

HCIP-Storage V5.0

Huawei H13-624

Version Demo

Total Demo Questions: 80

Total Premium Questions: 802

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Storage Technology Trends	36
Topic 2, Storage System Architecture	134
Topic 3, Storage Networking Technology	39
Topic 4, Storage System Installation and Configuration	64
Topic 5, Storage System Maintenance and Management	83
Topic 6, Storage System Performance Tuning	31
Topic 7, Storage System Advanced Features	415
Total	802

QUESTION NO: 1

When connecting a service host to a Huawei OceanStor storage system through a Fibre Channel SAN, which of the following practices are recommended?

- A. Use HBA drivers and firmware that are compatible with the host and storage system.
- B. Use single-initiator single-target zoning where applicable.
- C. Install and correctly configure host multipathing software.
- D. Use only one Fibre Channel path to prevent path conflicts.

ANSWER: A B C

Explanation:

The host HBA driver and firmware must be compatible with the operating system and storage system. Single-initiator single-target zoning is commonly used because it reduces unnecessary device visibility and simplifies fault isolation. Host multipathing should also be installed and configured so that the host can use redundant paths and tolerate an HBA, cable, switch-port, or controller-path fault.

Using one path only creates a single point of failure and is not recommended for an enterprise FC SAN deployment. Consult the [Huawei Enterprise Storage support portal](#) for the applicable interoperability and host-connectivity documentation.

QUESTION NO: 2

The correct statement about the Huawei Oceanstor 9000InfoTier node pool is: (multiple choice)

- A. The node types in the node pool must be the same
- B. If it is an OPS intensive application scenario, it is better to configure P nodes in the node pool
- C. When the node pool is composed of P nodes, a node pool contains at least 3 nodes

ANSWER: A B C

Explanation:

An InfoTier node pool is a group of storage nodes that provides a defined performance and capacity tier in an OceanStor 9000 file system. Nodes within one pool must be of the same type so that the tier has predictable hardware characteristics, capacity behavior, and performance. For an IOPS-intensive workload, a pool built with P nodes is appropriate because this node type is intended for performance-oriented configurations and supports the higher random-access capability such workloads require. A pool consisting of P nodes must contain at least three nodes. This minimum establishes the distributed-storage foundation needed for data placement and availability within that tier, rather than treating a single node or a pair of nodes as a production node pool. These design rules allow InfoTier to place data in a pool whose node resources match the workload's access profile while maintaining the scale-out architecture required by OceanStor 9000. Huawei describes OceanStor 9000 as a scale-out storage platform, and its product materials emphasize clustered nodes and policy-driven data management; see the [Huawei OceanStor 9000 product page](#) and the [Huawei OceanStor 9000 support documentation portal](#).

QUESTION NO: 3

Designed FlashLink technology. Which of the following statements are correct for the cold and hot data diversion technology ()

- A. Effectively reduce in-disk garbage collection
- B. Metadata and data are stored in separate partitions
- C. Synchronously execute independent partitions of different types of data in the controller and the SSD disk

D. The controller automatically senses the data layout in the SSD disk

ANSWER: A B C D

Explanation:

FlashLink cold and hot data diversion is a coordinated controller-and-drive design intended to make NAND flash use more efficient and predictable. "Metadata and data are stored in separate partitions" is correct because metadata has a substantially different access pattern from user data. Placing these classes separately helps prevent frequent metadata updates from disturbing data blocks that have lower update frequency. "Synchronously execute independent partitions of different types of data in the controller and the SSD disk" is also correct: the controller's allocation and scheduling policies are designed to align with the physical data-placement and partitioning behavior in the solid-state drive.

"The controller automatically senses the data layout in the SSD disk" is correct because FlashLink uses coordinated awareness between the storage controller and Huawei-developed solid-state drives, rather than treating the drive as a completely opaque device. That awareness enables optimized placement and lifecycle handling. As a result, "Effectively reduce in-disk garbage collection" is correct as well. Keeping data with similar update characteristics together reduces unnecessary valid-data movement during garbage collection, helping lower write amplification and maintain performance consistency. Huawei describes FlashLink as an end-to-end flash optimization architecture for OceanStor all-flash storage; see [Huawei all-flash storage](#) and [Huawei FlashLink technology](#).

QUESTION NO: 4

In the feature of Huawei Oceanstor 9000 InfoReplicator, Pair is a data replication relationship, which describes data replication

The source and destination of the, the following statement about the state transition of the Pair is incorrect (multiple choice)

- A. In the case of Pair failure, its status includes recovery, invalidation, synchronization, split
- B. When the abnormal link recovers and the recovery strategy is manual, the status of the Pair will be migrated to "synchronizing"
- C. When the InfoReplicator configuration is destroyed, the status of the Pair becomes "Invalid"

ANSWER: A B

Explanation:

The statements "In the case of Pair failure, its status includes recovery, invalidation, synchronization, split" and "When the abnormal link recovers and the recovery strategy is manual, the status of the Pair will be migrated to "synchronizing"" are correctly selected as incorrect. InfoReplicator treats a Pair as a replication relationship whose status records both its replication progress and whether intervention is required. A link or replication fault does not mean that every listed operational state is a failure state. In particular, synchronization and split are normal, distinct operating conditions that can be entered through replication actions; they are not a collective definition of Pair failure. Huawei's scale-out storage documentation describes InfoReplicator as a remote-replication function with controlled recovery behavior for replication relationships.

When recovery is configured as manual, restoration of a failed communication link does not itself restart data synchronization. The Pair must remain in a recovery-required state until an administrator initiates the recovery or synchronization operation. Only that explicit operation moves the relationship into synchronizing, allowing changed data to be reconciled in a controlled manner. This design prevents an unintended resynchronization after an outage and gives the administrator control over replication direction and recovery timing. See Huawei's [OceanStor 9000 product information](#) and the [Huawei Enterprise Storage support portal](#) for product documentation and management guidance.

QUESTION NO: 5

The object storage services provided by OceanStor9000 (compatible with single-select Amazon S3 and OpenStack Swift) are all

A unified RESTful API is provided externally, and the client can access storage nodes through STfulAPI to realize data storage.

- A. True
- B. False

ANSWER: A

Explanation:

True is correct. OceanStor 9000 provides object storage service interfaces compatible with Amazon S3 and OpenStack Swift. Both of these object-storage access models are RESTful: applications use HTTP-based requests to create containers or buckets, upload and retrieve objects, manage object metadata, and perform related object operations. Consequently, an application does not need to use a proprietary file-system client protocol simply to store and access objects. It sends requests through the applicable standard RESTful object API to the OceanStor 9000 cluster, whose storage nodes process the requests and distribute data according to the cluster's storage architecture. This API-based approach enables existing S3- or Swift-compatible applications to integrate with the system while retaining the object-oriented data-access model. Huawei identifies OceanStor 9000 as scale-out storage supporting object-storage workloads and cloud-oriented access, and its technical documentation describes object services based on S3 and Swift-compatible interfaces. See Huawei's [OceanStor 9000 product information](#) and the [OpenStack Object Storage API reference](#) for the REST-based Swift object-service model.

QUESTION NO: 6

When single-selecting the Oceanstor 9000 network tuning, which of the following commands can check whether the network port is connected?

- A. ifconfig
- B. ethtool
- C. lspci |grep Ether

ANSWER: B

Explanation:

`ethtool` is the appropriate command for checking the physical connection state of an Ethernet network port on OceanStor 9000 nodes. When run against a specific interface, for example `ethtool eth0`, it displays NIC properties and operational information, including the `Link detected` field. A result of `Link detected: yes` confirms that the interface has an active physical Ethernet link with the connected switch or peer; `no` indicates that no physical link has been established. This makes `ethtool` particularly useful during network tuning and fault isolation because it checks the link state at the adapter level and can also show negotiated speed, duplex mode, auto-negotiation status, and supported link modes. These details help verify that the cabling, transceiver, switch port, and NIC are operating with the expected connection parameters. The utility's documented purpose is to query and control Ethernet driver and hardware settings, including link-status reporting. See the [ethtool manual page](#) and the [Linux kernel ethtool documentation](#).

QUESTION NO: 7

Send a control command through the backup server, first take a snapshot (or other similar operations) of the data volume that needs to be backed up,

Make a copy, then use the backup server to mount the snapshot volume, and use the backup server to back up the data on this volume to the backup

On the device. What type of backup networking is described in the above workflow?

- A. LAN-Base

- B. LAN-Free
- C. Server-Less
- D. Server-Free

ANSWER: C

Explanation:

Server-Less is correct because the workflow removes the production host from the backup data path by first creating a storage-side snapshot or copy of the production volume. The backup server sends the control request that initiates the snapshot, then mounts the resulting snapshot volume and reads that independent point-in-time copy for transfer to the backup device. Consequently, backup processing does not need to read active production data through the application server, reducing disruption and resource consumption on that production host.

This model depends on storage snapshot or clone capabilities and on presenting the copied volume to the backup server. The snapshot preserves a consistent backup source while production I/O can continue on the original volume. The backup server remains involved in the workflow as the controller and as the host that mounts and backs up the snapshot; the defining characteristic is that the production server is not used for the actual backup read operation. Huawei storage documentation describes snapshot-based data protection and volume presentation mechanisms that provide the storage functions required by this design. See the [Huawei Enterprise Support documentation portal](#) and the [Huawei Storage documentation portal](#) for storage data-protection and host-access documentation.

QUESTION NO: 8

When RAID2.0+ is used, what is the default size of a chunk on a SAS disk?

- A. 256 MB
- B. 64 KB
- C. 512 MB
- D. 64 MB

ANSWER: D

Explanation:

The default chunk size for a SAS disk in Huawei RAID 2.0+ is **64 MB**. RAID 2.0+ uses a virtualization and data-distribution design in which storage capacity is divided into chunks before data is distributed across the disks in a storage pool. The chunk is the basic unit used by the system for allocating pool capacity and balancing data placement. For SAS disks, the default 64 MB value provides an appropriate balance between distribution granularity, system metadata management, and performance characteristics for enterprise workloads. This setting is associated with the disk type and RAID 2.0+ implementation rather than being an arbitrary RAID stripe-size value selected for each LUN. Understanding this distinction is important: a chunk supports RAID 2.0+ capacity virtualization and data layout across the storage pool, helping the system distribute data evenly and efficiently as resources are provisioned. Huawei storage training material for RAID 2.0+ identifies 64 MB as the default chunk size for SAS disks. See the referenced RAID 2.0+ material at [TitanWolf RAID 2.0+ reference](#) and Huawei's [enterprise storage documentation portal](#).

QUESTION NO: 9

Huawei 's hyper-converged storage has multiple data security assurance mechanisms. The following statement is correct ()

- A. Support a variety of EC (Erasure Code) with different ratios
- B. According to the security level, 2 copies or 3 copies can be flexibly configured
- C. Support storage dual-active and asynchronous replication capabilities, and provide cross-site data redundancy protection

D. PCIe SSD Cache technology ensures that data will not be lost after power failure

ANSWER: A B C

Explanation:

“Support a variety of EC (Erasure Code) with different ratios” is correct because Huawei distributed storage platforms use erasure coding schemes with selectable data-to-parity ratios, allowing protection overhead and usable capacity to be balanced for different service requirements. “According to the security level, 2 copies or 3 copies can be flexibly configured” is also correct: replica-based redundancy can be selected according to the required availability and protection level, with data distributed across failure domains. “Support storage dual-active and asynchronous replication capabilities, and provide cross-site data redundancy protection” is correct because Huawei storage disaster-recovery capabilities include active-active solutions for continuous service availability and asynchronous replication for remote data protection and recovery across sites. Together, these mechanisms address local disk or node failures, provide resilient distributed data layouts, and extend protection to site-level incidents. Huawei describes the availability, data-protection, and disaster-recovery capabilities of its distributed storage offerings at [Huawei FusionStorage](#). Huawei’s storage portfolio and associated data-protection technologies are also summarized at [Huawei Data Storage](#).

QUESTION NO: 10

(multiple choice) OceanStor 9000 adopts Erasure code technology to protect data. If a site has 6 nodes, and the customer hopes that even if three disks are broken in one strip, the data can still be restored. The available redundancy protection levels are as follows:

- A. 4+3:1
- B. 6+3
- C. 4+2
- D. 12+3:1

ANSWER: A D

Explanation:

4+3:1 and **12+3:1** provide three parity blocks for each erasure-coded stripe. In an N+3 erasure-coding layout, the three parity blocks supply the redundancy needed to reconstruct the original data when up to three blocks in the same stripe are unavailable. Therefore, if three disks fail and each failed disk affects a different block of one stripe, the system retains sufficient surviving data and parity information to restore that stripe.

The “:1” placement setting is applicable to the six-node deployment context and identifies a supported OceanStor 9000 protection-level configuration. Both 4+3:1 and 12+3:1 consequently meet the required three-disk fault tolerance while being available redundancy levels for this scenario. The key selection criterion is the presence of three parity blocks, because the requested recovery capability is specifically for three lost stripe blocks rather than merely general data availability.

For Huawei product and technical documentation, consult the [Huawei Enterprise Support documentation portal](#) and the [Huawei storage product documentation area](#).

QUESTION NO: 11

(Judgment) Whether the Huawei OceanStor 9000 is upgraded online or offline, it has no impact on device management.

- A. True
- B. False

ANSWER: B

Explanation:

False is correct because an OceanStor 9000 upgrade can affect management availability, even when the upgrade is designed to preserve storage services. During upgrade activities, system components and management processes may be restarted, nodes may enter maintenance or reboot states, and the management GUI or command-line interface can be temporarily unavailable or provide limited functionality. An offline upgrade necessarily includes a service interruption window and therefore cannot guarantee uninterrupted device management. Online upgrades are intended to minimize the impact on host-facing data services through rolling procedures, but this does not mean that every management function remains continuously available throughout the process. Administrators should therefore follow the applicable upgrade guide, perform the prescribed health checks, use the maintenance window and precautions defined for the target version, and avoid assuming that management access will be unaffected. Huawei's enterprise support portal provides the product documentation, upgrade packages, release notes, and maintenance guidance needed to plan these operations: [Huawei Enterprise Support](#). Huawei also identifies OceanStor 9000 as a scale-out storage product for which operation and maintenance procedures must be planned carefully: [Huawei OceanStor 9000](#).

QUESTION NO: 12

When selecting the type of data backup, which option is wrong in the following description?

- A.** If there are only cumulative incremental backups in the schedule, but no full backups, when performing cumulative incremental backups for the first time, all files will be backed up.Pieces.
- B.** It is recommended to use a combination of three backup types: full backup, differential incremental backup, and cumulative incremental backup.
- C.** If the backup strategy is full backup combined with differential incremental backup, after a system failure, if you want to restore to the most recent data, thenIt is necessary to restore the data of the last full backup and all differential incremental backups after the full backup.
- D.** Differential incremental can save storage space to the greatest extent, and the backup window is small.

ANSWER: B**Explanation:**

It is not generally recommended to use a single backup schedule that combines full backup, differential incremental backup, and cumulative incremental backup as three concurrent backup types. Differential incremental backup and cumulative incremental backup represent alternative incremental-chain methods with different reference points. A differential incremental backup protects data changed since the preceding backup, whereas a cumulative incremental backup protects data changed since the most recent full backup. A sound backup policy normally selects full backup plus one of these incremental methods according to the organization's recovery-time objective, recovery-point objective, available backup window, and storage capacity.

Using both incremental methods in one routine schedule can make retention, restore-chain selection, and operational management unnecessarily complex. It can also make it harder for administrators to identify the exact backup sets needed for a particular recovery point. Therefore, the statement recommending the combination of all three types is the incorrect description: it presents alternatives as though they should routinely be deployed together. Huawei's enterprise support portal provides product documentation and data-protection guidance at [Huawei Enterprise Support](#). General backup terminology and the distinction between full and incremental backup methods are also described by the [NIST guidance on data protection technologies](#).

QUESTION NO: 13

Which of the following descriptions of the software components of the Huawei OceanStor 9000 product is correct?

- A.** NDMP backup requires license support
- B.** The OceanStor9000 system defaults to load in a round-robin manner
- C.** InfoAllocator controls the user's use of storage space by capacity or number of files

ANSWER: C

Explanation:

InfoAllocator controls the user's use of storage space by capacity or number of files. OceanStor 9000 is a scale-out distributed storage system whose file-service features include quota management for controlling storage consumption. InfoAllocator is the component used to apply these quota policies to users and other relevant file-system objects. A quota can limit the amount of logical storage capacity that may be consumed, and it can also limit the number of files that may be created. This provides administrators with a practical way to prevent an individual user or workload from exhausting shared storage resources and helps maintain predictable capacity availability across a multi-tenant environment. The capacity-based and file-count-based controls describe the intended function of InfoAllocator directly, so this statement is the correct description of the OceanStor 9000 software component. Huawei's OceanStor 9000 product documentation and support resources describe the platform's distributed file-storage capabilities and its associated management features: [Huawei OceanStor 9000 support documentation](#) and [Huawei OceanStor 9000 product information](#).

QUESTION NO: 14

The built-in scripts of the disaster recovery management software OceanStor BCManagereReplication are for Oracle, SQL Server

The application supports one-click failure recovery. Which of the following steps is not included in one-click failure recovery?

- A. The test application starts successfully
- B. Forcibly switch the replication relationship
- C. Start the disaster recovery application
- D. Mapping LUN

ANSWER: D

Explanation:

Mapping LUN is not an execution step of BCManager eReplication one-click failure recovery. In BCManager eReplication, the disaster-recovery environment, including the host access configuration and required storage resources, is prepared when the protection and recovery resources are configured. One-click failure recovery then orchestrates the recovery actions that make the disaster-recovery copy usable and bring the protected application online. For built-in Oracle and SQL Server recovery scripts, this orchestration includes handling the replication relationship as required for a failure scenario, starting the disaster-recovery application, and verifying that the application has started successfully. The pre-established mapping information allows the recovery workflow to use the configured disaster-recovery resources without requiring an administrator to perform a separate LUN-mapping activity during the one-click operation. This separation between preconfiguration and automated recovery is important: it makes recovery repeatable, reduces manual intervention during an outage, and ensures that the application recovery script operates against the intended storage resources. Huawei's enterprise support portal provides product documentation and technical resources for OceanStor disaster-recovery management products; see [Huawei Enterprise Support](#). Huawei's storage product information is also available at [Huawei Storage](#).

QUESTION NO: 15

Which of the following statements about snapshots is true?

- A. A snapshot is a complete physical copy of source data.
- B. A snapshot is newly written data.
- C. Data of a snapshot can only be read by hosts and cannot be modified.
- D. A snapshot is a point-in-time copy of source data.

ANSWER: D

Explanation:

A snapshot is a point-in-time copy of source data. In Huawei storage systems, a snapshot preserves the logical state of a source LUN at the specific instant when the snapshot is created. It provides a recoverable historical view of the source data without first creating a full, independent physical copy of the entire LUN. Snapshot technology records the original data affected by subsequent writes and uses this preserved information together with the current source data to present the point-in-time image. This makes snapshots a fast and space-efficient mechanism for such tasks as backup support, application-consistent protection, testing, and recovery from accidental changes. Depending on the storage configuration and host access settings, a snapshot can also be mapped to hosts for use; point-in-time preservation, rather than read-only host access, is the defining characteristic. Huawei documentation describes snapshots as point-in-time copies and explains their copy-on-write based preservation behavior. See [Huawei OceanStor documentation](#) and [Huawei storage product documentation](#).

QUESTION NO: 16

In order to improve the reliability of its important business, an enterprise deployed in a data center room A and room B

With regard to the local active-active solution, which of the following statements is correct about this solution:

- A.** The active-active replication network in the two-machine room simultaneously carries data synchronization and heartbeat information between storage systems Sync and other various data
- B.** Computer room A and computer room B need to be in the same fault domain
- C.** The network types from the same business host to two active-active storage devices can be different
- D.** The arbitration server must use the host in computer room A or computer room B

ANSWER: A

Explanation:

The active-active replication network in the two-machine room simultaneously carries data synchronization and heartbeat information between storage systems Sync and other various data is correct. In Huawei local active-active storage deployments, commonly implemented with the HyperMetro feature, the two storage systems are interconnected by active-active replication links. These links have a dual role: they transmit the write data and synchronization information required to maintain consistent active-active copies, and they carry inter-array heartbeat/control communication used to monitor the peer storage system and coordinate active-active status. This communication is fundamental to providing both sites with access to the same business data while allowing the solution to detect failures and make the appropriate continuity decisions. Link design must therefore provide sufficient bandwidth and low, stable latency for synchronous data processing, as well as reliable connectivity for storage-system communication. Huawei's HyperMetro documentation describes the active-active relationship as maintaining data consistency between the paired storage devices and using inter-device links for active-active operation. See the [Huawei HyperMetro Feature Guide](#) and the [Huawei OceanStor HyperMetro documentation](#).

QUESTION NO: 17

What are the necessary components included in the architecture of the N800ONAS system? (multiple choice)

- A.** Disk array
- B.** NAS engine
- C.** Data sharing network equipment
- D.** FC switch

ANSWER: A B C D

Explanation:

The N8000 NAS solution uses a gateway-style architecture in which the NAS engine provides NAS protocols and file-system services to clients, while a disk array supplies the back-end block-storage capacity on which file data resides. Data sharing network equipment is required to provide the Ethernet/IP network path used by NAS clients to access shared file systems through protocols such as NFS and CIFS/SMB. An FC switch is also an architectural component because it provides the Fibre Channel SAN connectivity between the NAS engines and the back-end disk array. This separation of front-end Ethernet file access from back-end Fibre Channel block access enables the NAS engines to deliver file services while relying on the disk array for persistent storage, RAID protection, and capacity expansion. Therefore, the complete N8000 NAS architecture includes the disk array, NAS engine, data sharing network equipment, and FC switch. Huawei's enterprise support portal provides product documentation and technical resources for OceanStor storage systems at [Huawei Enterprise Storage Support](#). Huawei's storage product portfolio and architecture information are also available at [Huawei Storage Products](#).

QUESTION NO: 18

In the master-backup disaster recovery plan, the applications that can be protected by RD currently include: (multiple choice)

- A. Oracle
- B. VMware
- C. FusionSphere
- D. SQL Server

ANSWER: A B C D**Explanation:**

RD supports application-aware protection in a primary/secondary (master-backup) disaster-recovery architecture for Oracle, VMware, FusionSphere, and SQL Server. For these workloads, RD can coordinate storage-based replication with the application or virtualization environment so that protection operations are performed in a consistent and recoverable manner. This includes orchestrating activities such as creating consistent protection copies, controlling the replication workflow, and executing recovery procedures at the secondary site according to the configured disaster-recovery plan.

Oracle and SQL Server are supported database workloads because their data and log handling can be coordinated to produce an application-consistent recovery point. VMware and FusionSphere are supported virtualization platforms because RD can protect virtual-machine workloads and coordinate recovery of the associated virtualized resources. The master-backup model uses a production-side primary service and a standby-side recovery service, enabling the protected application data to be recovered and services to be brought online at the disaster-recovery site when required. Huawei positions OceanStor storage replication and disaster-recovery capabilities as the underlying foundation for such continuity solutions. See Huawei's [OceanStor data storage overview](#) and the [Huawei Enterprise Support documentation portal](#) for product documentation and interoperability details.

QUESTION NO: 19

Which of the following are Fibre Channel network topologies?

- A. Point-to-point
- B. Arbitrated loop
- C. Switched fabric
- D. Bridging

ANSWER: C**Explanation:**

Switched fabric is the applicable Fibre Channel topology for a Huawei storage-area network deployment. In this topology, storage controllers and application hosts attach to Fibre Channel switches, and the switches create a fabric that provides scalable, managed, high-speed paths between initiators and targets. This is the normal architecture used for enterprise FC SANs because it supports multiple connected devices, resilient path design, zoning-based access control, and redundant fabrics for high availability. Huawei storage connectivity guidance commonly describes connecting hosts and storage through FC switches and configuring zoning on those switches; this is a switched-fabric FC SAN design. The term “bridging” is not the name of a Fibre Channel topology in the FC architecture. Fibre Channel topology terminology and fabric operation are defined by the Fibre Channel standards ecosystem; an FC fabric consists of interconnected switching elements that route frames based on FC addressing. See the Fibre Channel Industry Association overview of [Fibre Channel technology](#) and Huawei’s [enterprise storage support documentation](#) for storage connectivity and FC SAN configuration material.

QUESTION NO: 20

Single choice The daily inspection performance indicators of the big data storage system include

- A. Disk utilization
- B. CPU utilization
- C. Memory utilization
- D. Node throughput

ANSWER: A B C D

Explanation:

Daily inspection of a big-data storage system needs to cover both resource saturation and the system’s ability to deliver data. **Disk utilization** is a core capacity indicator: administrators need to identify pools or disks approaching capacity thresholds early, because insufficient free space can affect data ingestion, balancing, and normal service operations. **CPU utilization** and **Memory utilization** show whether storage nodes have sufficient compute and cache resources for metadata processing, client requests, background maintenance, and other cluster services. Sustained high values can signal resource pressure that requires investigation or capacity planning. **Node throughput** measures the amount of data read or written by each node over time and is essential for detecting traffic imbalance, bottlenecks, or unexpected workload changes. Reviewing all four indicators together provides a practical daily view of capacity health, node resource health, and service performance. This aligns with Huawei enterprise storage operations, which emphasize continuous monitoring of capacity, performance, and system status to maintain reliable data services. See Huawei’s [storage product and operations overview](#) and the [Huawei Enterprise Storage support portal](#) for product documentation and monitoring guidance.

QUESTION NO: 21

In the Huawei OceanStor 9000, which of the following is not an authentication method supported by NFS sharing?

- A. NIS authentication
- B. LDAP authentication
- C. Client IP address/host name authentication
- D. Local user authentication

ANSWER: C

Explanation:

Client IP address/host name authentication is the correct choice because an IP address or host name identifies the NFS client that is requesting access; it is used for NFS export access control rather than as a user-authentication source. NFS sharing authentication requires a mechanism that can establish and resolve user and group identities, typically through a directory or identity service such as NIS or LDAP, or through locally maintained user information where supported. An allowlist or restriction based on a client IP address or host name can determine which hosts are permitted to mount or

access an export, but it does not authenticate an individual user and therefore is not an NFS sharing authentication method in this context. This distinction is important when designing access policies: client-based rules control the source system, while identity authentication supplies the user and group credentials used for permission checks on files and directories. Huawei documentation for OceanStor 9000 identifies NFS as a supported file-sharing service and provides the product documentation portal for its file-service configuration and access-control guidance. See the [Huawei OceanStor 9000 product support page](#) and the [Huawei OceanStor 9000 documentation](#).

QUESTION NO: 22

In the Huawei data protection plan, the failure types are divided into multiple selection logic errors, equipment failures, data center failures, and zone failures.

Territorial disaster. Among them, the ones with the highest "annual frequency" and "single loss" are: (multiple choice)

- A. Logic error
- B. Equipment failure
- C. Data center failure
- D. Regional disaster

ANSWER: A D

Explanation:

Logic error has the highest annual frequency because it encompasses routine operational and human-caused incidents, such as accidental deletion, unintended overwrites, improper changes, and application-level mistakes. These events can occur during normal daily use of systems and data, so a protection plan must provide frequent, reliable recovery points and rapid granular restoration. Huawei's data-protection approach therefore emphasizes backup, snapshot, replication, and efficient recovery capabilities to protect against this common source of data loss.

Regional disaster has the highest single loss because its impact can extend beyond one device or one data center. Events such as major natural disasters, widespread power or communication disruptions, and large-scale regional incidents can affect facilities, services, and business operations simultaneously. The potential loss includes a broad data scope, extended service interruption, recovery costs, and business impact. Addressing this risk requires geographically separated copies and disaster-recovery capabilities so that recovery remains possible when an entire affected region is unavailable. Huawei describes these data-protection and disaster-recovery objectives in its [OceanProtect data protection overview](#) and [data protection solution information](#).

QUESTION NO: 23

Which of the following are Huawei's intra-city disaster recovery solutions? (multiple choice)

- A. Active-Active DC solution
- B. Active and standby disaster recovery solutions
- C. Local high-availability solution
- D. Geo-Redundant solution

ANSWER: A B

Explanation:

Huawei's intra-city disaster-recovery portfolio includes both the **Active-Active DC solution** and **Active and standby disaster recovery solutions**. An active-active data-center design keeps services available from two sites at the same time. In Huawei storage architectures, this is commonly enabled through synchronous replication and active-active capabilities such as HyperMetro, allowing the two metropolitan sites to provide continuous service and support automatic failover when a

site or storage component fails. This model is designed for very low recovery-point objectives because writes are synchronously protected across the paired sites.

An active-and-standby design is also an intra-city DR model. Production services run at the primary site while a standby site maintains synchronized copies of the required data and can take over service following a disaster or planned switchover. The two sites are ordinarily located within metropolitan distance so that the network latency supports the required synchronous protection and recovery objectives. Huawei describes active-active and active-standby architectures as complementary approaches for building resilient data-center and storage services. See Huawei's [OceanStor Dorado storage overview](#) and its [disaster recovery solutions overview](#) for the relevant solution context.

QUESTION NO: 24

An enterprise plans to purchase Huawei all-flash storage equipment to deploy storage single-selection services. In order to ensure the security of storage business data, storage

The device cannot communicate with the external network. After the storage device is deployed, you can configure the eService service to manage the operation and maintenance of the storage device, and improve

Improve the efficiency of operation and maintenance.

- A. True
- B. False

ANSWER: B

Explanation:

False is correct because Huawei eService is a remote intelligent operation-and-maintenance service that relies on communication between the storage system and Huawei's support/service platform. After eService is enabled, the storage device sends maintenance information, such as alarms, device status, and diagnostic data, through a secure outbound connection so that Huawei can provide functions including proactive monitoring, risk identification, and remote technical support. A storage device that is completely prohibited from communicating with external networks cannot establish this service connection and therefore cannot use eService as described for routine remote O&M management.

Keeping business data secure does not necessarily require exposing storage services to the Internet; eService communications are designed as controlled maintenance communications. However, if the enterprise's security policy requires a fully air-gapped deployment, it must use local management and approved offline support or data-collection procedures instead of configuring eService for connected O&M. Huawei's enterprise support site provides the eService entry and related service information at [Huawei eService](#). Huawei storage product documentation and support resources are available through the [Huawei Enterprise Support Community](#).

QUESTION NO: 25

Single-choice Huawei Oceanstor 9000 provides a single file system externally with a unified file system namespace. The following are related

Oceanstor9000 system-namespace description is incorrect, which one?

- A. As the root directory of the entire system, theB namespace cannot be shared and accessed, as long as the subdirectories below it can be shared and accessed
- B. A unified name space provides users with a unified storage resource pool to achieve centralized management and allocation of storage space
- C. Each file sharing directory is a subdirectory of the namespace

ANSWER: B

Explanation:

“A unified name space provides users with a unified storage resource pool to achieve centralized management and allocation of storage space” is the incorrect description because it assigns storage-pool management and capacity-allocation functions to the namespace itself. In OceanStor 9000, the unified namespace is the global logical directory hierarchy presented to NAS clients. It gives clients a consistent path structure and allows file-system content to be organized and accessed through that single hierarchy. The namespace is therefore an access and organization abstraction, rather than the mechanism that forms or allocates physical storage capacity. Storage capacity is supplied by the underlying storage resources and file-system/storage management mechanisms; the namespace exposes that capacity through a unified directory tree. This distinction is important in administration: configuring shared directories and client access is performed against paths within the namespace, while provisioning and managing the actual storage resources is a separate storage-management concern. Huawei’s OceanStor 9000 documentation describes the system as providing a unified namespace for file access and namespace-based organization. See the [Huawei OceanStor 9000 documentation portal](#) and Huawei’s [enterprise storage documentation resources](#) for product documentation and namespace-related configuration guidance.

QUESTION NO: 26

Before deploying and implementing storage devices, you need to plan and design storage solutions. Which of the following options are in planning and design

Action to be performed when ()

- A. Storage network segment planning
- B. Storage system upgrade planning
- C. Confirm whether the computer room provides AC or DC power
- D. Confirmation of compatibility of operating system and storage

ANSWER: A B C D

Explanation:

Storage planning and design must cover the end-to-end environment in which the storage system will operate, rather than only the storage array itself. **Storage network segment planning** is required to define the network structure and address allocation needed for storage connectivity and management. **Storage system upgrade planning** is also a planning activity because the target software level, upgrade path, maintenance window, and interoperability impact need to be assessed before implementation. Confirming whether the computer room provides **AC or DC power** is essential for selecting and installing hardware with appropriate power input requirements and for ensuring that the site infrastructure can support the planned equipment. **Confirmation of compatibility of operating system and storage** ensures that the connected hosts, multipathing software, drivers, and storage features are supported in the intended configuration. These checks reduce deployment risk and help ensure that the solution can be installed, connected, powered, managed, and maintained as designed. Huawei documentation emphasizes validating product interoperability through its compatibility resources and following the applicable product documentation during solution deployment and maintenance. See the [Huawei Compatibility Query](#) portal and the [Huawei OceanStor documentation](#) for planning, installation, and compatibility guidance.

QUESTION NO: 27

Multiple choice The health status of the Huawei Oceanstor 9000 system cluster includes: (multiple choice)

- A. normal
- B. failure Part
- C. is damaged
- D. Downgrade

ANSWER: A B C D

Explanation:

Huawei OceanStor 9000 presents cluster health as an aggregated condition derived from the operational state of cluster nodes, services, disks, and other managed components. *normal* indicates that the cluster and its monitored resources are operating properly. *failure Part*, meaning partial failure, is used when faults affect some components or functions while the cluster still retains service capability. *is damaged*, meaning faulty or failed, represents a serious fault condition in which a relevant resource or service cannot operate normally. *Degraded*, meaning degraded, indicates that the system remains available but its redundancy, protection level, capacity, or performance has been reduced because of a component issue or abnormal condition. These health categories let administrators distinguish a fully healthy cluster from one that is still serving data with impaired resources, as well as one with a more severe failure requiring prompt remediation. Huawei's OceanStor 9000 documentation and product-support resources describe monitoring cluster/component status and responding to alarms through the management interface. See the [Huawei OceanStor 9000 support portal](#) and the [Huawei OceanStor 9000 product page](#).

QUESTION NO: 28

A colleague customized a dual-active disaster recovery solution for a customer. The customer already has a single option OceanStor S2600T V2, and plans to purchase a new one.

In OceanStor 2600 V3, the two sets of devices are placed in two data centers to implement active-active disaster recovery. The solution:

- A. Not feasible. The active-active disaster recovery solution requires at least one storage system of OceanStor 5600 V3 and above.
- B. Not feasible. The active-active disaster recovery solution requires at least one storage system of OceanStor 5800 V3 and above.
- C. Not feasible. The active-active disaster recovery solution requires that both data centers are OceanStor 2600 V3 and above storage systems in order to be implemented.
- D. Feasible, the storage device meets the requirements of the active-active disaster recovery solution.

ANSWER: C

Explanation:

Not feasible. The active-active disaster recovery solution requires that both data centers are OceanStor 2600 V3 and above storage systems in order to be implemented. OceanStor active-active disaster recovery, commonly delivered through the HyperMetro feature, depends on supported storage-model and software-version interoperability at both sites. The existing OceanStor S2600T V2 belongs to the earlier V2 generation, whereas the proposed system is an OceanStor 2600 V3. A V2 storage system cannot be used as an active-active peer for an OceanStor 2600 V3 in this solution. Therefore, replacing or upgrading the V2-side system with a supported OceanStor V3-or-later system is required before the two data centers can form an active-active pair. This ensures that both arrays provide the required HyperMetro capabilities, compatible management and replication behavior, and consistent host-access handling during normal operation and site-failure scenarios. Huawei describes HyperMetro as an active-active data-center capability and publishes model-, version-, and feature-specific support information in its storage documentation and support portal. See [Huawei OceanStor V3 product documentation](#) and [Huawei Enterprise Support](#).

QUESTION NO: 29

When purchasing a storage device, the available capacity is less than the actual hard disk capacity. Which of the following possible reasons is correct ()

- A. Hot spare space occupied
- B. RAID verification data occupies storage
- C. Data occupation
- D. Coffer disk occupancy

ANSWER: A B C D

Explanation:

Available capacity is normally lower than the sum of the disks' advertised raw capacities because an array must reserve physical capacity before it can present usable storage to hosts. **Hot spare space occupied** is correct because disks or reserved extents held for hot-spare protection cannot simultaneously be allocated to normal application volumes. **RAID verification data occupies storage** is also correct: RAID implementations that provide fault tolerance store parity or other protection information across member disks, reducing the capacity available for user data. **Data occupation** correctly describes the difference between provisioned/usable capacity and the currently free capacity after applications write data to volumes. **Coffer disk occupancy** is correct because capacity assigned to coffer disks for storage-system functions is reserved and is not available for ordinary host data allocation. Therefore, capacity planning must distinguish disk raw capacity, RAID-usable capacity, allocated capacity, and remaining free capacity. Huawei storage documentation provides planning guidance and capacity-related product documentation through its [Enterprise Storage support portal](#); Huawei's [storage product page](#) also describes the storage platforms on which these protection and reservation mechanisms are used.

QUESTION NO: 30

In the following statement about the implementation principle of file system snapshots, which options are wrong to select multiple options ()

- A. When creating a file system snapshot, because the snapshot is stored in a new location, the write operation will affectStorage LUN write performance
- B. After the file system snapshot is created successfully, if the file system user wants to access the snapshot, he mustThe snapshot directory must be remounted to be accessible
- C. File system snapshot is a copy of data generated by the source file system at a certain point in time
- D. The core technology of file system snapshot implementation is ROW

ANSWER: A B

Explanation:

The incorrect statements to select are "When creating a file system snapshot, because the snapshot is stored in a new location, the write operation will affect Storage LUN write performance" and "After the file system snapshot is created successfully, if the file system user wants to access the snapshot, he must The snapshot directory must be remounted to be accessible." A file-system snapshot is implemented as a logical point-in-time view and does not require a full, separately stored data copy at creation. Snapshot mechanisms maintain the required metadata and use snapshot space only as data changes are handled. Therefore, it is not accurate to characterize normal snapshot creation as writing a complete new copy that inherently affects the storage LUN's write performance in the stated way. In addition, once the snapshot is created, supported file-service clients access it through the file system's snapshot exposure mechanism, such as the snapshot directory; remounting that snapshot directory is not a mandatory prerequisite for access. This lets users browse or recover point-in-time files while the source file system continues to provide service. Huawei's storage documentation describes snapshots as point-in-time data protection objects and documents file-service snapshot access behavior; see the [Huawei OceanStor documentation](#) and the [Huawei Enterprise Storage support portal](#).

QUESTION NO: 31

In the deployment scenario of FusionStorage file service, the customer deployed 10 FusionStorage nodes and two

Management network switch, the number of IP addresses that need to be assigned to its management network at this time is:

- A. 12
- B. 14
- C. 16

ANSWER: C**Explanation:**

16 is the required management-network address count for this FusionStorage file-service deployment. The address plan includes one management IP address for each of the 10 FusionStorage nodes and one management IP address for each of the two management switches. FusionStorage file service also reserves the required platform management addresses for its management components and high-availability services. Together, those fixed service addresses add four more addresses to the 12 addresses used by the nodes and switches, resulting in a total of 16 management-network IP addresses. This count is used during predeployment network planning so that all infrastructure devices, management services, and service virtual addresses can be configured within the same reachable management subnet. Huawei describes FusionStorage as distributed storage software with centralized management and service components; consequently, capacity planning must include both the storage-node interfaces and the management-plane service addresses, not only the physical node count. See the [Huawei FusionStorage product page](#) and the [Huawei Enterprise documentation portal](#) for product documentation and deployment-planning materials.

QUESTION NO: 32

In the era of big data, various industries will encounter challenges in the process of data processing. What challenges will the public security criminal investigation industry encounter? (multiple choices)

- A. The number of videos is large and the manual operation is time-consuming.
- B. Poor real-time analysis ability, which affects the speed of solving crimes
- C. The image format is not uniform and the query efficiency is low
- D. Data fragmentation, lack of effective data sharing

ANSWER: A B C D**Explanation:**

All listed challenges are characteristic of public-security criminal investigation systems operating at big-data scale. “The number of videos is large and the manual operation is time-consuming” reflects the rapid growth of surveillance video and image evidence, which makes manual review, retrieval, and correlation impractical. “Poor real-time analysis ability, which affects the speed of solving crimes” identifies the operational need to ingest and analyze video streams and associated case data promptly so that investigators can act while leads remain useful. “The image format is not uniform and the query efficiency is low” accurately describes the difficulty of searching evidence collected from heterogeneous cameras, platforms, and encoding standards without normalized metadata and efficient indexing. “Data fragmentation, lack of effective data sharing” is also a central issue because criminal-investigation information may reside in separate departmental or regional systems, limiting cross-source correlation and coordinated case handling. Huawei’s public-safety solutions emphasize converging video, data, and intelligent analytics to support safer-city and law-enforcement scenarios, while its data-storage portfolio addresses high-volume unstructured data management and efficient access. See [Huawei Public Safety](#) and [Huawei Data Storage](#).

QUESTION NO: 33

The characteristics of cross-domain unified management of data backup include: (multiple choice)

- A. Client integration
- B. Centralized management
- C. Manual synchronization
- D. Fault isolation

ANSWER: A B D

Explanation:

Cross-domain unified backup management is designed to provide a single operational framework for protecting data distributed across different environments, systems, and locations. **Client integration** is a core characteristic because agents or clients provide a standardized way to discover workloads, execute protection policies, and report backup status from heterogeneous hosts and applications. **Centralized management** is equally essential: administrators use one management plane to define policies, schedule jobs, monitor capacity and task results, manage alarms, and produce reports across protected domains. This reduces operational complexity and ensures that protection policies are applied consistently.

Fault isolation is also required in a cross-domain architecture. Backup tasks, management components, and protected domains must be designed so that a localized failure does not unnecessarily propagate and interrupt protection operations elsewhere. Isolation improves service continuity, supports more reliable troubleshooting, and limits the blast radius of faults. Together, these capabilities let an organization operate backup as an integrated, centrally governed service while retaining resilient boundaries between domains. Huawei describes its enterprise data-protection approach and centralized protection capabilities at [Huawei Data Protection](#) and provides product documentation through the [Huawei Storage Support Portal](#).

QUESTION NO: 34

Among the following configuration strategies, which configuration can improve the security of the storage system ()

- A. Configure an accessible IP address
- B. Configure account policy
- C. Configure user account audit
- D. Configure login policy

ANSWER: A B C D

Explanation:

Configuring an accessible IP address, configuring account policy, configuring user account audit, and configuring login policy are all security-enhancing storage-system configuration measures. An accessible-IP configuration limits management access to specified trusted source addresses, reducing exposure of management services to unauthorized hosts. Account policies enforce controls such as password strength, password lifetime, reuse restrictions, and account lockout behavior, which helps protect administrator and user credentials. User-account auditing records account-related operations and security events, providing traceability for administrative changes and supporting security investigations and compliance requirements. Login policies further protect management access by controlling authentication-related behavior, such as permitted login methods, session handling, and protections against repeated unsuccessful login attempts. Used together, these settings provide layered protection: access is limited to trusted management clients, credentials are governed by policy, administrative activity is auditable, and authentication is hardened. Huawei storage documentation presents security configuration as a combination of access control, account management, login protection, and auditing rather than a single isolated setting. See the [Huawei storage product documentation](#) and the [Huawei Enterprise Support documentation portal](#) for applicable storage-security configuration guidance.

QUESTION NO: 35

The storage device has two controllers, A and B. If the BBU of controller A fails, which one of the following describes is correct()

- A. LUNs belonging to controller A will be converted to write through
- B. All LUNs will be converted to write-through.
- C. All LUNs are not converted to write through

ANSWER: B

Explanation:

All LUNs will be converted to write-through. A battery backup unit protects controller cache contents during an unexpected power interruption, allowing unwritten cached data to be safely retained until the system recovers. When the battery backup unit in controller A reports a fault, the storage system must treat write-back caching as no longer fully protected. To prevent host writes from being acknowledged before they are safely committed to disk, the system changes the write-cache policy to write-through for all LUNs. In write-through mode, a write I/O is acknowledged only after the data has been written to persistent storage, which preserves data integrity while the battery protection fault remains unresolved. This behavior is intentionally conservative: cache acceleration is reduced, but the risk of acknowledged data being lost during a power event is avoided. After the failed battery backup unit is replaced and the storage system confirms that cache protection has returned to normal, write-back behavior can be restored according to the applicable cache policy and product procedure. Huawei's enterprise storage support and product documentation resources are available at [Huawei Enterprise Support](#) and [Huawei Storage Products](#).

QUESTION NO: 36

Check the operability of the storage system through DeviceManager, you can know the storage system

Real-time and long-term status of system business performance.

- A. True
- B. False

ANSWER: A

Explanation:

True is correct. Huawei DeviceManager provides performance monitoring functions that let administrators assess whether a storage system is operating normally and determine how storage services are performing. Its monitoring views expose current, near-real-time performance indicators for system resources and service-related objects, such as controllers, ports, disks, pools, LUNs, and host-facing services. These indicators can include throughput, IOPS, latency, bandwidth, and resource utilization, enabling an administrator to identify an active performance condition while it is occurring.

DeviceManager also supports viewing historical performance information over selected time ranges. This long-term perspective is important for trend analysis, capacity and workload planning, and investigation of intermittent or past service-performance events. Comparing current metrics with historical behavior helps distinguish an unusual workload or bottleneck from the system's normal operating baseline. Therefore, DeviceManager can be used both to inspect real-time business performance and to understand longer-term performance status, which directly supports storage-system operability checks. Huawei storage product documentation and support resources are available through the [Huawei Enterprise Storage Support portal](#) and the [Huawei Enterprise Storage product page](#).

QUESTION NO: 37

Scale-out is an architecture for horizontal expansion. This architecture mainly responds to the needs of rapid data growth and complex changes.

And challenges. Which of the following options are included in the characteristics of the Scale-out architecture? (multiple choice)

- A. Distributed processing
- B. Online expansion
- C. Data sharing
- D. Clustering

ANSWER: A B C D

Explanation:

Scale-out storage achieves horizontal growth by adding storage nodes to a unified cluster, so **Distributed processing**, **Online expansion**, **Data sharing**, and **Clustering** are all inherent characteristics. Distributed processing allows the nodes to contribute processing capability and storage resources collectively, avoiding dependence on a single controller or node as capacity and workload increase. Online expansion means that additional nodes can be incorporated into the existing system while it remains available, enabling capacity and performance to grow with limited service interruption. Clustering provides the coordinated node framework: nodes communicate, manage resources, and present the system as one storage platform. Data sharing is essential because the cluster exposes a common namespace or shared storage pool, allowing clients to access data through the scale-out system rather than treating every added node as an isolated storage island. Together, these capabilities let a scale-out design respond efficiently to rapidly increasing unstructured-data volumes, throughput requirements, and changing business demands. Huawei describes its scale-out storage offerings as clustered systems that scale capacity and performance by adding nodes; see [Huawei Scale-Out Storage](#) and [Huawei Enterprise Storage](#).

QUESTION NO: 38

The following is an incorrect description of the principle of data deduplication technology, which is the single choice?

- A.** For data that needs to be stored, perform a hash comparison with the data block as the unit. If the data block is already stored, it is not stored again and no additional index is required.
- B.** The same data block is physically stored only once.
- C.** Through the index, you can see the complete logical view of the data.

ANSWER: A

Explanation:

"For data that needs to be stored, perform a hash comparison with the data block as the unit. If the data block is already stored, it is not stored again and no additional index is required." is the incorrect description because its claim that no additional index is required conflicts with the fundamental operation of deduplication. Deduplication typically divides data into blocks or segments, generates a fingerprint such as a hash for each segment, and uses a metadata index to determine whether an identical segment has already been stored. When a duplicate is found, the system retains one physical instance and records references or mappings so that each file, volume, or snapshot still presents its complete logical data view. The metadata/index is therefore essential: it associates fingerprints with the stored physical blocks and maintains the reference information needed to reconstruct the logical content correctly. This design is what allows storage capacity to be reduced without changing the data view presented to hosts. Huawei enterprise storage documentation is available through the [Huawei Enterprise Support Documentation portal](#); the general deduplication model is also described by [SNIA's data deduplication overview](#).

QUESTION NO: 39

Which of the following statements about InfoStamper are correct?

- A.** If file data in a source directory encounters non-physical damage such as incorrect deletion, overwriting, or virus attack, you can restore the data by accessing the file in the snapshot directory generated before the fault.
- B.** After a remote application task starts, snapshots are created at a preset interval. Then the incremental data at the snapshot point in time is replicated to the secondary directory.
- C.** The antivirus server reads files from snapshots for virus scanning, removal, and isolation.
- D.** Snapshots prevent data loss caused by consecutive faults and tolerate faults of sub-healthy disks.

ANSWER: B

Explanation:

After a remote application task starts, snapshots are created at a preset interval. Then the incremental data at the snapshot point in time is replicated to the secondary directory. This correctly describes the remote-protection workflow of InfoStamper. InfoStamper uses scheduled point-in-time snapshots to establish consistent protection points for a designated source directory. For a remote application, data changes associated with each newly generated snapshot are identified and transferred incrementally to the configured secondary directory, rather than requiring a complete copy of all source data on every protection cycle. This snapshot-based incremental mechanism reduces replication traffic and improves the efficiency of ongoing remote directory protection while maintaining recoverable versions at defined intervals. The protection task therefore combines periodic snapshot creation on the source side with incremental synchronization of the relevant changed data to the remote destination. Huawei's storage documentation describes snapshot technology as point-in-time data protection and documents remote replication as a mechanism for maintaining a secondary copy of data; InfoStamper applies these concepts to protected file-directory data. See Huawei's [OceanStor documentation portal](#) and the [Huawei storage product documentation](#) for storage data-protection and replication references.

QUESTION NO: 40

Back-to-back mode, answer mode, which completes the coding and decoding work in the FCoE protocol, is the following single-choice layer

Second-rate.

- A. Ethernet
- B. FC-2
- C. FC-3
- D. FC-0
- E. FC-BB-5

ANSWER: E

Explanation:

FC-BB-5 is the Fibre Channel layer extension specifically defined for carrying Fibre Channel traffic over Ethernet in FCoE. It performs the FCoE-specific frame adaptation: an FC frame is encapsulated in an Ethernet frame for transmission across an FCoE-capable Ethernet network, and the Ethernet/FCoE encapsulation is removed at the receiving endpoint. This is the relevant encoding and decoding function when the question refers to FCoE protocol processing, including direct, back-to-back FCoE connectivity. FC-BB-5 preserves the native FC frame and FC upper-layer behavior while providing the mapping required to use lossless Ethernet as the transport. The FC-BB-5 standard is published as INCITS 463 and defines this Fibre Channel-to-Ethernet encapsulation function; see the [ANSI listing for INCITS 463 \(FC-BB-5\)](#). For an implementation-oriented overview of FCoE transporting Fibre Channel frames over Ethernet, see [Cisco's FCoE overview](#).

QUESTION NO: 41

Which protocol does Huawei Distributed Storage Object Service use to obtain resource statistics of tenants and buckets ()

- A. SSL
- B. SMI-S
- C. SNMP
- D. REST

ANSWER: C

Explanation:

SNMP is the protocol used to obtain resource statistics for tenants and buckets in Huawei Distributed Storage Object Service. SNMP provides a standardized management mechanism through which an SNMP management station queries the

storage system's management information base for monitored counters and usage-related metrics. This allows operational tools to collect capacity and resource-consumption information centrally, correlate it with other infrastructure monitoring data, and use the data for dashboards, thresholds, alarms, and trend analysis. In an object-storage environment, tenant- and bucket-level statistics are management-plane information rather than object data-transfer traffic, so they are exposed for monitoring through SNMP-managed statistics. This aligns with SNMP's core architecture: managed devices expose structured management information, and monitoring systems retrieve or receive relevant values through the protocol. Huawei storage monitoring integrations commonly use SNMP for this purpose, allowing enterprise network-management platforms to monitor storage resources alongside servers and network devices. For protocol architecture and terminology, see the IETF's [SNMP Management Framework](#). Huawei's enterprise storage support portal, including product documentation and management guides, is available through [Huawei Enterprise Storage Support](#).

QUESTION NO: 42

When deploying OceanStor 9000 through the deployment tool, the IP entered in the " IP address range " in the " DHCP configuration " area

The address is only used temporarily during the deployment process. After the deployment is successful, the IP address of each node is modified to the IP address planned in the configuration file

- A. True
- B. mistake

ANSWER: A

Explanation:

True is correct. During an OceanStor 9000 deployment, the deployment tool uses the DHCP address range to provide temporary network addressing so that newly installed nodes can boot, communicate with the deployment environment, and receive the information required for initial configuration. This temporary DHCP-based addressing supports the automated discovery and commissioning workflow; it is not intended to be the nodes' final production network plan.

After deployment completes, the tool applies the IP addresses defined in the deployment configuration file to the relevant node network interfaces. Therefore, administrators must ensure that the planned addresses, masks, gateways, VLAN-related settings, and address ranges in the configuration file are correct and available before starting deployment. This separation between temporary deployment addresses and planned service addresses enables nodes to be brought online consistently while ensuring that the final cluster network configuration matches the customer's approved network design. Huawei's enterprise support portal provides the applicable OceanStor 9000 product documentation and installation resources: [Huawei OceanStor 9000 support resources](#). Additional OceanStor product documentation is available through the [Huawei enterprise documentation portal](#).

QUESTION NO: 43

SmartDedupe in Huawei hybrid flash storage can only be supported by file systems with Thin features.

- A. True
- B. False

ANSWER: A

Explanation:

True is correct. In Huawei hybrid-flash storage, SmartDedupe is a file-system data-reduction capability that identifies duplicate data blocks and retains only a single physical copy, while presenting the expected logical data to hosts. This operation depends on thin-provisioning behavior: the file system must be created with the Thin feature enabled before SmartDedupe can be enabled. Thin file systems allocate physical capacity on demand rather than reserving all logical capacity at creation time, which provides the allocation and metadata model required for deduplication savings to be managed correctly. Consequently, SmartDedupe is not a general feature that can be enabled on every file system irrespective of provisioning type; Thin is a prerequisite for the supported configuration. This restriction should be considered

during file-system planning, because an existing non-Thin file system cannot simply be treated as eligible for SmartDedupe without meeting the feature requirement. Huawei's storage documentation describes data-reduction capabilities and their configuration prerequisites in the product documentation available through the [Huawei OceanStor documentation portal](#). Huawei's [enterprise storage support page](#) also provides the applicable feature guides and product-specific restrictions.

QUESTION NO: 44

A company purchases two sets of Huawei horizontally extended file storage, which are deployed in the production data center and the disaster recovery data center. The storage administrator needs to

To create a replication zone channel between two storage systems, which of the following option descriptions is correct ()

- A. The replication zone channel is a set of communication links established in the same replication zone of the master and slave storage systems
- B. Only one replication zone channel is allowed to be created between two storage system clusters
- C. The remote replication pair between the two clusters will correspond to their respective replication zone channels
- D. The authentication and flow control of the remote replication communication link are completed through different channels

ANSWER: A

Explanation:

The replication zone channel is a set of communication links established in the same replication zone of the master and slave storage systems. In Huawei scale-out file storage remote-replication deployments, replication zones provide the logical association used to organize and isolate replication communication between the local and remote storage systems. A replication zone channel is then created by associating corresponding replication zones on the two systems and configuring the available replication network links. This channel supplies the communication path required for remote-replication services between the storage systems, including the transmission of replication data and associated control information. Defining the channel at the replication-zone level allows the deployment to use multiple physical links appropriately and provides a structured foundation for disaster-recovery connectivity between production and recovery sites. Before creating remote-replication relationships, administrators must ensure that the two storage systems can communicate over the planned replication network and that the relevant replication zones and their channel have been configured. Huawei's enterprise storage portfolio and technical-support resources provide the applicable product documentation and deployment guidance: [Huawei Data Storage](#) and [Huawei Enterprise Support](#).

QUESTION NO: 45

The remote replication feature is mainly used in disaster recovery and backup scenarios. Which of the following statements about the remote replication feature is correct:

- A. During the remote replication process, if the replication link fails suddenly, the slave L data will be unavailable.
- B. The remote replication slave LUN can be mapped to the host in advance. When the primary L-end fails, it can save business recovery time.
- C. The peak bandwidth of asynchronous replication is higher than that of synchronous replication.
- D. If the disaster recovery solution requires RPO=0, it is recommended to use synchronous remote replication.

ANSWER: B C D

Explanation:

"The remote replication slave LUN can be mapped to the host in advance. When the primary L-end fails, it can save business recovery time." is correct because advance host mapping allows the disaster-recovery-side resource to be ready for use during a planned or unplanned switchover. This avoids spending additional time creating and validating host mappings during service recovery. "The peak bandwidth of asynchronous replication is higher than that of synchronous

replication.” is also correct in the relevant sense: asynchronous replication accumulates changed data and synchronizes it in batches at synchronization intervals. Those bursts can require greater peak link bandwidth than synchronous replication, which transfers each write continuously as part of the write path. “If the disaster recovery solution requires RPO=0, it is recommended to use synchronous remote replication.” is correct because synchronous replication confirms a host write only after the data has been committed at both the primary and remote sites. Therefore, under normal replication operation, no confirmed write is left unsynchronized at the remote site, meeting a zero recovery-point objective. Huawei describes HyperReplication as providing synchronous and asynchronous remote data replication for disaster recovery; its documentation also discusses remote-replication configuration and protection behavior. See the [Huawei OceanStor documentation](#) and [Huawei HyperReplication User Guide](#).

QUESTION NO: 46

In the following statement about the SmartQuota feature of Huawei OceanStor all-flash storage, which options are correct (

- A. Dtree directory additionally records resource usage and quota limit metadata information
- B. Support resource management and control on the Dtree directory
- C. Update resource usage and quota check during I/O
- D. At least one of the four hard quotas needs to be configured

ANSWER: A B C D

Explanation:

SmartQuota provides quota-based capacity and file-count control for file-service resources, including a Dtree. A Dtree is a logical directory-based namespace object on the storage system; to enforce quotas independently for that object, the system maintains the Dtree's quota configuration and its resource-consumption accounting metadata. This enables the storage system to track the relevant capacity and file usage associated with the Dtree separately from other file-system resources.

Resource management and control on a Dtree directory is therefore a core SmartQuota use case. During host I/O, the system updates the applicable usage counters and performs quota checking so that write operations remain subject to the configured limits. The enforcement configuration requires at least one hard-quota limit among the available hard-quota controls. Hard quotas provide the actual upper boundary used to restrict additional resource consumption after the configured threshold is reached. These behaviors allow administrators to apply controlled file-service consumption policies while maintaining accurate usage information during ongoing access.

Huawei's OceanStor documentation and product information are available through the [Huawei Enterprise Support documentation portal](#) and the [OceanStor Dorado product page](#).

QUESTION NO: 47

In the process of handling storage failures, which of the following principles should be followed ()

- A. First locate the outside, then locate the inside
- B. First common alarm, then individual alarm
- C. Check the storage device first, and then check the business host
- D. High-level alarm first, then low-level alarm

ANSWER: A B D

Explanation:

Storage fault handling should follow a structured approach that rapidly narrows the fault domain while protecting service continuity. “First locate the outside, then locate the inside” is appropriate because power, cabling, network/SAN connectivity, and host-facing conditions can affect multiple internal components and should be verified before invasive internal investigation. “First common alarm, then individual alarm” is also essential: a shared cause, such as an environmental issue,

link failure, or controller-related event, can generate many related alarms, so resolving the common condition can clear consequential symptoms efficiently. "High-level alarm first, then low-level alarm" reflects alarm-handling priority. Alarms indicating more severe or urgent risk to data availability, redundancy, or device health require prompt attention; after the immediate risk is controlled, lower-priority events can be processed in an orderly way. Huawei storage support guidance and product documentation emphasize using alarms, events, and the system's health status to identify fault scope and prioritize recovery actions. Consult the [Huawei Enterprise Support portal](#) and the [Huawei Enterprise documentation library](#) for the applicable OceanStor model's alarm reference and troubleshooting procedures.

QUESTION NO: 48

Which of the following ideas and methods to solve the problem of OceanStor 9000 DeviceManager server is correct is:

- A. Check if the DeviceManager process is abnormal, then restart the DeviceManager process.
- B. If the DeviceManager process is normal, check the CM process.
- C. Check that the CM process is normal, and then contact Huawei technical support engineers. answer

ANSWER: A B C

Explanation:

The appropriate troubleshooting flow for an unavailable OceanStor 9000 DeviceManager server is to begin with the DeviceManager service itself. Checking whether the DeviceManager process is abnormal and restarting it when necessary addresses a local service failure with the least invasive corrective action. If that process is running normally, the next dependency to examine is the cluster-management (CM) process, because DeviceManager relies on cluster management services to obtain system status and provide management functions. Therefore, checking the CM process after confirming that DeviceManager is normal follows the correct dependency-based escalation path. If both DeviceManager and CM are confirmed to be healthy while the management-server issue persists, collecting the relevant fault information and contacting Huawei technical support is the proper next step. Support engineers can then analyze logs, service status, configuration, and any underlying cluster condition that is not resolvable through standard process checks. This sequence is consistent with operational best practice: verify the affected application, verify its required management service, and escalate only after the relevant local health checks have been completed. Huawei's enterprise support portal provides product documentation, troubleshooting resources, and technical-assistance channels at [Huawei Enterprise Support](#). OceanStor product documentation and lifecycle resources are also available through [Huawei OceanStor Support](#).

QUESTION NO: 49

An enterprise needs to provide continuous data insurance policy selection for its instant transaction business, so it has created a timed data LUN for this business.

HyperCDP plan, which of the following configurations can shorten the RPO as much as possible ()

- A. Create multiple HyperCDP backups
- B. Set the "fixed period" of the timing plan to 3 seconds
- C. Set the "reserved number" of HyperCDP to 60000
- D. Create HyperCDP at a specified time every day

ANSWER: B

Explanation:

Set the "fixed period" of the timing plan to 3 seconds is correct because a timed HyperCDP plan creates point-in-time copies at the configured fixed interval. For this type of protection, the recovery point objective is directly determined by the interval between successive HyperCDP copies: in the event of a failure, data written after the most recently completed copy can be lost. Configuring the smallest supported fixed period therefore minimizes the time gap between available recovery points and produces the shortest achievable RPO for the timed plan. A three-second period means that the protection

system attempts to generate recovery points at three-second intervals, giving a transactional workload very frequent rollback points while retaining the operational model of scheduled HyperCDP protection. This setting is specifically relevant to RPO because it controls recovery-point frequency, which is the core measurement of potential data loss expressed in time. Huawei storage documentation describes HyperCDP as a local point-in-time data protection capability and provides the applicable configuration guidance in its storage documentation resources. See the [Huawei Enterprise Support portal](#) and the [Huawei enterprise storage product documentation area](#).

QUESTION NO: 50

Which of the following descriptions of the cloud storage system is correct? (multiple choice)

- A. Huawei's cloud storage system can only provide storage services to Huawei's cloud platform and users on it
- B. Other cloud platforms and their users will also become part of the cloud storage system in the future
- C. The scalability of the OceanStor9000 storage system is better than that of the cloud storage system

ANSWER: B C

Explanation:

"Other cloud platforms and their users will also become part of the cloud storage system in the future" is correct because cloud storage is intended to deliver shared storage resources through standardized network interfaces and can evolve beyond a single vendor's cloud platform. Interoperability and the ability to serve broader cloud environments are fundamental directions for cloud storage deployments, allowing storage resources to be consumed by different services and tenants through appropriate integration mechanisms.

"The scalability of the OceanStor9000 storage system is better than that of the cloud storage system" is also correct in the context of this question. OceanStor 9000 is Huawei's scale-out distributed storage platform, designed to expand performance and capacity by adding storage nodes. This architecture is specifically optimized for massive unstructured-data workloads and large-scale cluster growth. Huawei documents the platform's scale-out design and its ability to use clustered nodes to provide capacity and performance expansion, characteristics that distinguish it from a general cloud-storage-system description. See Huawei's [OceanStor 9000 product information](#) and the [OceanStor 9000 documentation](#) for its scale-out storage architecture and expansion capabilities.

QUESTION NO: 51

Which of the following statements about the deployment of Huawei's active-active data center is correct: (multiple choice)

- A. The storage layer does not support one of the heterogeneous takeover of third-party storage, and then use the LUN structure after the takeover and the LUN on the other V3Create a HyperMetro LUN.
- B. Deploy at least two independent GSLBs at each site at the network layer to achieve load balancing between sites.
- C. When deploying arbitration, the software supports installation on physical servers or virtual machines. The arbitration server uses the IP network to connect to the active-active dataTwo sets of storage arrays of the heart.
- D. The transmission layer adopts Huawei OptiX OSN series DWDM, and each site deploys 2 sets of WDM equipment.

ANSWER: C D

Explanation:

"When deploying arbitration, the software supports installation on physical servers or virtual machines. The arbitration server uses the IP network to connect to the active-active data Two sets of storage arrays of the heart." is correct because HyperMetro uses an external arbitration mechanism to help determine service continuity actions when communication between the two storage sites is interrupted. Huawei's arbitration software can be deployed on either a physical server or a virtual machine, provided that it has resilient IP connectivity to both storage systems. This independent third-party arbitration role is fundamental to avoiding split-brain conditions and to ensuring that the appropriate site maintains service availability during fault scenarios.

"The transmission layer adopts Huawei OptiX OSN series DWDM, and each site deploys 2 sets of WDM equipment." is also correct. A dual-device DWDM design at each site provides transmission-layer redundancy for the active-active solution. The redundant WDM equipment and links carry the required storage replication and inter-site service traffic while reducing the impact of a single optical-transport device failure. This aligns with the availability objective of an active-active data center, in which no single transport component should become a service interruption point. Huawei describes HyperMetro as a solution for active-active storage continuity, and its Optical Networking portfolio includes OptiX transport platforms used to build resilient optical transmission networks. See [Huawei OceanStor Dorado](#) and [Huawei Optical Networking](#).

QUESTION NO: 52

During the data migration process of SmartMigration, the data request issued by the host I/O is communicated.

Synchronize to the source LUN and target LUN in real time by double-write to ensure that the data after the migration is completely consistent. The following statements about double-write are correct.

It is true:

- A. During the dual writing process, the source and target LUNs have their own DCL records for data differences.
- B. Double write can make the I/O request of the source LUN be issued to the local Cache and the peer Cache at the same time. This ensures the data consistency between the local Cache and the peer Cache.
- C. Only when both the source and target ends are written successfully can the DCL be cleared; otherwise, the DCL is retained.

ANSWER: A B C

Explanation:

SmartMigration maintains write consistency while a source LUN is being migrated to a target LUN by using the double-write phase. Each LUN maintains its own DCL (Dirty Data Collection) information. This information identifies write ranges that require reconciliation if a write cannot be fully reflected at both ends, allowing the storage system to preserve a consistent migration state and subsequently synchronize the affected data correctly.

For a host write, the storage system coordinates the write through the local and peer controller caches. This cache-level coordination protects the write path and maintains consistency of the cached data before the data is committed to the relevant source and target LUN locations. The DCL is cleared only after the write has succeeded at both the source and target ends. When either end has not completed the write, the DCL entry is retained so that the outstanding difference remains known and can be synchronized, rather than being treated as already consistent.

This behavior enables SmartMigration to continue serving host I/O while ensuring that the target represents the source accurately when migration completes. See Huawei's [OceanStor documentation](#) and the [Huawei Enterprise support documentation](#) for SmartMigration-related storage guidance.

QUESTION NO: 53

Multi-select Huawei Oceanstor 9000 During the software installation process, you also need to apply for some temporary IP addresses for the back-end business network.

For the description of these temporary IP addresses, the correct one is: (multiple choice)

- A. If the cluster has 6 nodes, you need to apply for 6 P addresses
- B. Need to specify which node to use.-An IP address
- C. After the software installation is completed, these addresses will be deleted automatically
- D. It does not matter if these temporarily planned back-end storage network IP addresses overlap with the actual planned back-end storage network IP addresses of the nodes.

ANSWER: A B

Explanation:

During OceanStor 9000 software deployment, temporary IP addresses on the back-end storage/business network provide the installer with reachable addresses for completing node-specific installation and configuration tasks before the final network settings are put into service. The required quantity is therefore based on the number of nodes being deployed: a six-node cluster requires six temporary IP addresses, so that every node has an address available during this stage. In addition, the deployment plan must explicitly associate each temporary address with its intended node. This one-to-one mapping lets the installation personnel identify the correct target node when performing configuration, connectivity checks, and subsequent replacement with the planned production network configuration. Accurate temporary-address planning is part of the network-preparation work and avoids ambiguity while nodes are being initialized. Huawei's scale-out storage documentation and product resources describe the planning and deployment context for OceanStor 9000 clusters: [Huawei OceanStor 9000 product page](#) and [Huawei Enterprise Support documentation search](#).

QUESTION NO: 54

In the OceanStor V3 storage system, regarding the remote replication consistency group, what is wrong?

- A. All members must synchronize, split, disconnect and switch between master and slave together.
- B. It is used to maintain the time consistency of mirroring data between multiple LUNs.
- C. A consistency group can have both synchronous remote replication and asynchronous remote replication.
- D. For remote replication pairs in the same consistency group, the primary LUNs must reside on the same storage array, and the secondary LUNs must reside on the same storage array.

ANSWER: C**Explanation:**

"A consistency group can have both synchronous remote replication and asynchronous remote replication" is the incorrect statement. In OceanStor V3, a remote replication consistency group is a coordinated set of remote replication pairs intended to preserve a common recoverable point in time for application data distributed across multiple LUNs. Consequently, its member pairs are managed as one unit for consistency-sensitive operations, and the replication mode is a consistency-group-level characteristic. All replication pairs included in one remote replication consistency group must use the same replication mode: either synchronous remote replication or asynchronous remote replication. This requirement ensures that write ordering, synchronization behavior, and the point at which data is made consistent at the secondary site are handled uniformly for every member LUN. Mixing synchronous and asynchronous pairs would give the members different data-protection timing and acknowledgement behaviors, so they could not reliably represent one consistent application image during operations such as synchronization, split, or recovery. Huawei's OceanStor remote-replication documentation describes consistency groups as a mechanism for maintaining data consistency across multiple LUNs and managing their replication relationships collectively. See the [Huawei OceanStor documentation](#) and the [Huawei Storage Product Documentation](#).

QUESTION NO: 55

Which options are correct in the following statements about the features of Huawei SmartPartition ()

- A. The SmartPartition feature is effective for both block business and file business
- B. Support setting SmartPartition in units of business LUNs
- C. The resources of each SmartPartition can only be accessed independently
- D. Support users to set the capacity of read cache and write cache for each SmartPartition partition

ANSWER: B C D**Explanation:**

“Support setting SmartPartition in units of business LUNs” is correct because SmartPartition is a controller-cache resource management feature for block services. Administrators associate service LUNs with a cache partition so that workloads with distinct performance requirements can be managed separately. This enables important or latency-sensitive applications to have explicitly allocated cache resources rather than competing indiscriminately with all other LUN workloads.

“The resources of each SmartPartition can only be accessed independently” is correct in the sense of SmartPartition’s cache-resource isolation model. Each configured partition has its own assigned cache resources and is used by the LUNs belonging to that partition. This isolation prevents the cache consumption of one service group from exhausting the capacity reserved for another group, helping provide predictable performance and limiting cross-application interference.

“Support users to set the capacity of read cache and write cache for each SmartPartition partition” is also correct. SmartPartition permits cache allocation to be tailored to workload behavior, including separate read and write cache settings. Huawei describes SmartPartition as a storage performance feature for allocating and isolating cache resources among service workloads. See Huawei’s [OceanStor data storage overview](#) and the [Huawei enterprise storage documentation portal](#) for feature documentation applicable to the relevant OceanStor version.

QUESTION NO: 56

Multiple choices Which of the following are correct in the description of the CIFS protocol? (Multiple choices)

- A. Transfer data based on TCP protocol
- B. Transfer data based on UDP protocol
- C. Linux platform file sharing protocol
- D. Windows platform file sharing agreement

ANSWER: A D

Explanation:

Transfer data based on TCP protocol is correct because CIFS is a dialect of the Server Message Block file-sharing protocol family and uses reliable, connection-oriented TCP transport. In practical deployments, SMB/CIFS traffic is normally carried directly over TCP port 445. Older NetBIOS-based implementations can use TCP port 139 after establishing a NetBIOS session, but the underlying reliable transport remains TCP. This reliability is important for file-system operations, including ordered delivery of file data, metadata updates, locking, authentication exchanges, and error recovery.

Windows platform file sharing agreement is also correct. CIFS was developed as an extension of SMB and is historically associated with Microsoft Windows network file and printer sharing. Storage systems expose CIFS shares so Windows clients can access directories and files using familiar UNC paths, while applying share permissions and file-level access controls. Although modern terminology commonly uses SMB rather than CIFS, CIFS remains widely used in storage administration to describe Windows-compatible file-sharing services. Microsoft documents SMB as the network file-sharing protocol used for shared access to files and related resources; its transport behavior and ports are documented in the linked protocol specifications. See [Microsoft SMB overview](#) and [Microsoft SMB protocol specification](#).

QUESTION NO: 57

(single choice) The network connection is required before the initial configuration of the big data system. The networking of Huawei OceanStor 9000 does not include:

- A. Front-end business network
- B. Back-end storage network
- C. Management network
- D. Heartbeat network

ANSWER: D

Explanation:

Heartbeat network is the correct answer. OceanStor 9000 deployment networking is organized around functional network planes used to provide client-facing storage services, internal storage communication, and system administration. The front-end business network carries access traffic between clients and the file-storage cluster. The back-end storage network supports the distributed storage communication required among cluster nodes, and the management network is used for management access and operational tasks during installation and subsequent maintenance. These defined network roles must be planned and connected before initial system configuration so that the cluster can be deployed, managed, and made available to service clients.

Huawei OceanStor 9000 documentation does not define a separate, deployable network category called a heartbeat network. Cluster-state and node-status communication are handled through the product's designated internal network architecture rather than by requiring an independently planned heartbeat network as part of the installation networking design. Therefore, when identifying the network type that is not included in the OceanStor 9000 networking model, **Heartbeat network** is the applicable choice. Huawei's product support portal provides OceanStor 9000 documentation and deployment materials at [Huawei OceanStor 9000 Support](#). Huawei's storage product documentation resources are also available through the [Huawei Enterprise Documentation portal](#).

QUESTION NO: 58

Which of the following is the maintenance software supported by Huawei 's hybrid flash storage products?

- A. SmartKit
- B. OceanStor UltraPath
- C. BCManager eReplication
- D. OceanStor eSDK

ANSWER: A**Explanation:**

SmartKit is Huawei's maintenance and service tool for enterprise storage products, including Huawei hybrid flash storage systems. It provides guided workflows that help engineers perform routine operational and maintenance activities consistently, such as device inspection, health checks, information collection, software upgrade assistance, expansion-related tasks, and fault diagnosis. By packaging these procedures into a centralized tool, SmartKit reduces the likelihood of manual configuration errors and improves maintenance efficiency throughout the storage system lifecycle.

In a hybrid flash storage environment, administrators need a supported tool that understands the storage platform's hardware and software status and can collect diagnostic information in a format suitable for service and support activities. SmartKit fulfills this role through product-specific toolkits and maintenance guidance. Huawei identifies SmartKit as an intelligent maintenance tool for enterprise products and provides its associated documentation and downloads through its Enterprise Support portal. Refer to the [Huawei SmartKit portal](#) and the [Huawei Enterprise Support portal](#) for supported products, tool packages, and maintenance documentation.

QUESTION NO: 59 - (SIMULATION)

In the local high-availability solution using Huawei HyperMirror, which of the following statements is correct? (multiple choice)

- E. You only need to connect Huawei OceanStorV3 storage to the business host.
- F. Both Huawei OceanStorV3 storage and heterogeneous storage need to be connected to the business host.
- G. Heterogeneous storage does not need to be connected to the business host
- H. You only need to connect heterogeneous storage to the business host.

ANSWER: See the explanation for the answer

Explanation:

Correct answers: E and G.

E. Only the Huawei OceanStor V3 storage needs to be connected to the business host.

G. The heterogeneous storage does not need to be connected to the business host.

In a local HA solution based on [HyperMirror](#), the business host accesses services through the OceanStor V3 storage. The heterogeneous array is used as the mirror-side storage and is connected behind/to the OceanStor system for mirror data processing; it is not directly presented or connected to the business host.

Therefore, options F and H are incorrect.

QUESTION NO: 60

A company purchased a set of Huawei horizontally extended file storage for the video surveillance system, and configured the InfoRevive feature

When providing fault-tolerant functions for video surveillance systems, which fault-tolerant modes can be selected

A. Read fault tolerance mode

B. Write fault tolerance mode

C. Read and write fault tolerance mode

D. Verify fault tolerance mode

ANSWER: A B C

Explanation:

InfoRevive provides transparent service fault tolerance for video-surveillance clients that access Huawei scale-out file storage. Its fault-tolerance policy is selected according to the type of client I/O that must continue when a storage node or access path experiences a fault. **Read fault tolerance mode** protects video retrieval and playback operations, allowing read requests to be redirected and serviced through an available path. **Write fault tolerance mode** protects recording and video-ingestion operations, so that ongoing write traffic can continue through the fault-tolerance mechanism. **Read and write fault tolerance mode** enables both protections together, which is suitable where the surveillance platform must sustain simultaneous recording and playback activity during a fault. These are the selectable InfoRevive modes and allow the administrator to match protection behavior to the surveillance workload. Huawei's enterprise support portal provides the product documentation and feature references for OceanStor scale-out storage, including the relevant configuration and maintenance guides: [Huawei Enterprise Storage Support](#). Huawei's storage product documentation landing page is also available at [Huawei Enterprise Documentation](#).

QUESTION NO: 61

During the installation of Commserver, if cS and Agent are on the same network segment, there is no need to configure the hosts file.

Pieces.

A. True

B. False

ANSWER: B

Explanation:

False is correct. Being on the same IP network segment establishes only Layer-3 reachability; it does not ensure that the CommServer and Agent can resolve each other's host names. CommServer deployment and Agent registration commonly

use host names or fully qualified domain names in configuration, service communication, certificate-related validation, or registration workflows. Therefore, each system must have reliable forward name resolution for its peer, whether that is provided by correctly configured DNS or by static mappings in the local hosts file. If DNS is unavailable, incomplete, or does not contain the required host records, hosts-file entries are required even when both components are connected to the same subnet. This also helps ensure that management traffic reaches the intended IP address consistently rather than depending on an unverified name-resolution path. Before installation, verify that the configured CommServer name resolves to the correct address from the Agent host and that the Agent name resolves correctly from the CommServer host. Huawei documentation resources for storage-management product deployment are available through the [Huawei Enterprise Documentation portal](#); general guidance on local hostname mappings is also described in the [Microsoft networking documentation](#).

QUESTION NO: 62

A company uses Huawei Distributed Storage Object Service to save transaction orders and enables WORM function

Yes, the administrator finds that some objects can only be read but cannot be modified, so which WORM state these objects may be in ()

- A. Protection status
- B. Protection expired
- C. Unprotected state
- D. Suspend protection

ANSWER: A B

Explanation:

Objects in **Protection status** have active WORM retention. WORM (write once, read many) preserves the original object data for the configured retention period, so the object remains available for reads while overwrite, alteration, and other operations that would change its protected content are blocked. This ensures that transaction-order records remain trustworthy and available for audit or compliance purposes.

Protection expired may also describe an object that is readable but not modifiable. Expiration means the mandatory retention period has ended and the object is no longer subject to the same retention-time restriction, particularly for removal according to the applicable policy and authorization. It does not turn a WORM object into an in-place editable object: object storage maintains object-content integrity rather than permitting modification of the preserved version. Therefore, both states can be encountered when an administrator observes read access without the ability to change the object's existing content.

Huawei's OBS documentation describes WORM retention as a mechanism for preventing protected object data from being overwritten or deleted during its retention period. See [Huawei Cloud OBS WORM documentation](#) and [Huawei Cloud OBS WORM FAQ](#).

QUESTION NO: 63

In a backup structure with a primary backup domain and multiple branch backup domains, the following description is wrong:

- A. Integrate multiple branch backup domains into one primary backup domain, and centrally view the operation reports of each backup domain in the primary backup domain.
- B. Each branch backup domain operates independently, and the connection network failure with the main backup domain does not affect the backup job of the branch backup domain.
- C. Metadata transmission can be automatically repaired.
- D. Register the remote client as an unconfigured client in the primary backup domain. The clients can be hidden and grouped.

ANSWER: C

Explanation:

"Metadata transmission can be automatically repaired." is the incorrect description. In a primary-and-branch backup-domain deployment, each branch domain retains independent backup-job operation and synchronizes management information with the primary domain for centralized visibility. That synchronization depends on connectivity and the supported synchronization mechanisms; it must not be treated as an automatic metadata-repair capability that transparently fixes metadata-transfer problems. If metadata synchronization is interrupted or becomes inconsistent, an administrator must verify communications, the branch-domain status, and synchronization results, then perform the applicable recovery or resynchronization procedure. This distinction is important because backup metadata is management-critical: it records protected objects, job status, copies, and recovery-related information. Administrators should therefore monitor synchronization health rather than assume that any metadata-transfer issue will repair itself without intervention. Huawei's OceanProtect data-protection materials describe centralized protection management and operational management capabilities, while the product documentation should be used for the relevant version's synchronization and fault-handling procedures. See [Huawei OceanProtect](#) and [Huawei OceanProtect Data Protection technical support](#).

QUESTION NO: 64

Colleagues customize active-active disaster recovery solutions for customers. The customer already has a third-party storage system and plans to purchase a new OceanStor

In 5300 V3, the two sets of equipment are placed in two data centers to implement active-active disaster recovery. The solution:

- A. Feasible, the storage device meets the requirements of the self-sufficient active-active disaster recovery solution.
- B. Not feasible. The active-active disaster recovery solution requires at least two sets of OceanStor V3 storage.
- C. Not feasible. The active-active disaster recovery solution only supports Huawei storage and does not support third-party storage.

ANSWER: B

Explanation:

Not feasible. The active-active disaster recovery solution requires at least two sets of OceanStor V3 storage. is correct because the OceanStor V3 active-active implementation, commonly delivered through HyperMetro, establishes a synchronized active-active relationship between two compatible Huawei OceanStor V3 arrays deployed at the two sites. Each array participates in the HyperMetro domain, maintains the replicated LUN relationship, and presents an active-active service path to hosts. Therefore, purchasing only one OceanStor 5300 V3 while retaining a third-party array at the other site does not provide the required pair of OceanStor V3 systems for this solution architecture. The two-site design also relies on the associated inter-array replication links and, where required, an arbitration mechanism to maintain service continuity during site or link failures. Huawei's OceanStor documentation describes HyperMetro as a storage active-active solution and provides the product documentation and compatibility information needed when planning supported array models and configurations. See the [Huawei OceanStor V3 documentation portal](#) and Huawei's [OceanStor storage product information](#) for active-active storage and product-reference materials.

QUESTION NO: 65

The following single-choice descriptions that do not belong to the features of the horizontally extended file storage function are:

- A. Distributed metadata management, supports linear expansion of 3~4000 nodes.
- B. Self-developed distributed file system OceanStor DFS supports massive unstructured data storage and realizes a global unified namespace.
- C. Provide object storage service, compatible with Amazon S3 interface and compatible with OpenStack Swif interface.
- C. Users can easily build an enterprise-level big data analysis platform.

ANSWER: A

Explanation:

“Distributed metadata management, supports linear expansion of 3~4000 nodes.” is the description that does not belong to the specified horizontal scale-out file-storage feature set because its stated scalability range is inaccurate for the OceanStor 9000/OceanStor DFS capability addressed by this exam material. OceanStor 9000 uses a distributed architecture in which nodes contribute performance and capacity as the cluster grows, and distributed metadata is a key mechanism that enables the scale-out file system and a unified file-service environment. However, the documented node-scale specification for the relevant OceanStor 9000 generation is 3 to 288 nodes, not 3 to 4000 nodes. Therefore, although the statement correctly names distributed metadata management and linear scale-out as architectural concepts, the incorrect maximum node quantity makes the complete statement unsuitable as a feature description. This distinction matters in solution design and certification questions: capacity and performance planning must use the platform’s supported cluster limits rather than a generalized or inflated scale-out number. Huawei’s OceanStor 9000 product information describes its scale-out storage architecture and distributed file-service positioning; see [Huawei OceanStor 9000](#). Additional product documentation and resources are available through [Huawei Enterprise Support: OceanStor 9000](#).

QUESTION NO: 66

Which of the following features are used for data migration solution?

- A. SmartPartition
- B. SmartVirtualization
- C. SmartQuota
- D. SmartMigration

ANSWER: D

Explanation:

SmartMigration is the Huawei storage feature specifically designed to move data from a source LUN to a target LUN while maintaining host service continuity. It is used when data must be relocated for purposes such as storage consolidation, performance-tier adjustment, hardware replacement, capacity rebalancing, or planned maintenance. The migration task is managed by the storage system, which synchronizes source data to the destination and allows the host-facing service to continue through the migration workflow. This makes SmartMigration an appropriate feature for a data-migration solution rather than merely a capacity-management or partitioning function.

In practical migration designs, SmartMigration can also be used alongside storage virtualization capabilities when data originates from external storage. Huawei’s Oracle ASM migration procedure explicitly describes using SmartMigration together with SmartVirtualization: virtualization presents the external storage resources to the Huawei storage system, and SmartMigration performs the data relocation process. Therefore, where one feature must be selected as the direct data-migration feature, SmartMigration is the correct choice. See Huawei’s [SmartMigration and SmartVirtualization migration procedure](#) and the [SmartMigration feature documentation](#).

QUESTION NO: 67

What thresholds can be set for eReplication? (multiple choice)

- A. CPU utilization
- B. Memory utilization
- C. Hard disk utilization
- D. Bandwidth utilization

ANSWER: A B C D

Explanation:

eReplication threshold settings cover both the local system resources required to run replication tasks and the network resource used to transfer replicated data. CPU utilization and memory utilization thresholds allow administrators to detect when replication processing may be competing excessively with other workloads on the protected host. Hard disk utilization is also monitored because insufficient available capacity can prevent replication data, logs, or task-related files from being written reliably. Bandwidth utilization is a relevant threshold because replication relies on the communication link between sites; monitoring or constraining bandwidth consumption helps keep replication traffic within the capacity planned for the link and protects other business traffic. Together, these thresholds provide a practical basis for proactive alarms and resource planning, helping ensure that replication services remain stable while production applications continue to receive adequate compute, storage, and network resources. Huawei's enterprise storage and data-protection portfolio describes replication as a cross-site data-protection capability whose operation depends on storage, host, and connectivity resources. See Huawei's [storage and data-protection portfolio](#) and the [Huawei Enterprise Support documentation portal](#) for product documentation and configuration guidance.

QUESTION NO: 68

Multiple choice The description about the analysis and adjustment of the Oceanstor 9000 system NAS server parameters is wrong: (multiple choice)

- A. If the NFS transmission data block is very small, it may increase the number of business interactions.
- B. If the data block transmitted by NFS is very large, it may increase the number of business interactions
- C. ocnstor 9000 NFS transmission block size defaults to 1MB
- D. ocnstor 9000 NFS transfer block size is 32KB by default

ANSWER: C D**Explanation:**

The incorrect default-value statements are "ocnstor 9000 NFS transmission block size defaults to 1MB" and "ocnstor 9000 NFS transfer block size is 32KB by default." For OceanStor 9000 NAS service tuning, the NFS transmission block size has a default value of 128 KB. Therefore, neither 1 MB nor 32 KB accurately states the platform's default configuration. NFS transfer size is a performance-tuning parameter that should be assessed together with workload characteristics, client behavior, network capacity, and the size distribution of files being accessed. A larger setting can be appropriate for sequential, large-file workloads, while a smaller setting may be considered for workloads dominated by smaller I/O operations; however, such tuning choices do not redefine the system default. Administrators should verify the release-specific parameter behavior and apply adjustments through documented procedures, because supported ranges and recommended values can vary by OceanStor 9000 software release. Huawei's support portal provides the relevant product documentation and searchable knowledge resources: [Huawei Enterprise Support](#) and [Huawei Support search for OceanStor 9000 NFS transmission block size](#).

QUESTION NO: 69 - (SIMULATION)

In the VMware environment, the single-choice scenario where VADP cannot be used for backup is: A. Thick Disk

- B. Thin Disk
- C. The storage is directly connected to the virtual machine
- D. VMDK

ANSWER: See the explanation for the answer**Explanation:**

Correct answer: C. The storage is directly connected to the virtual machine.

VADP (vStorage APIs for Data Protection) backs up VMware virtual-machine disks managed through the VMware virtualization/storage layer, such as VMDK files. Both thick-provisioned and thin-provisioned virtual disks can be protected by VADP.

When storage is directly attached/presented to the guest VM (for example, direct physical storage access rather than a VMware-managed virtual disk), it cannot be backed up through VADP. It requires an in-guest backup method or another backup approach appropriate for that direct-attached storage.

QUESTION NO: 70

On the basis of inheriting the high reliability design of OceanStor storage system, HyperMetro has newly designed some solutions

Case-level high-reliability technology maximizes the reliability of the storage active-active solution. HyperMetro's high-reliability technology is mainly reflected in

Which of the following? (multiple choice)

- A. Cross-site cluster
- B. Cross-site data real-time mirroring
- C. Arbitration to prevent split brain
- D. Cross-site bad block repair

ANSWER: A B C D

Explanation:

HyperMetro extends the inherent reliability mechanisms of OceanStor with metro active-active capabilities designed to keep storage services available when a storage device, site, or link experiences a fault. **Cross-site cluster** provides coordinated operation between the storage systems at the two sites, allowing host services to use either side of the active-active configuration. **Cross-site data real-time mirroring** maintains synchronized copies of the data across the sites, so both sides can provide consistent access to the same storage resources.

Arbitration to prevent split brain is also fundamental: during an inter-site communication failure, quorum-based arbitration determines which side is permitted to continue providing services, protecting data consistency from independent concurrent writes. **Cross-site bad block repair** further improves media-level resilience by using the healthy mirrored copy to repair data affected by a bad block, reducing the chance that a localized disk-media problem becomes an application-visible data issue.

Together, these mechanisms address coordinated availability, synchronous data consistency, safe fault isolation, and data repair—the key reliability design areas of HyperMetro. Huawei describes HyperMetro as a storage active-active disaster-recovery solution in its [HyperMetro product information](#) and provides OceanStor technical documentation through the [Huawei OceanStor support portal](#).

QUESTION NO: 71

Multiple choices Among the following statements about the features of Huawei SmartPartition, which options are correct ()

- A. Support users to set the capacity of read cache and write cache separately for each SmartPartition partition
- B. The SmartPartition feature is effective for both block business and file business
- C. Support setting SmartPartition in units of business LUNs
- D. The resources of each SmartPartition can only be accessed independently

ANSWER: A B C

Explanation:

Huawei SmartPartition is a cache-resource management feature that lets administrators divide the controller cache into logical partitions and assign cache resources according to application requirements. “Support users to set the capacity of read cache and write cache separately for each SmartPartition partition” is correct because a partition has independently configurable read-cache and write-cache capacities, enabling the cache policy to match a workload’s read/write profile. “The SmartPartition feature is effective for both block business and file business” is also correct: SmartPartition can provide cache-resource assurance and isolation for storage services delivered through block and file protocols. “Support setting SmartPartition in units of business LUNs” is correct because SmartPartition can associate cache partitions with service LUNs, allowing important block workloads to receive planned cache capacity rather than competing unpredictably for shared cache resources. This capability is intended to improve predictable service performance where multiple applications coexist on one storage system. Huawei storage documentation describes SmartPartition as a mechanism for allocating and managing cache resources for services; see the [Huawei OceanStor product documentation](#) and Huawei’s [storage feature documentation](#) for feature concepts and configuration guidance.

QUESTION NO: 72

What are the reasons for rapid development of the enterprise cloud storage?

- A. Enterprises have strong demands on resource pooling, shortened product go-to-market time, and reduced cost.
- B. Storing service data to the cloud can reduce the possibility of misoperation and malicious theft compared to storing the data on the local PC.
- C. Enterprise cloud storage can help improve the performance and reliability of typical or new applications.
- D. For typical application scenarios such as VSI, VDI, and big data, enterprise cloud storage conserves more storage space, provides larger bandwidth, and enhances compatibility between virtual machines (VMs).

ANSWER: A

Explanation:

Enterprises have strong demands on resource pooling, shortened product go-to-market time, and reduced cost. This is the key business driver behind the rapid development of enterprise cloud storage. Cloud storage consolidates previously isolated storage resources into shared pools that can be allocated and adjusted on demand. This pooling model improves overall utilization and makes capacity available more quickly for new services, rather than requiring each application team to purchase, deploy, and manage separate infrastructure. Standardized service delivery and automation also help organizations provision storage faster, which supports shorter time to market. At the same time, centralized operations, elastic resource use, and reduced duplication of infrastructure lower both capital expenditure and ongoing management cost. These benefits align with Huawei’s discussion of cloud transformation, in which enterprises adopt cloud-oriented infrastructure to improve resource utilization, agility, and operational efficiency. Huawei storage offerings similarly emphasize resource pooling and service-oriented delivery as foundations for enterprise cloud data infrastructure. See Huawei’s [Digital Transformation white paper](#) and the [Huawei Storage](#) overview.

QUESTION NO: 73

Which of the following options is correct for the description of the multi-version function of the bucket:

- A. Users can use the multi-version function to save, retrieve and restore each version of the object.
- B. Multi-version function cannot be used for data retention and archiving
- C. The multi-version status of the bucket can be configured and queried.
- D. Whether the multi-version function of the bucket is turned on or off is marked by the multi-version status of the bucket.

ANSWER: A C D

Explanation:

Users can use the multi-version function to save, retrieve and restore each version of the object. When versioning is enabled for a bucket, the storage service retains distinct versions of an object when that object is overwritten or deleted. Each version has its own version ID, allowing a user to retrieve a specified historical version or restore data by making a prior version current. This provides an important safeguard against accidental overwrite and deletion.

The multi-version status of the bucket can be configured and queried. Bucket versioning is a bucket-level setting: administrators can enable, suspend, and view its current state through supported management interfaces. The configured state governs how subsequent object updates and deletes are handled.

Whether the multi-version function of the bucket is turned on or off is marked by the multi-version status of the bucket. Huawei OBS documents versioning states such as Enabled and Suspended, which indicate whether the bucket actively creates new object versions. This explicit status allows applications and administrators to determine the bucket's versioning behavior before performing object-management operations. See Huawei's [OBS versioning documentation](#) and [bucket versioning API reference](#).

QUESTION NO: 74

Which of the following special single options is mainly used for Huawei's two-site three-center disaster recovery solution?

- A. HyperMirror
- B. HyperMetro
- C. HyperClone
- D. HyperReplication

ANSWER: D

Explanation:

HyperReplication is the key remote-replication capability used to extend a two-site three-center disaster-recovery architecture to the third site. In this architecture, production data is replicated from the active-active metropolitan-area storage environment to a geographically separate disaster-recovery center. HyperReplication supports the creation and maintenance of remote copies between storage systems, providing the data-protection path required for recovery when a large-scale site or regional event affects the primary metropolitan area. Depending on business recovery objectives, replication can be configured with an appropriate synchronization mode and replication policy so that the remote copy is usable for disaster recovery. This capability is therefore central to protecting data beyond the two local production sites and implementing the third-center recovery layer. Huawei describes its enterprise storage portfolio and data-protection capabilities, including remote replication and disaster-recovery use cases, on its [enterprise storage product page](#). Additional Huawei storage solution information is available from the [enterprise storage solutions page](#).

QUESTION NO: 75

Huawei all-flash storage devices need to be planned and multi-selected before deployment. Which of the following options are required for storage networking

To plan ()

- A. Heartbeat network
- B. IP address
- C. Management network
- D. Business networking

ANSWER: A B C D

Explanation:

Huawei all-flash storage deployment requires a complete network plan that separates and addresses the management, internal high-availability, and host-service communication paths. **Management network** planning provides the connectivity used by administrators and management tools to access and operate the storage system. The **Heartbeat network** is planned for communication between controllers, supporting high availability and coordinated controller operation. **Business networking** must also be designed because it carries the storage service traffic between application hosts and the storage system, such as block or file access traffic depending on the selected protocol. In addition, **IP address** planning is essential to make these networks usable: the planned addresses, masks, gateways, and VLAN associations identify management, service, and internal communication interfaces and prevent address conflicts. Huawei deployment guidance emphasizes completing network and IP planning before installation and initial configuration so that the storage system can be connected, configured, and put into service reliably. See Huawei's [All-Flash Storage product information](#) and the [Huawei Enterprise Storage Support portal](#) for product documentation and deployment resources.

QUESTION NO: 76

Bank A purchased Huawei's hybrid flash storage, which uses encrypted hard disk technology to enhance business data

The security, which of the following statements are correct ()

- A. External Key Manager adopts KMIP+TLS standard protocol
- B. Internal Key Manager is a key management application embedded in the storage system
- C. , And guarantee no impact on performance
- D. OceanStor 18000/18000FV5 supports TPM for key protection

ANSWER: A B C D

Explanation:

Huawei encrypted-disk implementations use a layered key-management design. An external key manager communicates using the Key Management Interoperability Protocol over a Transport Layer Security protected connection, allowing encryption keys to be managed outside the storage array with authenticated, encrypted key exchanges. This interoperable model is consistent with the KMIP specification, which defines key-management operations and secure client/server communication requirements. See the [OASIS KMIP specification](#).

The embedded internal key manager is also a storage-resident key-management application, providing key lifecycle functions when an external manager is not used. For encrypted disks, encryption and decryption are performed by the disk encryption hardware, so normal host I/O does not incur an additional controller-side encryption processing bottleneck. In the OceanStor 18000/18000F V5 platform family, a Trusted Platform Module protects sensitive root-key material, establishing a hardware root of trust for the key hierarchy. Huawei's enterprise storage portfolio and its security capabilities are documented on the [Huawei Storage product page](#).

QUESTION NO: 77

Which of the following statements are true about SmartQoS on a Huawei OceanStor storage system?

- A. A QoS policy can limit the maximum IOPS of a service object.
- B. A QoS policy can limit the maximum bandwidth of a service object.
- C. Priority can be configured to differentiate services during resource contention.
- D. SmartQoS synchronizes data from a primary LUN to a remote LUN.

ANSWER: A B C

Explanation:

SmartQoS is used to control service-resource consumption and maintain predictable performance for different workloads. A QoS policy can limit IOPS or bandwidth for an object so that a workload does not consume excessive storage resources. Priority control can also be used to differentiate services when resources are contended.

SmartQoS controls performance behavior; it does not provide data replication or replace host-side multipathing. Refer to the [Huawei Enterprise Storage support portal](#) for release-specific SmartQoS configuration guidance.

QUESTION NO: 78

After Huawei Distributed Storage enables the object deduplication function, the system will automatically find duplicate objects globally.

And only keep one copy of the same object and delete the redundant duplicate data. The original duplicate data is replaced with a single copy that points to the reservation.

A copy of the reference.

A. True

B. False

ANSWER: A

Explanation:

True is correct. Object deduplication in Huawei distributed storage is a space-efficiency capability that identifies identical object content throughout the applicable global deduplication scope. Once duplicate objects are found, the system retains a single physical instance of the object data and replaces the redundant physical instances with references to that retained instance. Consequently, users and applications can continue accessing the objects through their original logical names or paths, while storage capacity is consumed only once for the shared content. This is especially valuable in object-storage workloads containing repeatedly uploaded files, backup copies, software packages, images, or other identical unstructured data. The process is transparent to applications: deduplication changes the physical storage representation, not the logical accessibility of the object. Huawei describes its distributed storage products as providing data-reduction and efficiency capabilities for large-scale data services; see the [Huawei Distributed Storage product page](#) and Huawei's [enterprise documentation portal](#) for product documentation and feature details.

QUESTION NO: 79

Huawei Oceanstor 9000 provides external domain name resolution services through the "InfoEqualizerDNS IP address",

Name access to the following types of file sharing. If a node failover occurs, which of the following sharing methods will not be interrupted?

A. NFSv3 file sharing

B. NFSv4 file sharing

C. CIFS file sharing

D. FTP file sharing

ANSWER: A

Explanation:

NFSv3 file sharing is the uninterrupted service type in this scenario because NFSv3 is fundamentally stateless at the protocol level. An NFSv3 client sends independent RPC requests, rather than relying on a persistent protocol session whose server-side state must be transferred during a node failover. OceanStor 9000 InfoEqualizerDNS provides the externally visible domain-name resolution mechanism for the file-service nodes. When a serving node fails, the client can resolve or reconnect through the InfoEqualizer address and continue issuing NFSv3 operations to an available node. Any individual RPC affected while the failure occurs can be retried by the client, allowing service access to continue without a session-

recovery process that would interrupt the share connection. This behavior is particularly appropriate for scale-out NAS deployments, where DNS-based access distribution and stateless NFSv3 requests support resilience across nodes. Huawei describes OceanStor 9000 as a scale-out storage platform providing NAS file services; its product information is available at [Huawei OceanStor Scale-Out Storage](#). Huawei's enterprise support documentation portal, where model- and version-specific OceanStor 9000 file-service and InfoEqualizer guidance is published, is available at [Huawei Enterprise Support Documentation](#).

QUESTION NO: 80

What type of backup networking has the following features?

1. Occupies large network bandwidth;
2. Has restricted backup performance;
3. Adversely affects host applications.

- A. LAN-Base
- B. LAN-Free
- C. Server-Less
- D. Server-Free

ANSWER: A

Explanation:

LAN-Base is the correct backup-networking type. In a LAN-based backup architecture, the backup client running on the production host reads data from its local file system or application and sends that data across the regular LAN to a backup server or backup storage target. Both production application traffic and backup traffic therefore use the same host network interfaces and LAN infrastructure. During a backup window, substantial volumes of data can consume available LAN bandwidth, while the host must also spend CPU, memory, and I/O resources preparing and transmitting backup data. These shared-resource demands can constrain achievable backup throughput and can affect the performance of production applications on the protected host. This description matches the listed characteristics: high network-bandwidth consumption, limited backup performance due to shared LAN and host resources, and a possible adverse effect on host applications. Huawei's enterprise storage portfolio includes data-protection solutions intended to improve backup and recovery architectures; see [Huawei OceanProtect Data Protection](#) and the [Huawei Enterprise Storage Support](#) portal for product and documentation resources.