

DUMPSBOSS.

HCSP-Presales-Storage

Huawei H19-338

Version Demo

Total Demo Questions: 7

Total Premium Questions: 78

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Huawei Storage Product Introduction	25
Topic 2, Huawei Storage Solution Introduction	17
Topic 3, Storage Project Planning and Design	5
Topic 4, Mix Questions	31
Total	78

QUESTION NO: 1

Which of the following statements about basic concepts in data center network is incorrect?

- A. A leaf node is a fabric network function access node and provides various network devices to access a VXLAN network
- B. Spine nodes are core nodes on the fabric network. They provide high-speed IP forwarding and conned to leaf nodes using high-speed interfaces in any scenario, spine nodes do not need to support VXLAN
- C. The overlay technology is a tunnel encapsulation technology, such as VXLAN The basic principle is that packets are encapsulated in tunnel encapsulation mode, transparently transmitted on the underlying network, and then decapsulated to obtain original packets after reaching the destination. This is equivalent to a network overlay (overlay) on the underlying network.
- D. The underlay network is a bearer network consisting of various physical devices, such as TOR switches, aggregation switch core switches, load balancing devices and firewalls

ANSWER: B

Explanation:

The statement “Spine nodes are core nodes on the fabric network. They provide high-speed IP forwarding and conned to leaf nodes using high-speed interfaces in any scenario, spine nodes do not need to support VXLAN” is incorrect because the absolute claim that spine nodes never need VXLAN support is not valid for every data-center fabric design. In a typical leaf-spine VXLAN fabric, a spine commonly acts only as an IP underlay transit device, forwarding encapsulated VXLAN packets between leaf nodes at high speed. However, network roles are deployment-dependent rather than fixed solely by the device’s physical position in the topology. A spine can also be assigned an overlay-related function, such as a VXLAN Tunnel Endpoint role, a centralized gateway role, or an interconnection role requiring it to terminate, originate, or process VXLAN-encapsulated traffic. In those deployments, VXLAN capability is required on that spine device. VXLAN itself encapsulates Layer 2 frames in UDP/IP packets so that overlay segments can be transported across an IP underlay; this behavior and the role of VXLAN tunnel endpoints are defined in [RFC 7348](#). Huawei’s data-center networking portfolio likewise supports VXLAN-based fabric architectures whose functional roles can be deployed according to service and design requirements: [Huawei Data Center Networking](#). Therefore, “in any scenario” makes the statement inaccurate.

QUESTION NO: 2

CloudEngine S12700E is a flagship product for Huawei’s CloudCampus Solution- it is designed especially for customers to transform from traditional networks to intelligent, simple experience-centric networks Which of the following statements about the features of CloudEngine S12700E is false?

- A. A single CloudEngine S12700E can manage up to 10K APs
- B. A single CloudEngine S12700E can manage up to 50K terminals
- C. CloudEngine S12700E supports PoE
- D. CloudEngine S12700E integrates hardware-based OAM/BFD

ANSWER: C

Explanation:

“CloudEngine S12700E supports PoE” is the false statement. The CloudEngine S12700E is positioned as a high-end, modular core and aggregation switch for large campus networks. Its hardware design focuses on high-capacity switching, service convergence, resilient campus fabrics, and intelligent management capabilities rather than providing edge-device power on access ports. Consequently, the platform does not provide Power over Ethernet (PoE) functionality. PoE is normally implemented by access-layer switches, where it can deliver electrical power over Ethernet cabling to endpoint devices such as wireless access points, IP telephones, cameras, and IoT devices. In a CloudCampus architecture, S12700E equipment can act as a core device and provide central WLAN and user-management capabilities, while PoE-capable access switches supply power to connected endpoints. Huawei’s CloudEngine S12700E product information describes the platform’s role in building high-performance, intelligent campus networks and its extensive service and management

capabilities; PoE is not listed as a supported interface power feature. See Huawei's [CloudEngine S12700E product page](#) and the [Huawei CloudCampus solution overview](#) for the platform positioning and campus-network architecture.

QUESTION NO: 3

The desktop USG6530E can compete with all desktop models of Fortinet and Cisco What is the large-packet throughput of the USG6530E?

- A. 1 Gbit/s
- B. 2 Gbit/s
- C. 3 Gbit/s
- D. 4 Gbit/s

ANSWER: C

Explanation:

The correct large-packet throughput for the USG6530E is **3 Gbit/s**. In this context, “large-packet throughput” refers to the firewall forwarding performance measured with 1500-byte Ethernet frames, rather than performance with small packets or with security services enabled. This distinction matters because packet size has a direct effect on packets-per-second processing requirements: a given bit rate made up of 1500-byte packets requires far fewer packet-processing operations than the same bit rate made up of short frames. Huawei’s performance specifications for the USG6530E identify 3 Gbit/s as its large-packet firewall throughput, making this the appropriate value for a comparison based on basic high-volume traffic forwarding. The figure should be interpreted as a platform performance specification under the stated packet-size test condition; actual throughput in a deployment can vary with enabled security inspection features, traffic mix, policy configuration, and interface conditions. Huawei positions the USG6000E series as next-generation firewalls with performance metrics specified for defined service and test scenarios. See Huawei’s [next-generation firewall portfolio](#) and the [USG6000E product documentation](#) for product and specification context.

QUESTION NO: 4

The CloudEngine 16800 series switches support 4, 8, 12, and 16 LPU slots.

- A. True
- B. False

ANSWER: B

Explanation:

False is correct because the CloudEngine 16800 chassis family is offered in models with 4, 8, and 16 line processing unit (LPU) slots, rather than a model supporting 12 LPU slots. The chassis naming reflects these supported capacities: the CE16804 provides four LPU slots, the CE16808 provides eight, and the CE16816 provides sixteen. Therefore, a statement listing 4, 8, 12, and 16 as supported LPU-slot counts is inaccurate: it introduces a 12-slot configuration that is not part of the CloudEngine 16800 series chassis portfolio. In presales sizing and solution design, the chassis selection should consequently be based on the required number of service-line cards and future expansion capacity using the available 4-slot, 8-slot, or 16-slot platforms. Huawei’s CloudEngine 16800 product information identifies this family as a high-density data-center core switch series with these chassis form factors. See Huawei’s [CloudEngine 16800 product page](#) and the related [CloudEngine 16800 hardware documentation](#) for chassis and slot specifications.

QUESTION NO: 5

A customer expects rapid data growth but wants to minimize the initial purchase capacity. The customer also has many similar virtual-machine images and database test copies. Which three capabilities should be proposed to improve capacity utilization? Select three.

- A. Thin provisioning
- B. Data deduplication
- C. Data compression
- D. Increasing the number of RAID member disks
- E. Synchronous remote replication

ANSWER: A B C

Explanation:

Thin provisioning allocates physical capacity as hosts actually consume it instead of reserving all logical capacity at LUN creation. **Deduplication** reduces repeated data blocks, which is particularly valuable for similar virtual-machine images and copies. **Compression** reduces the physical space required by compressible data.

Increasing the number of RAID disks can improve performance or resilience, depending on the design, but does not itself reduce the amount of capacity consumed by stored data. Synchronous replication creates an additional protected copy and therefore increases, rather than reduces, capacity requirements.

QUESTION NO: 6

An administrator must create local recovery points before database upgrades and provide writable copies for test activities without performing a full backup restore. Which three statements about HyperSnap are correct? Select three.

- A. It can create local point-in-time recovery points before a planned change.
- B. It can support rapid restoration after an application-level error.
- C. It can be used to provide copies for test and development workflows.
- D. It independently protects data if the entire production storage system is lost.
- E. It requires a full physical copy of the source LUN for every recovery point.

ANSWER: A B C

Explanation:

HyperSnap creates local point-in-time copies that can be used for rapid recovery after application errors or unsuccessful changes. Snapshots can also be used to provide writable copies for testing and development, subject to the selected snapshot and clone workflow. Because snapshots reside on the production storage system, they are useful for local operational recovery but do not replace an independent remote disaster-recovery or backup copy.

A snapshot is not a substitute for protection against complete array or site loss. It also does not require a full-volume copy operation each time a recovery point is created.

QUESTION NO: 7

A customer plans to consolidate VMware, database, and file-service workloads on Huawei enterprise storage. Which three protocol-to-workload associations are appropriate? Select three.

- A. FC for database block LUNs
- B. iSCSI for VMware block datastores

C. NFS for NAS datastores or file services

D. SMTP for host access to block LUNs

E. DNS for file-system mounting

ANSWER: A B C

Explanation:

FC and **iSCSI** are block-storage protocols and are appropriate for workloads that require block LUNs, such as database hosts and VMware VMFS datastores. **NFS** is a file-storage protocol and can be used to provide NAS datastores or file shares to supported workloads.

SMTP is an email-transfer protocol, not a storage access protocol. DNS provides name-resolution services and does not provide host access to block or file storage.