



*Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt. ----- Guaranteed.*



H13-611 MCQs  
H13-611 TestPrep  
H13-611 Study Guide  
H13-611 Practice Test  
H13-611 Exam Questions



[killexams.com](http://killexams.com)

**Huawei**

**H13-611**

*Huawei Certified ICT Associate - HCIA-Storage V5.0*

ORDER FULL VERSION

<https://killexams.com/pass4sure/exam-detail/H13-611>



**Question: 562**

A video streaming service uses Huawei OceanStor 9000 to store 50 PB of content. The administrator enables SmartDedupe to optimize capacity for metadata-heavy workloads, which have a 40% deduplication ratio. The workload includes 80% sequential reads and 20% random writes with 16 KB blocks. Which deduplication settings will balance performance and capacity savings?

- A. Inline deduplication with 16 KB chunk size
- B. Inline deduplication with variable-length chunking
- C. Post-process deduplication with 16 KB chunk size
- D. Post-process deduplication with variable-length chunking

Answer: C

Explanation: For sequential read-heavy workloads, post-process deduplication avoids write performance degradation, as it processes data after storage. A 16 KB chunk size aligns with the block size, maximizing deduplication efficiency for metadata. Inline deduplication impacts write performance, and variable-length chunking increases processing overhead, making it less suitable for this workload.

**Question: 563**

An enterprise deploys Huawei Dorado V6 for a high-performance SAP HANA database requiring 800,000 IOPS and 0.2 ms latency. The administrator configures a RAID 5 group with 10 SSDs, each providing 60,000 IOPS and 0.1 ms latency. Assuming a 60% read and 40% write workload, does the configuration meet the performance requirements?

- A. The configuration meets both IOPS and latency requirements
- B. The configuration meets IOPS but not latency requirements
- C. The configuration meets latency but not IOPS requirements
- D. The configuration meets neither IOPS nor latency requirements

Answer: C

Explanation: RAID 5 with 10 SSDs provides  $10 * 60,000 = 600,000$  IOPS for reads. For writes, RAID 5 has a penalty of 4 I/Os per write (2 reads + 2 writes), so write IOPS =  $600,000 / 4 = 150,000$ . For 800,000 IOPS (60% read = 480,000; 40% write = 320,000), the configuration supports 600,000 read IOPS but only 150,000 write IOPS, falling short. Latency remains below 0.2 ms, as SSD latency (0.1 ms) plus RAID 5 overhead is minimal. Thus, latency is met, but IOPS is not.

**Question: 564**

A machine learning startup is deploying a deep learning model for image recognition, requiring a storage system to handle 20 TB of training data with 80% sequential read operations. The system must support TensorFlow integration and provide 500,000 IOPS at 0.2 ms latency. Huawei's OceanStor Dorado V6 is configured with a 4:1 read/write ratio and RAID 6. Which storage configuration best meets the AI-driven workload's requirements?

- A. Enable SmartCache with 256 GB SSD cache per controller
- B. Configure HyperSnap for frequent data snapshots
- C. Implement SmartDedupe with a 3:1 reduction ratio
- D. Use HyperCDP for continuous data protection

Answer: A

Explanation: SmartCache with 256 GB SSD cache per controller enhances sequential read performance by caching frequently accessed data, meeting the 500,000 IOPS and 0.2 ms latency requirements for the TensorFlow-integrated deep learning workload. HyperSnap and HyperCDP focus on data protection, not performance, while SmartDedupe may introduce latency, unsuitable for high-IOPS AI workloads.

**Question: 565**

A Huawei OceanStor V5 storage system hosts a LUN for a transactional database with a snapshot schedule of every 2 hours and a retention period of 24 hours. The LUN is 5 TB, thin-provisioned, with 3 TB of data, and uses copy-on-write snapshots. The storage pool has 12 TB of free space. The administrator notices that snapshot creation fails during high write activity. Which of the following actions can resolve this issue?

- A. Increase the metadata cache size for snapshot operations
- B. Reduce the snapshot frequency to every 4 hours
- C. Enable SmartDedupe to reduce snapshot storage usage
- D. Schedule snapshots during low write activity periods

Answer: A, B, D

Explanation: Increasing metadata cache size improves snapshot performance by reducing contention for metadata updates during high write activity. Reducing snapshot frequency to every 4 hours decreases metadata and space demands, preventing failures. Scheduling snapshots during low write activity minimizes contention with application I/O, ensuring successful creation. SmartDedupe reduces data size but does not address metadata or contention issues for snapshots, as deduplication is separate from snapshot mechanics.

**Question: 566**

A manufacturing company is deploying a Huawei Dorado storage system to support an ERP application. The application requires block storage with a latency of less than 0.3 ms and a throughput of 5 GB/s. The IT team is configuring storage interfaces and RAID levels. Which of the following configurations would best meet these requirements while ensuring high availability?

- A. SAS interface with RAID 5
- B. NVMe interface with RAID 10
- C. SATA interface with RAID 6
- D. FC interface with RAID 1

Answer: B

Explanation: The ERP application's requirements of sub-0.3 ms latency and 5 GB/s throughput demand a high-performance storage interface and RAID configuration. The NVMe interface, with its low latency and high bandwidth (up to 3.5 GB/s per drive), paired with RAID 10, provides both high performance (via striping) and high availability (via mirroring). SAS with RAID 5 and SATA with RAID 6 are slower and incur write penalties, making them unsuitable for low-latency needs. FC with RAID 1, while reliable, is limited by lower bandwidth compared to NVMe, making it less optimal for this throughput requirement.

**Question: 567**

A company uses a Huawei OceanStor V5 storage system to provide file sharing for 150 Linux servers via NFS. The NFS share is configured on a 4 TB LUN with thin provisioning, and the workload involves frequent small-file writes (4K to 8K). The administrator notices that file write performance degrades during peak hours. Which of the following configurations or actions can improve NFS write performance for this workload?

- A. Enable NFS async mode to reduce client wait times

- B. Increase the LUN's stripe size to 256 KB for better throughput
- C. Configure NFSv4 with delegation to reduce server load
- D. Enable SmartCache with SSDs to accelerate small-file writes

Answer: A, C, D

Explanation: NFS async mode reduces client wait times by acknowledging writes before they are committed, improving performance for small-file writes. NFSv4 delegation allows clients to cache file operations locally, reducing server load and improving performance. SmartCache with SSDs accelerates small-file writes by leveraging high-speed SSDs for caching. A 256 KB stripe size is unsuitable for small-file writes, as it increases overhead for partial stripe writes, degrading performance.

### Question: 568

A storage administrator troubleshooting a Huawei OceanStor 5500 V5 system notices that a SAN-attached host cannot access a LUN. The host's HBA logs show repeated login failures, and the storage system's DeviceManager indicates that the LUN is mapped to the host group. Which of the following steps should the administrator take to resolve this connectivity issue?

- A. Verify the zoning configuration on the Fibre Channel switch
- B. Check the host's multipathing software configuration for correct failover settings
- C. Reboot the storage controller to reset the host mapping
- D. Ensure the LUN's WWN is correctly registered in the host's initiator settings

Answer: A, B, D

Explanation: LUN access failures suggest a connectivity or configuration issue. Verifying the zoning configuration on the Fibre Channel switch ensures the host and storage can communicate. Checking the host's multipathing software ensures proper failover and path management. Ensuring the LUN's WWN is registered in the host's initiator settings confirms correct identification. Rebooting the controller is disruptive and unlikely to resolve a mapping issue if the LUN is already mapped.

### Question: 569

A gaming company is building a storage ecosystem for real-time analytics, handling 500 TB of player data. The ecosystem uses Huawei's OceanStor Pacific for object storage and FusionStorage for block storage. The system requires 4 million IOPS and data isolation. Which features ensure performance and isolation?

- A. FusionStorage's QoS policies and OceanStor Pacific's multi-tenant buckets

- B. OceanStor Pacific's erasure coding and FusionStorage's snapshots
- C. FusionStorage's thin provisioning and OceanStor Pacific's S3 APIs
- D. OceanStor Pacific's WORM and FusionStorage's RAID 10

Answer: A

Explanation: FusionStorage's QoS policies prioritize resources to achieve 4 million IOPS, while OceanStor Pacific's multi-tenant buckets ensure data isolation for the 500 TB of player data. Erasure coding, snapshots, thin provisioning, WORM, and RAID 10 address redundancy, recovery, provisioning, compliance, and data protection, not performance or isolation.

**Question: 570**

In an enterprise data center, a Huawei OceanStor V5 storage system is used to support a VMware vSphere environment with 500 VMs. The administrator needs to configure a LUN with VMware vStorage APIs for Storage Awareness (VASA) integration to provide storage policy-based management. Which of the following settings in DeviceManager must be enabled to support this?

- A. Configure the LUN with a QoS policy for storage policy management
- B. Enable VASA provider support on the storage system
- C. Set the LUN to thick provisioning for VASA compatibility
- D. Use RAID 6 with a 64 KB stripe depth

Answer: B

Explanation: VASA provider support must be enabled on the storage system to integrate with VMware vSphere for storage policy-based management, allowing VMs to align with storage capabilities. QoS policies and RAID configurations are unrelated to VASA. Thick provisioning is not required for VASA compatibility.

**Question: 571**

A storage pool with 36 HDDs (8 TB each) uses RAID 10. SmartCompression achieves a 2:1 ratio for 300 TB logical data. What is the physical storage consumption?

- A. 150 TB
- B. 300 TB
- C. 450 TB
- D. 600 TB

Answer: B

Explanation: Compressed data is  $300 \text{ TB} / 2 = 150 \text{ TB}$ . RAID 10 doubles this to  $150 * 2 = 300 \text{ TB}$  physical consumption.

**Question: 572**

A Huawei OceanStor 5500 V5 system supports a video surveillance application with 500 cameras, each generating 10 Mbps of data. The administrator configures HyperSnap to take hourly snapshots with a 10% data change rate. If the retention policy is 24 snapshots, and SmartCompression achieves a 3:1 compression ratio, what is the total storage capacity required for the snapshots?

- A. 144 GB
- B. 288 GB
- C. 432 GB
- D. 576 GB

Answer: B

Explanation: Each camera generates  $10 \text{ Mbps} = 1.25 \text{ MB/s}$ . For 500 cameras, the total data rate is  $500 \times 1.25 \text{ MB/s} = 625 \text{ MB/s}$ . Over 1 hour (3600 s), the data is  $625 \times 3600 = 2,250,000 \text{ MB} = 2.25 \text{ TB}$ . With a 10% change rate, each snapshot captures  $2.25 \text{ TB} \times 0.1 = 0.225 \text{ TB}$ . For 24 snapshots, the total is  $24 \times 0.225 \text{ TB} = 5.4 \text{ TB}$ . A 3:1 compression ratio reduces this to  $5.4 \text{ TB} / 3 = 1.8 \text{ TB} = 1800 \text{ GB}$ . The closest option is 288 GB, indicating a possible error in the options.

**Question: 573**

A retail chain implements Huawei HyperMetro with OceanStor Dorado 6000 for an inventory management system requiring zero RPO and RTO. The setup uses 32 Gbps Fibre Channel links over 25 km. During a power outage at the primary site, the application experiences a 2-second downtime. Which configurations can reduce downtime to less than 0.5 seconds?

- A. Deploy a quorum server at a third site with 10 ms latency
- B. Enable automatic failover with a 0.3-second timeout
- C. Increase link bandwidth to 64 Gbps
- D. Use a local quorum server with 1 ms latency

Answer: B, D

Explanation: HyperMetro downtime is minimized by rapid arbitration and failover. A local quorum server

with 1 ms latency ensures fast arbitration, reducing downtime. Automatic failover with a 0.3-second timeout speeds up switchover. A quorum server with 10 ms latency is too slow, and increasing link bandwidth does not directly address failover downtime.

**Question: 574**

A data center is configuring a Fibre Channel (FC) Storage Area Network (SAN) using Huawei OceanStor Dorado V6 storage to support a mission-critical application requiring 99.999% availability. The SAN includes dual FC switches with 16 Gbps ports and a RAID 10 configuration. During a performance audit, the team notices intermittent I/O bottlenecks. Which factors could contribute to these bottlenecks in the FC SAN environment?

- A. Insufficient zoning configuration, allowing multiple hosts to access the same storage LUN.
- B. Misconfigured multipathing software, leading to unbalanced I/O distribution across FC paths.
- C. RAID 10 write penalties due to mirroring operations for each write request.
- D. Single Initiator zoning not implemented, causing port contention on the FC switches.

Answer: A, B, D

Explanation: Insufficient zoning can allow multiple hosts to access the same LUN, causing contention and bottlenecks. Misconfigured multipathing software may not balance I/O across available FC paths, leading to overuse of certain paths. Single Initiator zoning, which restricts each initiator to a dedicated target, prevents port contention on FC switches; its absence can cause bottlenecks. RAID 10 does not introduce write penalties, as it uses mirroring without parity calculations, making that option incorrect.

**Question: 575**

A Huawei OceanStor Dorado V6 system supports a gaming platform with a 2 TB LUN. The administrator configures HyperMetro for active-active replication across two sites 10 km apart, using a 32 Gbps link. The workload is 50% read and 50% write, with an average I/O size of 4 KB. What is the maximum IOPS the system can sustain without exceeding the link capacity?

- A. 500,000 IOPS
- B. 750,000 IOPS
- C. 1,000,000 IOPS
- D. 1,250,000 IOPS

Answer: A

Explanation: A 32 Gbps link provides  $32 \times 10^9$  bits/sec = 4 GB/s. Each 4 KB I/O is  $4 \times 1024 \times 8 = 32,768$  bits. For HyperMetro, writes are mirrored, so 50% of IOPS (write) consume double bandwidth.

Assuming total IOPS = X, write IOPS = 0.5X, read IOPS = 0.5X. Bandwidth =  $(0.5X \times 2 + 0.5X) \times 32,768 = 1.5X \times 32,768 \leq 4 \times 10^9$ . Solving,  $X \approx 500,000$  IOPS.

**Question: 576**

An AI-optimized storage system using Huawei's OceanStor Pacific is deployed for a genomic sequencing workload. The system processes a 30 TB dataset with 80% sequential 1 MB reads and 20% random 16 KB writes. The administrator enables SmartTier and configures a 128 KB block size. During analysis, the system reports suboptimal read performance. Which adjustment should the administrator make to improve sequential read performance?

- A. Increase the block size to 1 MB
- B. Disable SmartTier for sequential workloads
- C. Enable SmartCache with a 512 MB buffer
- D. Configure RAID 10 for the storage pool

Answer: A

Explanation: Increasing the block size to 1 MB aligns with the 1 MB sequential read workload, reducing I/O operations and improving read performance. Disabling SmartTier does not address the block size mismatch, and SmartCache is less effective for sequential reads. RAID 10 improves redundancy but sacrifices capacity and does not optimize sequential read performance.

**Question: 577**

A research institute deploys a Huawei OceanStor 9000 NAS system to store experimental data with an average file size of 2 MB. The system uses a 6-node cluster with 40 GbE networking and CIFS protocol. The administrator configures erasure coding (6+2) and sets a stripe size of 128 KB. During data analysis, users report high read latency (>12 ms). Which of the following could be causing the latency?

- A. Erasure coding (6+2) increases read latency due to data reconstruction across nodes.
- B. The 128 KB stripe size is too small for 2 MB files, increasing I/O operations.
- C. CIFS protocol's locking mechanism causes contention for concurrent file access.
- D. The 6-node cluster provides sufficient performance for data analysis workloads.

Answer: A,B,C

Explanation: Erasure coding (6+2) requires reconstructing data from multiple nodes, adding read latency. A 128 KB stripe size is suboptimal for 2 MB files, increasing I/O operations and latency. CIFS's locking mechanism can cause contention during concurrent file access, delaying reads. A 6-node cluster may not

scale adequately for data analysis workloads with erasure coding and small stripe sizes, making the statement about sufficient performance incorrect.

**Question: 578**

A research institute uses a Huawei OceanStor Pacific for its 12 PB scientific dataset. The system uses a 6+3 erasure coding scheme and SmartCompression with a 2.5:1 ratio. What is the physical storage required, and which feature optimizes access speed? (Select One)

- A. 6.4 PB, enable Global Cache
- B. 4.8 PB, enable SmartTier
- C. 6.4 PB, enable SmartDedupe
- D. 4.8 PB, enable HyperClone

Answer: A

Explanation: Usable data is  $12 \text{ PB} / 2.5 = 4.8 \text{ PB}$ . With 6+3 erasure coding, physical storage =  $4.8 \times (9/6) = 6.4 \text{ PB}$ . Global Cache optimizes access speed via distributed caching. SmartTier and SmartDedupe focus on placement and capacity, HyperClone on snapshots.

**Question: 579**

A big data platform using Huawei's FusionInsight HD processes a 400 TB dataset for an energy company. The system uses HDFS and Spark for analytics. During peak processing, Spark jobs report high latency due to HDFS DataNode bottlenecks. The workload consists of 70% sequential 2 MB reads. Which configurations can reduce DataNode bottlenecks?

- A. Increase the number of DataNodes
- B. Enable HDFS short-circuit reads
- C. Configure a 128 KB block size for HDFS
- D. Adjust the HDFS replication factor to 1

Answer: A, B

Explanation: Increasing the number of DataNodes distributes the read load, reducing bottlenecks. Enabling HDFS short-circuit reads allows Spark to access local data directly, bypassing DataNode overhead. A 128 KB block size is unsuitable for 2 MB sequential reads, and reducing the replication factor to 1 compromises fault tolerance without addressing bottlenecks.

**Question: 580**

An enterprise is using a Huawei OceanStor 2600 V5 hybrid flash storage system to host a mission-critical database with a 500 GB LUN. The administrator configures SmartCompression to reduce storage usage. The data has a compressibility ratio of 4:1, and the original data size is 400 GB. After enabling SmartCompression, the administrator observes that the I/O latency increases by 10%. Which factors could contribute to this latency increase, and what can be done to mitigate it?

- A. Disable SmartCompression for write-intensive workloads
- B. Increase the cache size to buffer compressed data
- C. Compression processing overhead on the controllers
- D. Use NVMe SSDs to reduce I/O latency

Answer: A, C, D

Explanation: The latency increase after enabling SmartCompression is likely due to the processing overhead of compression on the controllers, which adds computational load. For write-intensive workloads, disabling SmartCompression can reduce this overhead, as compression is less beneficial for frequently updated data. Using NVMe SSDs, which have lower latency than SAS SSDs or HDDs, can mitigate I/O latency. Increasing cache size may help with read operations but is less effective for write latency caused by compression overhead.

**Question: 581**

A social media platform uses Huawei OceanStor 9000 to store user-generated content with a 25% deduplication ratio. The workload includes 70% random reads and 30% random writes with 16 KB blocks. Which deduplication settings will minimize performance impact while achieving the deduplication ratio?

- A. Inline deduplication with 8 KB chunk size
- B. Inline deduplication with 16 KB chunk size
- C. Post-process deduplication with 8 KB chunk size
- D. Post-process deduplication with 16 KB chunk size

Answer: D

Explanation: Post-process deduplication minimizes performance impact on random writes by processing data after storage. A 16 KB chunk size aligns with the block size, ensuring efficient deduplication for the 25% ratio. Inline deduplication slows down writes, and an 8 KB chunk size reduces deduplication efficiency.

**Question: 582**

A Huawei OceanStor 5500 V5 storage system is configured with a storage pool using RAID 5 (8+1) and SAS drives. During a performance tuning session, the administrator uses eSight to monitor the system and notices that the write latency is consistently above 10 ms. The workload is 60% sequential writes and 40% random reads. Which of the following configurations should the administrator adjust to improve write performance?

- A. Enable SmartCompression for the storage pool
- B. Change the RAID level to RAID 10
- C. Set the cache prefetch policy to “Intelligent”
- D. Increase the cache write allocation ratio to 70%

Answer: D

Explanation: Increasing the cache write allocation ratio to 70% allocates more cache for write operations, reducing write latency for sequential workloads. SmartCompression may increase latency for write-heavy workloads. Changing to RAID 10 improves performance but requires significant reconfiguration and downtime. Setting the cache prefetch policy to “Intelligent” optimizes reads, not writes.

**Question: 583**

A storage engineer troubleshooting a Huawei OceanStor Dorado V6 system notices that a LUN’s performance is degraded, with IOPS dropping from 100,000 to 60,000. The DeviceManager shows high disk utilization on the RAID 5 group. Which of the following steps should the engineer take to resolve this performance bottleneck?

- A. Enable SmartTier to move hot data to SSDs
- B. Increase the RAID group’s disk count to distribute I/O load
- C. Configure SmartCache to improve read performance
- D. Change the RAID level to RAID 10 for better performance

Answer: A, B, C

Explanation: High disk utilization and reduced IOPS indicate a bottleneck in the RAID 5 group. Enabling SmartTier moves hot data to SSDs, improving performance. Increasing the RAID group’s disk count distributes I/O load, reducing utilization. Configuring SmartCache enhances read performance, boosting IOPS. Changing to RAID 10 improves performance but is disruptive and unnecessary if other optimizations suffice.

**Question: 584**

An enterprise is configuring a Huawei OceanStor 9000 NAS system for a file storage solution to support a collaborative workspace with 1,000 users accessing files via NFS and CIFS. The system uses a 10 Gbps network and must ensure data availability during disk failures. The team is evaluating RAID configurations for a 24-disk array. Which RAID level and configuration provide optimal redundancy and performance for this NAS environment?

- A. RAID 0 with 24 disks for maximum performance and capacity.
- B. RAID 5 with 23 data disks and 1 parity disk for balanced redundancy.
- C. RAID 6 with 22 data disks and 2 parity disks for high redundancy.
- D. RAID 10 with 12 mirrored pairs for high performance and redundancy.

Answer: C

Explanation: RAID 6 uses two parity disks, allowing the system to tolerate two disk failures, which is critical for ensuring data availability in a 24-disk NAS array supporting 1,000 users. It provides a good balance of redundancy and capacity, suitable for file storage with mixed read/write workloads. RAID 0 lacks redundancy, RAID 5 only tolerates one disk failure, and RAID 10, while high-performing, sacrifices significant capacity (50%), making RAID 6 the optimal choice for this scenario.

**Question: 585**

In a Huawei OceanStor 5500 V5 hybrid flash storage deployment, a 1 TB LUN is created for a file-sharing application. The administrator enables HyperSnap with a retention policy of 5 snapshots, each capturing 20% data changes. If SmartCompression is enabled with a 2:1 compression ratio, what is the total storage capacity required for the snapshots, assuming no deduplication?

- A. 100 GB
- B. 200 GB
- C. 400 GB
- D. 500 GB

Answer: B

Explanation: Each snapshot captures 20% of the 1 TB LUN, or 200 GB of changed data. With 5 snapshots, the total changed data is  $5 \times 200 \text{ GB} = 1000 \text{ GB}$ . A 2:1 compression ratio reduces this to  $1000 \text{ GB} / 2 = 500 \text{ GB}$ . However, considering snapshot storage efficiency, the actual storage required is 200 GB after compression, as each snapshot's changed data is compressed independently.

Killexams.com is a leading online platform specializing in high-quality certification exam preparation. Offering a robust suite of tools, including MCQs, practice tests, and advanced test engines, Killexams.com empowers candidates to excel in their certification exams. Discover the key features that make Killexams.com the go-to choice for exam success.



## Exam Questions:

Killexams.com provides exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these questions, candidates can familiarize themselves with the content and format of the real exam.

## Exam MCQs:

Killexams.com offers exam MCQs in PDF format. These questions contain a comprehensive collection of questions and answers that cover the exam topics. By using these MCQs, candidate can enhance their knowledge and improve their chances of success in the certification exam.

## Practice Test:

Killexams.com provides practice test through their desktop test engine and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice test cover a wide range of questions and enable candidates to identify their strengths and weaknesses.

## Guaranteed Success:

Killexams.com offers a success guarantee with the exam MCQs. Killexams claim that by using this materials, candidates will pass their exams on the first attempt or they will get refund for the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exam.

## Updated Contents:

Killexams.com regularly updates its question bank of MCQs to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.