

Huawei H23-211_V1.0

HCSA-Field-Scale-Out Storage V1.0

**Thank You for Downloading H23-211_V1.0 Updated Exam
Questions**

https://certs4sale.com/huawei/h23-211_v1.0-pdf-exam-dumps

Question: 1

Which task should be completed before a storage project is formally handed over to the customer?

- A. Removing all system logs and inspection results
- B. Completing health checks, service verification, acceptance records, and delivery documentation
- C. Disabling all redundancy mechanisms
- D. Disconnecting one power supply from each node

Answer: B

Explanation:

Before handover, the delivery team should confirm that hardware, networks, cluster services, storage resources, and client access are operating as designed. System alarms and health-check findings should be reviewed, and unresolved issues should be documented and corrected. The team should complete installation records, configuration documents, test results, asset information, and customer acceptance materials. Customer administrators may also require basic operation and maintenance guidance. Logs and reports should be retained because they establish the system's condition at delivery. Redundancy and power connections must remain properly configured to ensure service availability.

Question: 2

Compared with three complete data copies, what is a common characteristic of an appropriately designed erasure-coding policy?

- A. It provides no protection against disk failure
- B. It requires every fragment to be stored on the same node
- C. It can provide comparable fault tolerance with less raw-capacity overhead
- D. It eliminates all processing and network overhead

Answer: C

Explanation:

Erasur coding can protect data by storing data fragments together with calculated parity fragments. Depending on the selected scheme, the system can reconstruct the original data after one or more fragment failures. Because parity fragments are smaller than multiple complete copies, erasure coding often uses raw capacity more efficiently than three-copy replication. However, it does not eliminate processing or network overhead. Encoding, decoding, reconstruction, and data movement consume resources, so workload requirements must be considered. Replication may be preferred for some latency-sensitive or small-file workloads, while erasure coding can be advantageous for large, capacity-intensive datasets.

Question: 3

Which traffic normally travels across the storage back-end network?

- A. Only user web-browser traffic
- B. Internal communication, data distribution, reconstruction, and synchronization between storage nodes
- C. Only administrator login traffic

D. External email communication

Answer: B

Explanation:

The back-end storage network carries internal cluster traffic between storage nodes. Depending on the architecture, this traffic may include data placement, metadata operations, replication, erasure-code calculations, recovery, reconstruction, and cluster-control communication. Because large quantities of data can move across this network, it requires sufficient bandwidth, low latency, proper redundancy, and consistent configuration. A congested or incorrectly configured back-end network can reduce system performance and extend recovery times. Client service traffic and administrative traffic may use separate networks. During installation, engineers should verify interface bonding, VLAN configuration, link speed, switch configuration, and communication between every required node pair.

Question: 4

Which protocol allows a client to mount a remote directory and access it using file-system operations?

- A. NFS
- B. S3
- C. SNMP
- D. NTP

Answer: A

Explanation:

NFS provides file-level access by allowing a client to mount an exported directory from a storage server. After mounting the directory, users and applications can perform familiar file operations such as opening, reading, writing, renaming, and deleting files, subject to permissions. S3 uses an object-based access model rather than a traditional mounted file system. SNMP supports monitoring, while NTP synchronizes time. During deployment, engineers must check the NFS service version, export paths, client access rules, user identity mapping, network connectivity, firewall settings, and mount parameters to ensure reliable and secure access.

Question: 5

Which statement best describes scale-out storage architecture?

- A. Capacity can grow only by replacing all existing controllers
- B. Every workload must be stored on one server
- C. Multiple nodes operate as a cluster and can be added to expand resources
- D. The system cannot redistribute data after expansion

Answer: C

Explanation:

Scale-out storage combines multiple nodes into a unified cluster. Data and service workloads are distributed across these nodes, and additional nodes can be added to expand capacity, processing resources, bandwidth, or all three. The cluster may automatically rebalance data after expansion so that new resources are used effectively. This differs from scale-up storage, where expansion is often limited to adding components within one storage frame or replacing controllers with larger models. Scale-out architecture is well suited to environments with unpredictable growth, but deployment still requires proper

network, protection, compatibility, and failure-domain planning.

THANK YOU FOR DOWNLOADING H23-211_V1.0 UPDATED EXAM QUESTIONS

Note: Thanks For Trying The Demo Of Our H23-211_V1.0 Exam Product

Visit Our Site to Purchase the Full Set of Actual H23-211_V1.0 Exam Questions With Answers.

Money Back Guarantee



Click The Link Below

https://certs4sale.com/huawei/h23-211_v1.0-pdf-exam-dumps