

HCIA-Storage V4.5

Huawei H13-611 V4.5

Version Demo

Total Demo Questions: 55

Total Premium Questions: 556

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Storage Technology Trends	46
Topic 2, Basic Storage Technologies	175
Topic 3, Huawei Storage Products and Features	54
Topic 4, Storage Resource Management	35
Topic 5, Storage Data Protection	125
Topic 6, Storage Connectivity	92
Topic 7, Storage Performance Optimization	13
Topic 8, Storage O&M Management	12
Topic 9, Storage Troubleshooting	1
Topic 10, Mix Questions	3
Total	556

QUESTION NO: 1

Which of the following are advantages of NAS storage system over file servers? (Multiple Choice)

- A. High data security
- B. Diverse services and powerful functions
- C. Large storage space
- D. High stability and reliability

ANSWER: A B C D

Explanation:

All listed characteristics describe common advantages of a purpose-built NAS storage system compared with a conventional file server. **High data security** is provided through centralized access control, authentication integration, permissions, auditing, encryption capabilities, and storage-level protection features. **Diverse services and powerful functions** reflects NAS support for standardized file-sharing protocols and centralized data services, enabling clients and applications to access shared files efficiently while administrators manage the environment from a unified storage platform.

Large storage space is an important NAS benefit because dedicated storage systems are designed for scalable capacity expansion and consolidated file-data pools, rather than being constrained by the internal disks and host-centric architecture typical of many file servers. **High stability and reliability** comes from enterprise storage features such as redundant hardware components, RAID-based data protection, high-availability designs, and nondisruptive maintenance capabilities. Huawei enterprise storage platforms provide NAS/file services together with these capacity, availability, and data-protection capabilities. See the [Huawei OceanStor Dorado product information](#) and [Huawei OceanStor Pacific product information](#).

QUESTION NO: 2

Which of the following are deployment modes of cloud computing? (Multiple Choice)

- A. Public cloud
- B. Private cloud
- C. Hybrid cloud
- D. E-Government cloud

ANSWER: A B C

Explanation:

Public cloud, private cloud, and hybrid cloud are the standard cloud-computing deployment modes. A public cloud delivers shared, provider-operated resources and services to multiple customers over a network, with logical isolation between tenants. A private cloud provides cloud capabilities for the exclusive use of one organization; it may be deployed on the organization's premises or hosted and operated by another party. A hybrid cloud deliberately combines distinct private and public cloud environments, enabling data, applications, and workloads to be used across those environments according to requirements such as security, compliance, performance, and elasticity.

These deployment models describe who owns, operates, and consumes the cloud infrastructure, as well as how cloud environments are connected. This classification is distinct from cloud service models, such as IaaS, PaaS, and SaaS, which describe the level of service delivered. Huawei Cloud describes public-cloud services as resources provided through Huawei-operated cloud infrastructure, while its hybrid-cloud portfolio supports coordinated use of on-premises and public-cloud resources. See Huawei Cloud's [Hybrid Cloud Solution](#) overview and the NIST [cloud-computing definition](#) for the recognized deployment-model terminology.

QUESTION NO: 3

Which RAID level combines mirroring and striping, and can tolerate the failure of one disk in each mirrored pair?

- A. RAID 0
- B. RAID 1
- C. RAID 5
- D. RAID 10

ANSWER: D

Explanation:

RAID 10 is correct. RAID 10 combines RAID 1 mirroring with RAID 0 striping. Data is first duplicated across mirrored disk pairs and then striped across those pairs. This design provides high read and write performance together with redundancy.

Because each disk has a mirror copy, RAID 10 can tolerate a disk failure as long as both disks in the same mirrored pair do not fail. For example, in a four-disk RAID 10 group, two disks can fail if they belong to different mirror pairs. Huawei storage systems use RAID configurations to balance performance, usable capacity, and fault tolerance. See the [SNIA RAID overview](#) and [Huawei Enterprise Storage Support](#).

QUESTION NO: 4

Lifecycle Management Of Information and Data consists of 5 different stages. Which of the following is the second stage?

- A. Maintenance.
- B. Transmission.
- C. Generation.
- D. Usage.

ANSWER: B

Explanation:

Transmission. is the second stage in the five-stage information and data lifecycle used in Huawei storage training: generation, transmission, usage, maintenance, and destruction. After information is generated or captured by an application, user, device, or business process, it must be transferred to the system where it will be processed, stored, or made available to authorized users. That movement is the transmission stage.

Transmission is an essential lifecycle step because it establishes how newly created data enters the storage and service environment. In practice, this can include host-to-storage I/O, file transfers, database replication, ingestion through application interfaces, and transfer across network links. Storage designs must protect data while it is transmitted by applying appropriate connectivity, access controls, integrity protection, and network security measures. Once the data has arrived at its intended destination, it can enter the usage stage, where business applications and users access and process it.

This lifecycle perspective helps administrators apply suitable storage, protection, performance, and governance controls at the point data moves between systems. Huawei's enterprise-storage materials emphasize secure, reliable data infrastructure and lifecycle-oriented data management: [Huawei Enterprise Storage](#) and [Huawei Storage Products](#).

QUESTION NO: 5

When the disk is read and written, the sector is the minimum addressing unit, and a segment can be written to the first half of a sector

Data, and then write a piece of data to another sector.

- A. TRUE
- B. FALSE

ANSWER: B

Explanation:

FALSE is correct because a disk sector is the smallest physical addressable unit used for disk I/O. Although an application, file system, or database may handle a data segment that is smaller than a sector, the storage device does not independently address or write only a selected half of one sector as a separate physical disk operation. A write that modifies part of a sector is handled by the storage stack or device as an update to the relevant complete sector, potentially using read-modify-write processing where required. Consequently, data cannot be physically laid out by writing one independently addressed portion to the first half of a sector and then treating another portion as a separate sector-level write target. Sector-level addressing establishes the granularity at which the disk exposes physical logical block access, while higher-level software maps files and smaller logical data structures onto those blocks. Huawei storage documentation describes storage systems in terms of block-based data services and logical block access, consistent with sector/block granularity in the underlying media. See Huawei's [storage product overview](#) and Huawei's [storage documentation portal content](#) for storage and block-service concepts.

QUESTION NO: 6

What are components of an ICT infrastructure? (Choose all that apply.)

- A. Software
- B. Hosts and Servers
- C. Backup Devices
- D. Network switches

ANSWER: A B C D

Explanation:

ICT infrastructure is the integrated foundation used to deliver information and communications technology services. It includes physical compute resources, such as hosts and servers, which run workloads and applications. It also includes network switches, which provide the connectivity needed for hosts, storage, users, and other networked systems to communicate. Backup devices are infrastructure components because they support data protection, recovery, and service continuity; storage training commonly treats backup and recovery facilities as part of the wider data-center infrastructure.

Software is also a component of ICT infrastructure. Infrastructure is not limited to hardware: operating systems, virtualization software, management platforms, and other foundational software enable hardware resources to be configured, operated, monitored, and used to provide services. Huawei describes its ICT portfolio as spanning infrastructure domains including computing, storage, networking, and related software capabilities. Therefore, Software, Hosts and Servers, Backup Devices, and Network switches are all valid components of an ICT infrastructure. See Huawei's [ICT Insights](#) and its [Enterprise ICT solutions](#) overview.

QUESTION NO: 7

Which of the following statements about NFS are correct?

(Multiple Choice)

- A. NFS provides file-level access to remote file systems.
- B. NFS clients can mount exported file systems.
- C. NFS is commonly used by UNIX and Linux hosts.

D. NFS presents raw block LUNs to hosts.

ANSWER: A B C

Explanation:

NFS is a network file system protocol that enables a client to access files located on a remote server through an IP network. An NFS client mounts an exported file system and accesses its directories and files in a manner similar to a local file system. NFS is widely used by UNIX and Linux hosts in NAS environments.

NFS provides file-level access, not block-level LUN access. A client does not use NFS to receive a raw disk device for local formatting. NFS can operate over Ethernet and TCP/IP networks, allowing multiple clients to access file-system services provided by a NAS server. See the [NFS Version 4 Protocol specification](#) and [Huawei OceanStor Data Storage](#).

QUESTION NO: 8

After the multi-path software is installed on the host, which of the following descriptions are correct? (Multiple choice)

- A, LUN is automatically mapped to the host
- B. Assist the operating system to access the same LUN through multiple links to avoid logical errors
- C. Improve link reliability and prevent single link failure, thereby avoiding system failure
- D. Establish a SCSI connection between the host and the storage device

Answer: B, C

Explanation:

- A. Assist the operating system to access the same LUN through multiple links to avoid logical errors
- B. Improve link reliability and prevent single link failure, thereby avoiding system failure
- C. Establish a SCSI connection between the host and the storage device

ANSWER: A B

Explanation:

Multipathing software allows a host to use more than one physical I/O path to the same mapped LUN while presenting that storage consistently to the operating system. It discovers and manages the available paths, selects an active path for I/O, and coordinates path use so that the host can safely continue treating the storage as one logical device. This is why “Assist the operating system to access the same LUN through multiple links to avoid logical errors” is correct: the multipathing layer provides the operating-system-side path management needed when several routes reach identical storage.

“Improve link reliability and prevent single link failure, thereby avoiding system failure” is also correct. In a properly designed redundant SAN deployment, separate host bus adapters, switches, controllers, and cables can provide alternate paths. If one path becomes unavailable because of a component or connectivity failure, the multipathing driver detects the condition and fails I/O over to an available healthy path. This removes a single-path dependency and improves service continuity. Huawei describes host multipathing as a host-side capability for managing redundant storage access paths; its enterprise support portal provides product documentation and host-connectivity guidance at [Huawei Enterprise Support](#). Huawei’s storage portfolio and high-availability design context are also available at [Huawei Storage](#).

QUESTION NO: 9

After employee A of a company logs into the storage device of the company system, he finds that he has certain control rights. But not repairable

If the system time is changed, user management cannot be performed. May I ask which of the following roles does Employee A have?

- A. Super administrator
- B. Administrator
- C. ordinary users
- D. Read-only users

ANSWER: B

Explanation:

Administrator is the appropriate role. In Huawei storage management, an administrator has operational control permissions needed for routine storage administration, such as managing storage resources and carrying out day-to-day system management tasks. However, permissions that affect the fundamental security and governance of the storage system are reserved for the super-administrator level. Changing the system time is a system-level security-sensitive action because time affects event records, alarms, audit trails, certificates, and log correlation. User management is likewise a privileged security function because it controls identities, credentials, and authorization assignments. The stated combination of having meaningful control rights while being unable to change the system time or manage users therefore corresponds to the Administrator role. Huawei storage documentation distinguishes user roles by their management scope and reserves the highest-level system and account-management authority for the super administrator. For product documentation and role-based management guidance, consult the [Huawei Enterprise Support documentation portal](#) and the [Huawei OceanStor product page](#).

QUESTION NO: 10

Which statements are true about distributed storage? (Select All that Apply)

- A. Supports large-scale horizontal expansion
- B. Does not support elastic EC for data redundancy protection.
- C. Supports expansion by addition of nodes to increase both computing capability and storage space
- D. Supports storage resource pooling and virtualization.

ANSWER: A C D

Explanation:

Distributed storage is built on a scale-out architecture. "Supports large-scale horizontal expansion" is correct because capacity and performance can grow by adding storage nodes, rather than relying solely on increasingly powerful individual systems. This architecture is intended to accommodate continuously growing unstructured-data workloads and large numbers of concurrent users.

"Supports expansion by addition of nodes to increase both computing capability and storage space" is also correct. In a distributed storage cluster, each added node commonly contributes storage devices as well as processing, network, and service resources. The cluster therefore expands its aggregate capability while presenting resources through the distributed system.

"Supports storage resource pooling and virtualization" is correct because distributed storage combines physical disks and nodes into a shared logical resource pool. Administrators can provision and manage storage from this unified pool instead of managing each device independently. Huawei describes its scale-out OceanStor Pacific distributed storage as supporting elastic scale-out and unified resource management, characteristics that reflect horizontal expansion and pooled storage architecture. See the [Huawei OceanStor Pacific product page](#) and Huawei's [enterprise storage solutions overview](#) for further information.

QUESTION NO: 11

What are the advantages of SSDs over HDDs?

(Multiple Choice)

- A. Larger capacity
- B. Faster read/write speed for random access
- C. Better environmental adaptability
- D. Longer service life

ANSWER: B C

Explanation:

Faster read/write speed for random access is a core SSD advantage because an SSD accesses data electronically in flash memory rather than moving a mechanical actuator and rotating disk platters to the required location. This removes seek-time and rotational-delay overhead, producing substantially lower latency and much higher random I/O operations per second. These characteristics are especially valuable for database workloads, virtual machines, boot volumes, and other applications that issue many small, nonsequential reads and writes.

Better environmental adaptability is also an advantage because SSDs contain no spinning media or moving read/write heads. Their solid-state design provides greater resistance to shock and vibration and allows operation in mobile, edge, or equipment-room environments where physical movement and vibration can be concerns. This mechanical robustness complements the performance benefits of flash storage. Huawei's all-flash storage portfolio emphasizes flash-based performance and reliability characteristics for enterprise workloads, while Seagate's storage overview describes the low-latency, moving-part-free properties of SSD technology. See [Huawei All-Flash Storage](#) and [Seagate: SSD vs. HDD](#).

QUESTION NO: 12

Which of the following terms describes the delay caused by waiting for the required sector on an HDD platter to rotate beneath the read/write head?

- A. Seek time
- B. Rotational latency
- C. Cache hit ratio
- D. Transfer rate

ANSWER: B

Explanation:

Rotational latency. is correct. After the read/write head reaches the correct track, the disk may still need to rotate until the required sector is positioned under the head. The waiting time for this rotation is called rotational latency. It is one of the mechanical delays that affects HDD response time.

Seek time is the time required to move the read/write head to the required track. Transfer rate describes how much data can be moved per unit of time, while cache hit ratio indicates the percentage of requests served from cache. SSDs do not have rotational latency because they contain no rotating platters. See [Huawei Enterprise Storage](#) for storage performance information.

QUESTION NO: 13

Which of the following are the data transmission rates that can be configured on the FC ports on the FC switch in FC SAN given that your cable and SFP supports it? (Select All that Apply)

- A. 8Gb
- B. 11Gb
- C. 4Gb.
- D. 16Gb.

ANSWER: A C D

Explanation:

8Gb, **4Gb.**, and **16Gb.** are valid Fibre Channel port line rates. Fibre Channel uses defined, interoperable speed generations, and 4 Gbit/s, 8 Gbit/s, and 16 Gbit/s are established rates supported by corresponding generations of FC switch ports and optical transceivers. In an FC SAN, a port can be configured to a particular supported rate or, on many platforms, set to negotiate automatically with the connected port. The available choices ultimately depend on the switch port hardware, installed SFP, optical cable type and distance, and the capabilities of the peer port. For example, an FC port using a compatible 16 Gbit/s SFP can commonly operate at 16 Gbit/s and may support backward-compatible 8 Gbit/s or 4 Gbit/s operation when connected to suitable equipment. Selecting a mutually supported rate ensures that the physical FC link initializes correctly and provides the expected bandwidth. Fibre Channel's standardized generations and their respective rates are summarized by the [Storage Networking Industry Association](#) and the [Fibre Channel Industry Association](#).

QUESTION NO: 14

The server with Linux operating system installed accesses the disk array based on the RAID 5 strategy.

The order is:

- A. File system->Application system->Volume->I/O subsystem->RAID controller->Disk
- B. Application System->File System->Volume->I/O Subsystem->RAID Controller->Disk
- C. Application system->File system->I/O subsystem->Volume->RAID controller->Disk

ANSWER: C

Explanation:

The correct access path is **Application system->File system->I/O subsystem->Volume->RAID controller->Disk**. An application first issues a file-level request, which is interpreted by the Linux file system. The request then enters the operating system I/O subsystem, where it is converted and managed as block I/O before being directed to the appropriate logical volume. The volume presents storage capacity to the host while mapping host I/O to the storage resources underneath it. The RAID controller subsequently processes the request according to the RAID 5 layout, distributing data and parity operations across the member disks. Finally, the physical disks perform the required reads or writes. This sequence correctly places the Linux I/O subsystem before the storage volume, reflecting the transition from host-side file and block-I/O processing to array-side logical-volume and RAID processing. RAID 5 protection is implemented by the controller across the disks, but the host normally accesses the array through the volume exposed to it rather than addressing individual disks. The Linux kernel block-layer documentation describes the role of block I/O processing: [Linux Kernel Block Layer documentation](#). Huawei storage documentation is available through the [Huawei Enterprise Storage support portal](#).

QUESTION NO: 15

Statement 1: OceanStor Dorado systems are designed for high performance environments. Statement 2: OceanStor Dorado systems are so-called All Flash Arrays.

- A. Statement 1 and 2 are true.
- B. Statement 1 is true; Statement 2 is false.

C. Statement 1 is false; Statement 2 is true.

D. Statement 1 and 2 are false

ANSWER: A

Explanation:

Statement 1 and 2 are true. Huawei positions the OceanStor Dorado family as an all-flash storage product line intended for enterprise workloads requiring very high performance, low latency, and resilient, continuously available data services. Its architecture uses flash media rather than traditional spinning disks, which is precisely why these systems are categorized as all-flash arrays (AFAs). The Dorado platform is designed to deliver predictable low-latency I/O for demanding applications such as databases, virtualization, and mission-critical transactional workloads. Huawei documentation describes OceanStor Dorado as all-flash storage and highlights its performance-oriented design, including NVMe-based connectivity and intelligent storage functions that help sustain efficient processing under heavy workload conditions. Therefore, the two statements accurately describe both the intended high-performance use case and the all-flash-array classification of OceanStor Dorado systems. See Huawei's [OceanStor Dorado all-flash storage overview](#) and the [OceanStor Dorado product documentation](#) for product positioning and technical information.

QUESTION NO: 16

What applies to unstructured data? (Multiple Choice)

A. After being written, the data is not often modified.

B. Uncertain value.

C. Large capacity, rapid growth.

D. Long-term storage required.

ANSWER: A B C D

Explanation:

Unstructured data commonly includes files such as documents, images, audio, video, surveillance recordings, backups, and application log files. These data sets are typically written and then read, shared, analyzed, or retained far more often than they are modified in place. Their business value can be uncertain when first created: an individual file may have little immediate value, while later analysis, compliance needs, investigations, or AI workloads can make retained data valuable. Unstructured data also expands quickly because modern organizations continuously generate content from users, devices, applications, and digital services. This creates a need for storage platforms that can scale capacity efficiently while maintaining manageable access and protection. Long retention is another key characteristic, driven by archival, governance, audit, legal, backup, and historical-analysis requirements. Consequently, storage for unstructured workloads emphasizes scalable file or object access, capacity efficiency, durability, and lifecycle-oriented data management. Huawei describes its scale-out storage offerings as supporting massive unstructured-data workloads and growth, while its data-protection materials address long-term retention and protection needs. See [Huawei OceanStor Pacific](#) and [Huawei Data Protection](#).

QUESTION NO: 17

Statement 1: RAID 2.0+ offers better protection against data loss than traditional RAID and it performs much better

Statement 2: To rebuild a RAID 2.0+ protected failed drive takes a lot of time as all drives are involved in the rebuild of the spare disk.

A. Statement 1 is true; Statement 2 is false.

B. Statement 1 and 2 are false.

C. Statement 1 is false; Statement 2 is true.

D. Statement 1 and 2 are true.

ANSWER: A

Explanation:

“Statement 1 is true; Statement 2 is false.” is correct. Huawei RAID 2.0+ is a distributed RAID implementation that virtualizes storage resources into small data blocks and distributes those blocks across the available drives in a disk domain. This architecture avoids concentrating a conventional RAID group’s data and reconstruction workload on only a limited set of member disks. It improves overall I/O balancing and uses the aggregate performance of the disk domain, which can deliver better performance than traditional fixed RAID groups.

RAID 2.0+ also enhances reliability and data protection by distributing redundancy information and rebuild work across the drives in the disk domain. When a drive fails, the system reconstructs only the affected small data chunks using the available resources of multiple drives. Consequently, rebuild operations are accelerated rather than inherently taking a long time. A shorter reconstruction window reduces the period during which the system has reduced redundancy and is therefore an important contribution to protection against data loss. Huawei describes RAID 2.0+ as a key reliability and performance technology in its storage systems; see [Huawei OceanStor Dorado](#) and [Huawei Enterprise Support documentation](#).

QUESTION NO: 18

Which of the following is correct? (Multiple Choice)

- A. Only when a snapshot is activated will the system record all changes to the original LUN.
- B. A snapshot needs to be removed from the LUN group to be fully unmapped.
- C. SmartTier supports a High Performance, a Performance, a High Capacity and a Capacity Tier.
- D. To write data to the secondary LUN of a clone pair, the pair must be split.

ANSWER: A B D

Explanation:

Only when a snapshot is activated will the system record all changes to the original LUN. An activated Huawei storage snapshot provides a point-in-time view and causes the storage system to preserve the required original data when subsequent writes modify the source LUN. This change tracking and data preservation is what allows hosts to continue accessing a consistent historical image while production I/O proceeds.

A snapshot needs to be removed from the LUN group to be fully unmapped. In Huawei storage mapping, a LUN group is part of the mapping-view relationship used to present storage resources to hosts. Where a snapshot remains included through that group-level relationship, removing only an individual direct association does not completely remove the effective host mapping; the snapshot must be removed from the LUN group to eliminate that group-based presentation path.

To write data to the secondary LUN of a clone pair, the pair must be split. Before splitting, the secondary LUN is maintained as the clone target under the clone-pair relationship and is protected from independent host writes. Splitting ends that synchronized relationship and makes the secondary LUN independently writable, enabling it to be used for testing, reporting, or other separate workloads. Huawei’s storage portfolio and documentation resources are available at [Huawei Data Storage](#) and [Huawei Enterprise Storage Support](#).

QUESTION NO: 19

The hard disk domain supports online addition of hard disks, and the storage pool supports online expansion.

- A. TRUE
- B. FALSE

ANSWER: A

Explanation:

TRUE is correct. In Huawei storage systems, a disk domain (also called a hard disk domain in some product documentation and training materials) is a logical collection of disks that provides capacity to storage pools. It is designed to support nondisruptive capacity growth: eligible disks can be added while the system remains in service. This lets administrators increase the available physical capacity without taking the storage system offline or interrupting host access.

A storage pool obtains capacity from its associated disk domain and can likewise be expanded online when additional capacity is available. Online expansion is an important operational capability because it allows capacity planning to be implemented as workloads grow, while preserving service continuity for LUNs, file systems, and other resources provisioned from the pool. The controller manages the added capacity and makes it available for subsequent allocation, subject to the platform's disk compatibility, RAID/layout, and configuration requirements. This behavior aligns with Huawei's scale-out and nondisruptive management approach for enterprise storage. See Huawei's [OceanStor Dorado product information](#) and the [Huawei Enterprise Support documentation portal](#) for product-specific storage-pool and disk-domain expansion procedures.

QUESTION NO: 20

When using LAN-based backups it is best to create a dedicated backup network.

- A. TRUE
- B. FALSE

ANSWER: A

Explanation:

TRUE is correct. LAN-based backup sends production data across Ethernet networks from hosts or storage systems to backup infrastructure. Because backup jobs can transfer very large volumes of data, particularly during full backups, replication, or recovery operations, they can consume substantial bandwidth and create congestion. A dedicated backup network isolates this high-volume traffic from business application, user-access, and management traffic. This separation helps backup jobs complete within their intended backup window while maintaining predictable performance for production services. It also enables administrators to size links, configure quality-of-service policies, and monitor network utilization specifically for data-protection workloads. From an operational perspective, separate backup connectivity can simplify troubleshooting and reduce the risk that an unrelated network issue affects both production access and backup data movement. Huawei's storage and data-protection portfolio is designed around reliable protection, recovery, and management of enterprise data; network planning and sufficient bandwidth are important prerequisites for those functions. Further product and solution information is available from [Huawei Enterprise Storage](#) and [Huawei Data Center Solutions](#).

QUESTION NO: 21

What are the challenges facing traditional data centers? (Multiple choice)

- A. Traditional data centers only monitor but not manage
- B. Capacity expansion of traditional data centers in units of nodes
- C. Traditional data center has low cooling efficiency
- D. Traditional data center maintenance costs are high

ANSWER: A B C D

Explanation:

Traditional data centers only monitor but not manage is correct because conventional operations commonly collect alarms and performance information but depend heavily on manual interpretation and intervention. This monitoring-centric approach limits closed-loop management, makes fault handling slower, and raises the operational burden as infrastructure grows. Capacity expansion of traditional data centers in units of nodes is also a challenge: compute, storage, and network resources are often deployed and scaled as fixed hardware units instead of being flexibly pooled and allocated according to workload requirements. This can lead to stranded capacity and inefficient expansion planning.

Traditional data center has low cooling efficiency is correct because conventional facility designs and uneven equipment utilization can create inefficient airflow and cooling operation, increasing power usage. Traditional data center maintenance costs are high is likewise correct: diverse physical devices, fragmented management tools, and frequent manual onsite activities require substantial skilled labor throughout the infrastructure lifecycle. Huawei's data-center guidance emphasizes intelligent, automated management and energy-efficient facilities to address these operational and energy challenges. See [Huawei Data Center Network solutions](#) and [Huawei Intelligent Data Center solutions](#).

QUESTION NO: 22

What is a measure of storage system performance?

- A. IOPS
- B. Bandwidth
- C. capacity
- D. Delay

ANSWER: A B D

Explanation:

IOPS, Bandwidth, and Delay are core measurements used together to assess storage-system performance. IOPS measures the number of input/output operations completed per second. It is especially meaningful for transactional and small-block random workloads, where the storage system must process a high volume of individual requests efficiently. Bandwidth measures the volume of data transferred per unit of time, commonly expressed in MB/s or GB/s. It is important for workloads involving large, sequential transfers, such as backup, media, analytics, and bulk data movement. Delay, also called latency, measures the time required for an I/O request to be serviced and returned to the host. Low delay is essential for responsive applications and is often measured in milliseconds or microseconds. These metrics must be considered in the context of workload characteristics, including I/O size, read/write ratio, access pattern, queue depth, and host connectivity. Huawei storage platforms expose performance monitoring data so administrators can observe IOPS, throughput, and response-time behavior, establish baselines, and identify bottlenecks. See Huawei's [enterprise storage support resources](#) and the [Huawei storage product documentation portal](#) for product-specific monitoring guidance.

QUESTION NO: 23

It is common practice to create an extra copy of a backup tape that holds important data.

- A. TRUE
- B. FALSE

ANSWER: A

Explanation:

TRUE is correct. Creating an additional copy of a tape backup that contains important data is a standard data-protection practice. A second copy reduces the risk that a single event—such as media damage, tape-drive failure, accidental overwrite, loss during transport, or a site-level incident—will prevent recovery. The extra tape copy can be retained separately, ideally in a secure offsite location with environmental controls and documented handling procedures. This approach supports the core backup principle of maintaining multiple copies of important information and separating copies so that one local failure does not affect every recoverable version. For critical business data, backup policies commonly define copy-retention periods, offsite storage requirements, media labeling, and periodic restore testing. Restore tests are especially important because they verify that both the tape data and the recovery process remain usable when required. This practice aligns with established contingency-planning guidance that calls for alternate storage sites and protection of backup information. Huawei storage data-protection solutions likewise emphasize reliable backup, recovery, and protection against data-loss events. See [NIST SP 800-34 Rev. 1, Contingency Planning Guide](#) and [Huawei Data Protection](#).

QUESTION NO: 24

Which of the following are the 3 layers of SCSI Target Model? (Multiple Choice)

- A. Network Layer.
- B. Port Layer.
- C. Middle Layer.
- D. Device Layer.

ANSWER: B C D

Explanation:

The SCSI Target Model is organized into the Port Layer, Middle Layer, and Device Layer. This layered design separates external SCSI connectivity from command processing and the backend storage implementation. The Port Layer presents target ports to initiators through a transport or protocol interface, enabling host systems to establish SCSI access paths. The Middle Layer is responsible for common target-side SCSI processing, including handling commands and coordinating their flow between the presented target port and the logical storage object. The Device Layer represents the underlying logical unit and implements the device-specific operations that fulfill SCSI requests against the storage resources. Together, these layers provide a modular target architecture: front-end access is handled at the port, shared SCSI-target processing occurs in the middle, and data services are delivered by the device implementation. This organization aligns with the SCSI architectural model's separation of logical units, target ports, and command-processing behavior. The SCSI Primary Commands specification provides the standards context for SCSI devices, logical units, target ports, and command behavior: [SCSI Primary Commands \(SPC-5\)](#). Additional SCSI architecture material is available from the [T10 Technical Committee](#).

QUESTION NO: 25

Depending on the amount of data to be backed up, which of the types are actual backup methods?

(Multiple Choice)

- A. Real-time backup
- B. Full backup
- C. Incremental backup
- D. Differential backup

ANSWER: B C D

Explanation:

Full backup, incremental backup, and differential backup are the standard backup methods categorized by the amount and scope of data copied during each backup operation. A full backup creates a complete copy of all selected source data at a specified point in time. It provides a self-contained recovery point and commonly serves as the baseline from which subsequent backups are organized.

Incremental backup copies only the data that has changed since the most recent backup, whether that prior backup was full or incremental. This substantially reduces the amount of data transferred and stored for routine backup jobs. Differential backup copies data changed since the most recent full backup. As a result, each differential backup can grow over time until another full backup is performed, while recovery uses the full backup together with the selected differential backup.

These methods allow backup policies to balance backup-window duration, storage consumption, network use, and recovery requirements. Huawei's backup documentation describes full and incremental backup behavior in its backup policy and backup-management guidance. See [Huawei Cloud Backup User Guide](#) and [Huawei Cloud Backup management documentation](#).

QUESTION NO: 26

Is the following correct understanding of data and information?

- A. Data is a record that reflects the attributes of objective things and a specific form of information
- B. Data refers to information that can be stored and transmitted
- C. Information is data that has been processed to meet the needs of user decisions
- D. Data is the sum of information and data redundancy

ANSWER: A B C

Explanation:

“Data is a record that reflects the attributes of objective things and a specific form of information” is correct because data represents facts, observations, measurements, symbols, or records describing real-world objects and events. In a storage context, these records are the fundamental content that systems collect, retain, protect, and process. “Data refers to information that can be stored and transmitted” is also correct in practical IT terms: for information to be managed by a computer or storage system, it must be represented in a form that can be encoded, stored on media, copied, and communicated across networks.

“Information is data that has been processed to meet the needs of user decisions” correctly identifies the key distinction between raw data and information. When data is organized, interpreted, analyzed, or placed in context, it acquires meaning that can support operational work and decision-making. For example, individual capacity measurements are data; a summarized trend showing predicted capacity exhaustion is decision-useful information. This relationship is central to storage systems because their value is not merely retaining bits, but reliably preserving data so it remains available for processing into useful information. Huawei’s enterprise support resources provide storage documentation and terminology context at [Huawei Enterprise Support](#). See also the NIST definition and related terminology for information at [NIST CSRC Glossary](#).

QUESTION NO: 27

In high end storage systems, what is the main function of the service processor?

- A. To manage the I/O commands and complete disk operations.
- B. To process the network service.
- C. To maintain the interconnection of disk drives within the enclosures.
- D. To optimize the speed of I/O operations.

ANSWER: A

Explanation:

“To manage the I/O commands and complete disk operations.” is correct because the service processor is the core processing component that delivers the storage array’s data services. It receives host read and write requests through the front-end interfaces, interprets and schedules those requests, and coordinates cache, RAID, back-end disk access, and data-placement operations before returning results to the host. In an enterprise storage array, this I/O-processing responsibility is central to providing block-storage services reliably and with high availability. Service processors are normally deployed redundantly so that the system can continue servicing host I/O if a processor or path fails. Their processing also supports storage capabilities that operate in the I/O path, such as cache management, RAID protection, virtualization, and data-reduction functions where applicable. Huawei enterprise storage systems use controller/service-processing resources to handle host access and coordinate internal storage resources, rather than assigning this end-to-end responsibility solely to an enclosure component. See Huawei’s [enterprise storage product information](#) and [OceanStor documentation portal](#) for architecture and component documentation.

QUESTION NO: 28

Which of the following are common RAID types. (Multiple Choice)

- A. RAID 1
- B. RAID 5
- C. RAID 0+1
- D. RAID 50

ANSWER: A B C D

Explanation:

RAID 1, RAID 5, RAID 0+1, and RAID 50 are all recognized and commonly implemented RAID layouts. RAID 1 uses mirroring, maintaining duplicate copies of data across member disks to provide redundancy. RAID 5 uses block-level striping with distributed parity, combining usable capacity and fault tolerance for a disk-group configuration. RAID 0+1 is a nested RAID arrangement that combines striping and mirroring; it is widely recognized in storage terminology and is supported in appropriate storage-controller and disk-group designs. RAID 50 is another nested layout, combining multiple RAID 5 groups with RAID 0 striping across those groups, allowing systems to scale performance and capacity while retaining parity-based protection within each constituent group.

These names follow established RAID nomenclature: standard RAID levels can be combined into nested layouts to address differing availability, performance, and capacity objectives. The exact RAID levels available on a particular Huawei storage product depend on its model, software version, and supported disk-domain or storage-pool configuration, but each listed layout is a valid common RAID type. See the [SNIA online storage dictionary](#) and Huawei's [storage product documentation](#) for RAID terminology and implementation guidance.

QUESTION NO: 29

Which of the following are advantages of using a dual-controller storage system?

(Multiple Choice)

- A. They provide controller redundancy.
- B. They can maintain storage services after one controller fails.
- C. They can share I/O processing workload.
- D. They remove the need for backup protection.

ANSWER: A B C

Explanation:

A dual-controller storage system provides controller redundancy. If one controller fails, the surviving controller can continue providing storage services through the available paths, provided that the host connectivity, multipathing configuration, and storage-system design support failover. This reduces the risk that a single controller failure interrupts host access.

Dual controllers can also share I/O processing and cache resources, improving availability and potentially increasing performance according to the workload and controller architecture. However, dual controllers do not eliminate the need for backup or disaster-recovery protection. Controller redundancy protects against certain hardware faults, but it does not protect against logical corruption, accidental deletion, or a site-level disaster. See [Huawei OceanStor Data Storage](#).

QUESTION NO: 30

In a clustered NAS architecture, which of the following allows multiple file system to be consolidated and does not need to be managed separately?

- A. Global Namespace.
- B. Global Sharing.
- C. Global Cache.
- D. Global Concurrency.

ANSWER: A

Explanation:

Global Namespace is correct. In a clustered NAS architecture, a global namespace presents multiple underlying file systems as one unified, logical directory tree. Clients access data through a consistent path rather than needing to know which physical file system, storage node, or file server owns a particular directory. This consolidates otherwise separate file-system namespaces into a single view, simplifying client access and administrative operations. Administrators can organize data and scale storage across nodes while maintaining a common namespace, avoiding separate management from the user-access perspective. This is a fundamental clustered NAS capability: capacity and performance can grow by adding nodes or file systems, while the namespace remains unified and transparent to clients. Huawei describes scale-out NAS systems as providing a unified namespace for file services, enabling data to be accessed through a single logical file-system view. See Huawei's [OceanStor scale-out storage documentation](#) and the [Huawei scale-out storage overview](#).

QUESTION NO: 31

Which statements are true about the SAN network? (Select All that Apply)

- A. The FC SAN must be configured with dedicated iSCSI HBAs for servers
- B. The IP SAN can use the existing IP network.
- C. The IP SAN supports the iSCSI protocol.
- D. The FC HBA used by the FC SAN can convert SCSI packets into FC packets.

ANSWER: B C D

Explanation:

An IP SAN can use an organization's existing Ethernet/IP infrastructure, provided that the network is designed to deliver the required bandwidth, availability, and latency for storage traffic. This is a central benefit of IP-based storage: servers communicate with storage targets across standard TCP/IP networks rather than requiring a separate Fibre Channel fabric. iSCSI is specifically designed for this model. It transports SCSI commands and data over TCP/IP, allowing hosts and storage systems to provide block-storage services through IP networking.

An FC HBA used in an FC SAN also performs the host-side protocol processing needed to carry SCSI block commands over Fibre Channel. In the Fibre Channel protocol stack, SCSI is mapped at the FC-4 layer and is encapsulated for transmission as Fibre Channel frames. Therefore, describing the adapter as converting SCSI packets into FC packets accurately expresses its practical protocol-translation role. Huawei storage documentation describes FC and iSCSI as supported host connectivity protocols, while the IETF iSCSI specification defines the transport of SCSI over TCP/IP. See [Huawei OceanStor documentation](#) and [RFC 3720: iSCSI](#).

QUESTION NO: 32

The synchronous remote copy management module RM is software.

- A. TRUE
- B. FALSE

ANSWER: A

Explanation:

TRUE is correct. RM refers to the Remote Mirroring/remote replication capability used by Huawei storage systems to maintain a copy of data at a geographically separate storage system for disaster-recovery purposes. The synchronous remote-copy mode is implemented as a storage software feature: it coordinates host-write processing so that data is committed according to the synchronous replication policy at both the local and remote sites before the write is acknowledged. This functionality is configured, monitored, and managed through the storage system's management software and normally requires the relevant remote-replication feature license; it is not a separate physical hardware module installed in the array.

In practical deployment, the storage systems, their controllers, and the replication links provide the underlying infrastructure, while RM supplies the replication logic and management capability. This software implementation enables administrators to create and manage replication pairs, select synchronous behavior, observe pair status, and perform recovery-related operations. Huawei describes its enterprise storage offerings as providing replication and business-continuity capabilities through storage features and management functions. See Huawei's [OceanStor Dorado product information](#) and the [Huawei Enterprise Storage technical-support portal](#) for product documentation and feature references.

QUESTION NO: 33

Regarding the traditional model and cloud computing model, high reliability in the data center and cross-data center disaster recovery methods, the following says

What is wrong?

- A.** In the traditional mode, the data center relies on component redundancy to achieve high reliability
- B.** In the cloud computing mode, the data center relies on component redundancy to achieve high reliability
- C.** In the traditional model, cross-data center disaster recovery is mainly based on copying business data, and the recovery process is complicated and the recovery period is long
- D.** In the cloud computing mode, cross-data center disaster recovery is the replication of all data in the business data and operating environment.
Simple recovery process and short recovery period

ANSWER: B

Explanation:

In the cloud computing mode, the data center relies on component redundancy to achieve high reliability is the incorrect statement. Component redundancy—such as redundant power supplies, fans, network links, and individual hardware modules—is a fundamental availability approach in a traditional data center. Cloud computing data centers use this foundation but achieve their principal high-reliability capability through a broader cloud architecture: resource pooling, virtualization, distributed software, automated fault detection and recovery, workload scheduling, and redundancy at the service or virtual-machine level. These mechanisms enable workloads to be restarted, migrated, replicated, or rescheduled when underlying infrastructure fails, reducing dependence on any single physical component. This architectural distinction is important because cloud reliability is designed around service continuity and fault isolation across shared infrastructure, rather than simply making each hardware component redundant. Huawei Cloud describes its disaster recovery and replication capabilities as cloud services that protect workloads and data through managed replication and recovery mechanisms, reflecting this service-oriented model. See Huawei Cloud's [Data Replication Service overview](#) and the [Huawei Cloud documentation portal](#) for official service and architecture documentation.

QUESTION NO: 34

The dual switch network design is widely used in the FC SAN environment. The following description of the dual switch network redundancy design,

Which ones are correct? (Multiple choice)

- A. Deploy two FC switches in a SAN environment to cross connect to the two controllers of the storage device
- B. Each host requires 2 available FC ports, and connects the two ports to different switches
- C. The host port and storage device port on the switch that need to communicate must be in the same zone
- D. Although there are multiple paths, the dual-switch network only supports link redundancy and does not support load balancing

ANSWER: A B C

Explanation:

A dual-switch Fibre Channel SAN is designed around independent fabrics so that the loss of one switch, one host HBA port, one cable, or one storage-side connection does not remove host access to storage. Deploying two FC switches and connecting them independently to the storage system's controller ports creates the required redundant storage access paths. Likewise, a host needs separate available FC ports, with each port attached to a different switch fabric, so that there is no single-switch dependency.

Zoning is an essential part of this design. A host initiator port and the storage target port it must access have to be members of the same zone; this controls visibility and permits the intended FC communication. After the SAN paths are established, host multipathing software recognizes the paths to the same logical unit and manages path failover. Where supported by the storage array, host operating system, and multipathing policy, these paths can also be used for load balancing, in addition to redundancy. Huawei storage documentation describes redundant host links, zoning, and multipathing as core host-connectivity practices. See the [Huawei OceanStor host connectivity documentation](#) and the [Huawei SAN configuration documentation](#).

QUESTION NO: 35

Xiao Wang, an employee of a company, did the following in the company's storage system: 9:00 Take a snapshot of LUN A and activate it

Alive, all the data of the snapshot LUN and the source LUN were modified at 10:00, and the snapshot was rolled back at 11:00.

The following instructions are correct? (Multiple choice)

- A. The data in the copy space before COW is 9:00 LUN A
- B. The data in the copy space before COW is 10:00 LUN A data
- C. The data of the source LUN is the modified data at 10:00
- D. After the snapshot is rolled back, the data of the source LUN becomes the data of the modified snapshot LUN at 10:00

ANSWER: A D

Explanation:

"The data in the copy space before COW is 9:00 LUN A" is correct because a copy-on-write snapshot preserves the original source-LUN block contents when those blocks are first changed after the snapshot is created. Therefore, the copy space holds the pre-modification version of the affected blocks: the state of LUN A at the 9:00 snapshot point. This preserved data enables the snapshot point-in-time view to remain available even after subsequent writes occur.

"After the snapshot is rolled back, the data of the source LUN becomes the data of the modified snapshot LUN at 10:00" is also correct. A rollback restores the source LUN using the snapshot's data state at the time rollback is performed. Since the snapshot LUN was modified at 10:00, its current data is the rollback source; consequently, after rollback at 11:00, the source LUN reflects that modified snapshot-LUN data rather than retaining its independently modified source-LUN contents.

This behavior follows the fundamental COW snapshot principle: preserve old source data in copy space for the snapshot view, then use the snapshot's current data when executing rollback. See Huawei's [Snapshot documentation search](#) and the [OceanStor product documentation](#).

QUESTION NO: 36

In HUAWEI OceanStor, what unit is Thick LUN applying for space from the storage system?

- A. grain
- B. chunk
- C. stripe
- D. extent

ANSWER: D

Explanation:

extent is correct. In Huawei OceanStor block-storage architecture, a Thick LUN reserves its configured capacity from the storage pool when it is created. The storage system manages and allocates that LUN capacity in extents, which are the basic allocation and data-placement units used between the LUN and the pool's storage tiers. Thus, the LUN is constructed from extents obtained from a specified tier or from capacity managed by the pool. Extent-based allocation also enables OceanStor to place and relocate data at an appropriately granular level as policies require. For example, when SmartTier is in use, the system can analyze access characteristics and move data among tiers in extent units, helping place active data on higher-performance media while retaining less-active data on lower-cost media. This makes an extent the relevant unit both for Thick LUN space allocation and for tier-oriented data management. Huawei's OceanStor product information describes the platform's unified storage-pool and intelligent tiering capabilities: [Huawei OceanStor Dorado](#). Additional official product and technical resources are available through the [Huawei Enterprise Support](#) portal.

QUESTION NO: 37

What are the main functions of the business continuity solution? (Select All that Apply)

- A. Minimize the impact on the business service.
- B. Predict the disaster occurrence accurately.
- C. Zero data loss and zero service interruption.
- D. Improve the utilization rate of the disaster recovery equipment

ANSWER: A C D

Explanation:

Business continuity solutions are designed to keep critical services available through failures and site-level disasters while protecting the consistency and recoverability of production data. "Minimize the impact on the business service." is therefore a core function: continuity architecture uses redundant components, replication, automated switchover, and recovery procedures to reduce the business effects of an outage. "Zero data loss and zero service interruption." expresses the target capability of an active-active or equivalent high-availability continuity design, where synchronous replication and transparent service takeover can deliver an RPO of zero and, in suitable designs, an RTO approaching zero. "Improve the utilization rate of the disaster recovery equipment" is also a business-continuity function because active-active and active-standby designs can allow recovery-site resources to provide services, host workloads, or otherwise contribute value during normal operation rather than remaining entirely idle. Huawei's storage continuity capabilities include replication and active-active solutions intended to maintain application availability and data protection across failures. See Huawei's [data protection solutions](#) and [enterprise storage portfolio](#) for continuity and disaster-recovery capabilities.

QUESTION NO: 38

Which of the following is the end that receives and processes the SCSI commands?

- A. LUN.

- B. Storage.
- C. Target.
- D. Initiator.

ANSWER: C

Explanation:

Target. is correct because, in the SCSI client-server model, the target is the endpoint that receives SCSI commands sent over a storage protocol and processes those commands. A host-side SCSI initiator issues operations such as reads, writes, inquiries, and capacity requests. The target endpoint accepts the requests through a target port, interprets the command descriptor blocks, performs the requested operation against its backed storage resources, and returns status and any requested data. In a storage area network, the target function is commonly provided by a storage array's front-end ports or by a server exporting block storage through a protocol such as Fibre Channel or iSCSI. The logical units made available by that target are addressed as LUNs, allowing the host to access storage through the target's SCSI service. This initiator-target relationship is a fundamental part of SCSI architecture and is retained across common block-storage transports. The SCSI Architecture Model defines a target as the device-side entity that contains logical units and processes commands received from an initiator. See the [SCSI Architecture Model](#) and Huawei's [storage product documentation](#) for terminology and storage-network concepts.

QUESTION NO: 39

Which of the following descriptions of RAID technology are correct? (Multiple choice)

- A. Not every RAID technology can improve write performance
- B. In RAID 0, multiple hard drives are accessed simultaneously (parallel access), speeding up data access performance
- C. The current RAID technology only uses the parity mechanism to protect the data
- D. The performance of all RAID levels increases linearly with the increase in the number of member disks

ANSWER: A B

Explanation:

"Not every RAID technology can improve write performance" is correct because RAID levels use different data-placement and protection methods. Striping can distribute write I/O across several disks, but parity-based implementations must also maintain redundancy information. In particular, small writes may require additional reads and writes to update parity, so RAID should not be treated as a universal write-performance improvement. The effective result depends on the RAID level, write pattern, controller cache, stripe size, and number and type of member disks.

"In RAID 0, multiple hard drives are accessed simultaneously (parallel access), speeding up data access performance" is also correct. RAID 0 divides data into stripes and distributes those stripes among member disks. For suitably parallel sequential or large I/O workloads, several disks can transfer different portions of the data at the same time, increasing aggregate throughput. RAID 0 is designed for performance through striping and does not provide redundancy, making it appropriate only when the required availability is provided elsewhere or the data can be recreated. These characteristics are consistent with standard RAID terminology and storage implementation guidance. See the [IBM documentation on RAID levels](#) and the [SNIA RAID overview](#).

QUESTION NO: 40

A bank will launch a business system, which has very high performance and security requirements, and the database uses Oracle data

Database and plan to deploy the Oracle database to a storage system, then create a RAID group for the Oracle database

At this time, the optimal RAID strategy should be:

- A. RAID 1
- B. RAID 5
- C. RAID 6
- D. RAID 10

ANSWER: D

Explanation:

RAID 10 is the optimal choice for an Oracle database supporting a bank's critical business system because it combines disk mirroring with striping. Mirroring maintains a duplicate copy of every data block, giving strong fault tolerance and fast recovery when a drive fails. Striping distributes I/O across multiple disks, which increases throughput and lowers latency for the mixed, frequently random read/write workload typical of an online transaction database. RAID 10 also avoids parity-calculation overhead during writes, helping it deliver predictable write performance for database redo, data-file updates, indexes, and transaction processing. These characteristics are particularly important where service continuity and stable response times are priorities. In Huawei storage planning, RAID selection must balance performance, availability, capacity efficiency, and the workload characteristics; mission-critical databases commonly prioritize the performance and resiliency provided by mirrored storage. Oracle likewise recommends designing database storage for the required availability and I/O performance profile. See Huawei's [storage documentation portal](#) and Oracle's [database documentation library](#) for storage-planning and database deployment guidance.

QUESTION NO: 41

The cloudification of IT architecture is from "separation" to "convergence". What does convergence mean?

- A. Convergence of IT infrastructure
- B. Application and server integration
- C. Integration of hardware and software
- D. Multi-application sharing and open "cloud computer"

ANSWER: D

Explanation:

"Multi-application sharing and open 'cloud computer'" correctly describes convergence in a cloudified IT architecture. Convergence means moving away from isolated, purpose-built infrastructure stacks in which a specific application is tightly associated with dedicated compute, storage, and networking resources. Instead, resources are pooled, virtualized, and delivered through a common platform so that multiple applications can consume shared infrastructure according to changing workload requirements.

This model creates an open cloud-computing environment: resources can be provisioned and managed centrally, capacity can be assigned dynamically, and standardized infrastructure capabilities can support many services rather than a single silo. The essential outcome is broader resource sharing and coordinated delivery of computing services, enabling greater flexibility and efficiency as applications evolve. Huawei describes cloud infrastructure as a foundation for pooling and using ICT resources through cloud technologies, supporting agile and efficient service delivery. See Huawei's [cloud-computing solutions overview](#) and its [Elastic Cloud Server product documentation](#) for examples of cloud resources being centrally managed and provided to workloads on demand.

QUESTION NO: 42

In data backup and recovery, application testing and data analysis scenarios. HyperClone supports up to 1 master: 8 slave LUN mode. The knowledge points come from HCNA-Storage V3.0 textbooks and official product documentation.

The following descriptions for SAN and NAS are correct: (multiple choices)

- A. SAN is more used in the database scenario to store this structured data.
- B. Now NAS is for unstructured data storage, such as the sharing of key office documents and centralized storage by many department employees. file server.
- C. DAS, NAS, and SAN are all external storage, so storage resource sharing can be easily achieved.
- D. The characteristic of NAS is to put the file system of the server in the storage, so that the server needs to share data through the network.

ANSWER: A B D

Explanation:

“SAN is more used in the database scenario to store this structured data.” is correct because a storage area network commonly provides block storage to hosts. Block LUNs can be formatted and managed by the server or database platform, supporting the performance, low-latency access, and controllable data-management requirements typical of transactional databases and other structured-data workloads.

“Now NAS is for unstructured data storage, such as the sharing of key office documents and centralized storage by many department employees. file server.” is also correct. NAS provides file-level access over an IP network, commonly through file-sharing protocols such as NFS or SMB/CIFS. This makes it well suited to shared office files, home directories, file-server consolidation, media, and other unstructured data.

“The characteristic of NAS is to put the file system of the server in the storage, so that the server needs to share data through the network.” correctly describes the essential NAS model: the storage system manages and exports file systems, while clients access shared files through the network rather than directly managing a block device's file system. Huawei's enterprise storage portfolio and documentation describe both block-oriented SAN and file-oriented NAS services; see [Huawei Enterprise Storage](#) and [Huawei Enterprise Support Documentation](#).

QUESTION NO: 43

What is the main purpose of backup battery units (BBU) in storage controllers?

- A. To maintain the power for clock synchronization.
- B. To power the interface modules.
- C. To provide additional power to storage controller to speed up the I/O.
- D. To ensure that the non completed I/O is not lost during power outage.

ANSWER: D

Explanation:

To ensure that the non completed I/O is not lost during power outage is correct because a backup battery unit protects data held in a storage controller's cache when external power fails. Controllers commonly acknowledge host write requests after placing the write data in protected cache, before that data has been permanently written (destaged) to disks or SSDs. The BBU supplies temporary power to preserve the cache contents and, depending on the controller design, permits the controller to safely retain or flush those pending writes. Once power is restored, the controller can complete destaging the protected data to persistent storage, maintaining write consistency and preventing data loss caused by an abrupt outage.

This cache-protection role is fundamental to reliable write-back caching. It allows the system to improve write performance without sacrificing the durability expected by hosts and applications. Storage systems monitor battery health and charge state because a failed or insufficiently charged BBU can require the controller to limit or disable write-back caching to protect data integrity. Huawei storage documentation and support resources cover controller-cache protection, component status, and maintenance considerations: [Huawei Enterprise Storage Support](#) and [Huawei Support search for BBU documentation](#).

QUESTION NO: 44

What does the term RPO mean?

- A. Return Purchase Order
- B. Recover Plan Objective
- C. Restore Point Objective
- D. Recovery Point Objective

ANSWER: D

Explanation:

Recovery Point Objective is the correct expansion of RPO. In data protection and disaster-recovery planning, it defines the maximum acceptable amount of data loss measured in time. Put another way, it identifies how far back in time recovered data may be relative to the moment of a disruption. For example, an RPO of one hour means that backup, replication, or another protection mechanism must ensure that no more than one hour of newly created or changed data is lost if a failure occurs. RPO is therefore used to determine the required frequency of backups or the replication interval for an application or service. A lower RPO requires more frequent protection and commonly demands storage, network, and management resources capable of meeting that objective. It is distinct from the time required to restore service; RPO concerns data currency at recovery. Huawei's storage data-protection materials describe recovery objectives as key considerations when selecting backup and replication solutions. For general disaster-recovery terminology, see [NIST's Recovery Point Objective definition](#) and Huawei's [OceanStor data protection documentation](#).

QUESTION NO: 45

Which of the following are correct steps in a LAN based backup job?

(Multiple Choice)

- A. Agent on application server sends data to the backup device.
- B. Agent on application server sends data to backup media server,
- C. Backup media server sends data to the backup device.
- D. Backup server sends a "start-backup" command to the backup device.

ANSWER: B C

Explanation:

In a LAN-based backup architecture, the application server runs an agent that reads protected data and transfers the backup stream across the production LAN to the backup media server. Therefore, **Agent on application server sends data to backup media server**, describes the data-movement path from the protected host to the system responsible for media handling. The backup media server then receives that stream, performs its media-management role, and writes the data to the configured backup target. Consequently, **Backup media server sends data to the backup device**. is also a required step in the normal LAN backup data path.

This division of responsibilities is fundamental to enterprise backup design: application-side agents ensure that application or file-system data can be consistently collected, while a media server centralizes the connection to and handling of backup storage, such as tape libraries or disk-based backup devices. The backup server commonly coordinates policies, schedules, catalogs, and job control, whereas the data itself follows the application-server-to-media-server-to-backup-device flow. This design enables centralized management and permits multiple protected hosts to use shared backup infrastructure. See Huawei's enterprise support portal for storage and data-protection documentation at [Huawei Enterprise Support](#) and the SNIA overview of storage and data-protection concepts at [SNIA: What Is Storage?](#).

QUESTION NO: 46

A government unit adopts Huawei's two-site, three-center solution for disaster recovery construction, and its business is mainly concentrated from 8:00 to 18:00.

What is wrong with the following approach?

- A. Production centers and remote disaster recovery centers should use synchronous remote replication
- B. The time period for data synchronization between the production center and the disaster recovery center should be selected when the business volume at night is small
- C. When doing remote replication, a consistency group should be created according to the business of the unit, and related businesses are placed in a consistency group
- D. You can use the BCManager software to set the data synchronization period

ANSWER: A

Explanation:

Production centers and remote disaster recovery centers should use synchronous remote replication is the wrong approach. In a two-site, three-center design, the production site and the nearby disaster-recovery site are normally the pair suited to synchronous replication, because synchronous replication requires low latency and sufficient link bandwidth to ensure that writes are committed at both ends before the host receives acknowledgment. The remote disaster-recovery center is generally deployed farther away to protect against regional-level failures. It therefore normally receives data through asynchronous replication, which avoids making production application write latency dependent on a long-distance transmission link. For an organization whose main workload runs from 8:00 to 18:00, asynchronous replication or its synchronization activity can also be planned to reduce impact during periods of high business activity while still meeting the defined recovery point objective. Huawei storage disaster-recovery capabilities provide synchronous and asynchronous remote replication modes so that a design can match replication mode to distance, latency, bandwidth, and service recovery requirements. See Huawei's [OceanStor Dorado storage overview](#) and the [Huawei Enterprise Support documentation portal](#) for product documentation on remote-replication and disaster-recovery planning.

QUESTION NO: 47

LAN Free and Server Free are backup networking methods that have little impact on host performance.

- A. TRUE
- B. FALSE

ANSWER: A

Explanation:

TRUE is correct. LAN-free and server-free backup are SAN-oriented backup architectures designed to avoid the conventional pattern in which backup data travels from a production host across the Ethernet LAN to a backup server and then to backup media. In LAN-free backup, the host can send backup data to the backup target through the SAN, substantially reducing LAN traffic and avoiding the network congestion that can affect production services. The application host still participates in reading and supplying the data, but the backup path is more efficient and typically imposes less network-related overhead on the host environment.

Server-free backup takes this reduction further by using SAN-based third-party copy capabilities. The storage system or SAN data-movement function transfers data between source storage and the backup target without requiring the application host to carry the backup data stream. This minimizes host CPU, memory, and I/O involvement and helps preserve host resources for business workloads. These approaches require compatible SAN infrastructure and backup software, but their central purpose is to reduce the impact of backup processing and backup-network traffic on production hosts. Huawei's storage documentation and enterprise support resources describe SAN storage and data-protection capabilities relevant to such designs: [Huawei Enterprise Storage Support](#) and [Huawei Enterprise Storage Solutions](#).

QUESTION NO: 48

The following description about CIFS and NFS sharing is incorrect: (multiple choices)

- A. CIFS is a protocol based on network sharing, which has high requirements on the reliability of network transmission, so it is usually adopted Transmission using TCP protocol
- B. One disadvantage of CIFS is that Windows clients must install professional software
- C. NFS is a stateless protocol, and CIFS is a stateful protocol. NFS can recover from link failures. Recovery, and CIFS cannot
- D. Regardless of the CIFS or NFS protocol, the client needs to convert the file format

ANSWER: B D

Explanation:

“One disadvantage of CIFS is that Windows clients must install professional software” is an incorrect statement because Windows includes native SMB client functionality. CIFS is an older dialect within the SMB file-sharing family, and modern Windows systems can connect to SMB/CIFS shares through built-in Explorer, command-line, and operating-system networking components; a separate professional client is not inherently required. Microsoft documents SMB as Windows’ native network file-sharing protocol and its client-side use.

“Regardless of the CIFS or NFS protocol, the client needs to convert the file format” is also incorrect. CIFS/SMB and NFS are network file-system sharing protocols: they define how clients request operations such as opening, reading, writing, locking, and enumerating files over a network. They do not require universal conversion of the underlying file’s content format merely to access a share. An application may need to understand a particular document or media format, but that is separate from the CIFS/SMB or NFS sharing protocol. NFS specifies remote file access procedures rather than a mandatory client file-format conversion process. See [Microsoft’s SMB overview](#) and [RFC 1813: NFS Version 3 Protocol Specification](#).

QUESTION NO: 49

When does data become information?

- A. When it has been stored in the correct file format
- B. When the data has been stored in digital form
- C. When the enterprise understands the meaning of data and extracts the content needed for decision-making
- D. When the data is verified to be correct

ANSWER: C

Explanation:

Data becomes information when it is interpreted in context so that its meaning is understood and it can support a purpose such as decision-making. Raw data consists of facts, values, records, or observations without inherent business significance on its own. An enterprise turns those facts into information by organizing, processing, analyzing, and contextualizing them to extract content that is relevant to a user or business objective. For example, individual storage-capacity measurements are data; identifying a sustained capacity-growth trend and using it to plan an expansion converts those measurements into actionable information. This distinction is fundamental to storage and data-management work because the value of stored data is realized when it can be reliably interpreted and used by people, applications, and business processes. The option stating that the enterprise understands the meaning of data and extracts the content needed for decision-making therefore correctly defines the transition from data to information. This aligns with the widely used definition of information as data that has been processed, organized, or structured in a meaningful context; see [Encyclopaedia Britannica’s overview of information systems](#) and [TechTarget’s definition of information](#).

QUESTION NO: 50

The following two descriptions are correct:

Description 1: The maintenance cost of traditional PC is usually lower than that of desktop cloud.

Description 2: In cloud computing, data is stored and managed centrally.

- A. Descriptions 1 and 2 are correct
- B. Description 1 is correct, description 2 is incorrect
- C. Description 1 is wrong, description 2 is correct
- D. Both descriptions 1 and 2 are incorrect

ANSWER: C

Explanation:

Description 1 is wrong, description 2 is correct is correct. Cloud computing centralizes computing resources, applications, and data in a provider-managed data center or cloud platform. Rather than placing business data separately on each endpoint, organizations can store it in centralized storage services and administer access, backup, security controls, and lifecycle policies from a unified platform. This centralization is a core cloud characteristic because it makes resources easier to manage consistently, protect, and scale.

For desktop cloud specifically, user desktops commonly run as centrally hosted virtual desktops while users access them from endpoints. Centralized management reduces the need to maintain operating systems, applications, security configurations, and data individually across many physical PCs. Administrators can apply policies and perform updates centrally, which generally reduces operational effort and long-term maintenance costs compared with managing a large traditional-PC fleet. Huawei Cloud Workspace describes this centralized desktop-cloud approach and its centralized management capabilities. See [Huawei Cloud Workspace](#) and the [Huawei Cloud cloud-computing overview](#).

QUESTION NO: 51

Which are characteristics of the HUAWEI OceanStor hardware?

(Multiple Choice)

- A. SAN and NAS convergence
- B. SSD and HDD convergence
- C. No need for heterogeneous convergence anymore
- D. Primary and Backup convergence

ANSWER: A B D

Explanation:

Huawei OceanStor is designed as a converged storage platform, so **SAN and NAS convergence**, **SSD and HDD convergence**, and **Primary and Backup convergence** are valid hardware characteristics. SAN and NAS convergence means that a unified OceanStor system can provide both block-storage and file-storage services, allowing an organization to use one platform for workloads that require SAN access and workloads that require NAS access. SSD and HDD convergence reflects the use of heterogeneous media within a storage system: SSDs deliver high performance for latency-sensitive data, while HDDs provide economical capacity, with storage-tiering capabilities helping place data appropriately. Primary and Backup convergence describes the platform's ability to support production storage together with data-protection and backup-oriented service requirements, reducing the need for separate infrastructure silos. These converged capabilities simplify infrastructure design, improve resource utilization, and enable centralized management as service requirements evolve. Huawei's OceanStor portfolio information is available at [Huawei OceanStor Storage](#), and Huawei's enterprise storage documentation portal is available at [Huawei Storage Support Documentation](#).

QUESTION NO: 52

Clustered NAS can provide concurrent multi-node access of data to hosts.

- A. TRUE
- B. FALSE

ANSWER: A

Explanation:

TRUE is correct. Clustered NAS, also called scale-out NAS, combines multiple storage nodes into a unified NAS cluster and presents file-system services to clients through standard NAS protocols such as NFS and SMB. The cluster's nodes work together to process client requests, so hosts can concurrently access data while workload processing is distributed across multiple nodes. This architecture enables capacity and performance to scale by adding nodes, rather than being limited to the processing capability of a single controller pair. It also provides a single logical storage resource and namespace from the client perspective, even though the underlying data and service workload can be handled by different cluster nodes. In practical deployments, concurrent multi-node access improves throughput for large numbers of file clients and helps balance access load across the cluster. Huawei describes its scale-out storage offerings as supporting scale-out architecture and unified file services, which are the fundamental capabilities behind clustered NAS access. See Huawei's [OceanStor Dorado storage portfolio](#) and the [Huawei Scale-Out Storage support page](#) for product and technical information on clustered, scale-out storage capabilities.

QUESTION NO: 53

Which of the following are the 3 layers of components that storage systems are made up of?

(Multiple Choice)

- A. Storage Analysis
- B. Storage Solution
- C. Storage Software
- D. Storage Hardware

ANSWER: B C D

Explanation:

Storage systems are commonly described through the complementary layers of **Storage Hardware**, **Storage Software**, and **Storage Solution**. Storage Hardware is the physical foundation that provides capacity, connectivity, processing, and reliability. It includes elements such as storage controllers, disk enclosures, drives, interface modules, and power and cooling components. Storage Software operates on and manages that physical foundation. It delivers the logical storage functions that make the platform usable, including RAID or other data-layout mechanisms, volume and pool management, host access services, snapshots, replication, performance management, alarms, and system administration. Storage Solution is the higher-level layer that applies storage hardware and software capabilities to a business or application requirement. It combines the storage platform with design, deployment, data-protection, connectivity, and operational practices to meet requirements for availability, performance, scalability, and recovery. Together, these layers represent the progression from physical infrastructure, to storage functions, to an implementable service or application-oriented solution. Huawei's enterprise storage materials describe storage products and their software-enabled data services in this broader solution context: [Huawei Enterprise Storage](#) and [Huawei Storage Solutions](#).

QUESTION NO: 54

Statement 1: Multipathing is caused by adding redundant hardware.

Statement 2: Huawei offers a different version of Ultra Path for different operating systems.

- A. Statements 1 and 2 are both true.

- B. Statement 1 is true and statement 2 is false.
- C. Statement 1 is false and statement 2 is true.
- D. Both statements 1 and 2 are false.

ANSWER: A

Explanation:

Statements 1 and 2 are both true. Multipathing exists when a host has more than one physical I/O path to the same storage resource. Those paths are created by redundant hardware and connectivity components, such as multiple host bus adapters, network interface cards, storage-controller ports, switches, cables, and logical paths. This redundant design enables a host to continue accessing storage if an individual link, port, adapter, or controller path fails. Multipathing software also identifies the available paths and applies path-selection and failover policies, allowing the host to use the redundant infrastructure reliably and, where supported, distribute I/O across paths.

Huawei UltraPath is Huawei host multipathing software for storage systems. Because host operating systems have different kernel, driver, installation, device-management, and path-management requirements, Huawei supplies UltraPath packages and documentation tailored to supported operating systems rather than one interchangeable package for every platform. This OS-specific delivery ensures that the multipathing component integrates correctly with the host environment and Huawei storage arrays. Huawei's official support portal provides UltraPath software and host-connectivity documentation by platform: [Huawei Enterprise Support—UltraPath search](#). Huawei also describes its storage portfolio and host connectivity ecosystem at [Huawei Enterprise Storage](#).

QUESTION NO: 55

The maintenance network port is used, by the DeviceManager, for daily management.

- A. TRUE
- B. FALSE

ANSWER: B

Explanation:

FALSE is correct because DeviceManager daily administration uses the storage system's management network interface/IP address, rather than the maintenance network port. Administrators access DeviceManager through a browser over the management network to perform routine activities such as viewing system health, configuring storage resources, managing hosts, checking alarms, and monitoring performance. The management interface is designed to provide the normal management communication path between administrators and the storage system.

The maintenance network port has a separate service-oriented role. It is intended for maintenance operations and service access, typically by authorized support personnel or maintenance tools when performing deployment, diagnosis, troubleshooting, or repair-related tasks. Separating the management and maintenance communication paths helps keep routine administration distinct from specialized service procedures and supports controlled access to the storage device. Huawei's storage product documentation distinguishes management access from maintenance/service access in its hardware installation, cabling, and management guidance. See the [Huawei Enterprise Storage Support portal](#) and the [Huawei Enterprise documentation portal](#) for product-specific port definitions and DeviceManager administration documentation.