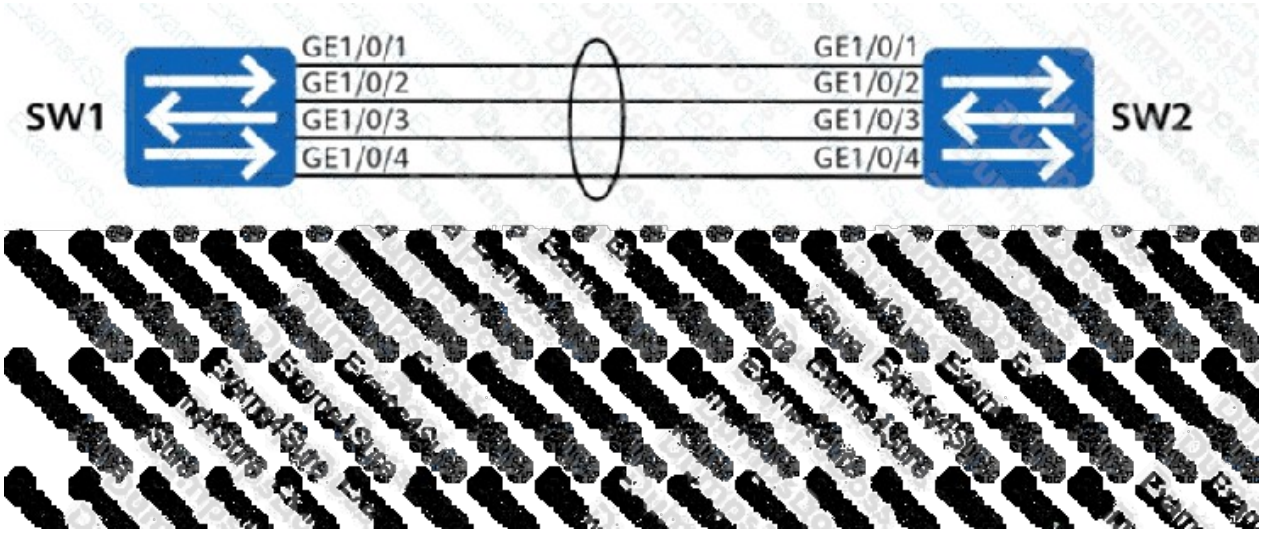

A. Set a lower port priority on SW2.

B. Set a higher system priority on SW2.

C. Set a lower system priority on SW2.

D. Set a lower root path cost on SW2.

ANSWER: C

Explanation:

Set a lower system priority on SW2. In STP, root bridge election is determined by comparing bridge IDs (BIDs) carried in BPDUs. The BID is composed of the bridge priority, the VLAN/system-ID extension where applicable, and the switch MAC address. STP elects the device with the numerically lowest BID as the root bridge. Therefore, configuring SW2 with a lower bridge system priority than the other switches causes its BID to be preferred, provided that no other device has an even lower resulting BID. If bridge priorities are equal, the MAC address is used as the tie-breaker, but relying on a MAC-address tie-break is not an intentional root-placement design. Huawei recommends explicitly planning and configuring root bridge priority so that the root is placed at an appropriate, stable location in the Layer 2 topology. On Huawei VRP devices, STP bridge priority can be configured globally or, in relevant STP modes, for a VLAN instance; the configured priority must use supported increments. This ensures SW2 advertises superior BPDUs and is elected root for the intended spanning-tree instance. See Huawei's STP basic function configuration documentation and the IEEE 802.1D spanning tree standard.

QUESTION NO: 7

An enterprise deploys 802.1X authentication for wireless users. A wireless STA sends authentication information through a Fit AP to a RADIUS server. Which role does the Fit AP perform in the 802.1X authentication process?

A. Authenticator

B. Supplicant

C. Authentication server

D. RADIUS client user

ANSWER: A

Explanation:

The Fit AP performs the authenticator role. In 802.1X, the supplicant is the endpoint requesting network access, such as the wireless STA. The authenticator controls the access port and relays authentication information between the supplicant and the authentication server. The RADIUS server commonly acts as the authentication server.

The Fit AP does not normally make the final user-credential decision itself when RADIUS authentication is deployed. It enforces the access result returned through the authentication process.