Huawei H13-629_V3.0 Dumps

Shared by Whitfield on 29-09-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

When creating a VM on the FusionCube platform, you can configure CPU QoS based on the CPU performance requirements of services to ensure that the minimum guaranteed computing power, memory, and I/O resources can be obtained during VM running.

Options:

A- TRUE
B- FALSE

Answer:

B

Explanation:

The statement is FALSE because it attributes guarantees for CPU, memory, and I/O resources collectively to CPU QoS.

CPU QoS applies to the allocation and scheduling of CPU computing resources. Memory and I/O are distinct resource classes and are controlled through their own resource-management mechanisms. A CPU QoS policy cannot by itself guarantee the VM's minimum memory allocation and storage/network I/O resources.

Huawei's FusionCompute architecture explicitly separates compute virtualization management, storage virtualization management, and network virtualization management, reflecting the fact that these resource domains are managed independently. The Lab Guide likewise demonstrates separate resource reservation parameters for CPU and memory, configuring ''Reserved CPU (%)'' and ''Reserved Memory (%)'' independently rather than treating memory as part of CPU QoS.

FusionCube combines compute, virtualization, networking, and scale-out storage into a unified platform, but convergence does not eliminate the distinction between their QoS controls.

Therefore, CPU QoS can protect or regulate CPU computing capability, but it does not independently guarantee CPU, memory, and I/O resources together. The overall statement is consequently FALSE.

Study Guide Topics: Data Storage Technology Presentation FusionCube FusionCompute VM Resource Management CPU QoS Resource Isolation.

Question 2

Question Type: MultipleChoice

Huawei's storage data migration system can be divided into multiple layers. Each layer corresponds to different migration solutions. Which of the following is the migration solution corresponding to the host file system layer?

Options:

A- Migration solution based on Migration Director for Block
B- Migration solution based on MigrationDirector for NAS Storage
C- Migration solution based on MigrationDirector for NAS
D- Migration solution based on MigrationDirector for SAN

Answer:

C

Explanation:

The correct answer is C. MigrationDirector for NAS.

Huawei's technical structure for storage migration separates migration solutions according to the layer at which data is accessed and moved. The official training diagram associates NAS storage migration with the file-system layer and identifies MigrationDirector for NAS as the dedicated migration tool.

NAS data exists as files and directories exposed through file protocols rather than as raw SAN block devices. MigrationDirector for NAS therefore works with file-system-oriented migration workflows. Huawei's NAS migration procedure uses a migration server, configures a migration task, performs initial synchronization, applies access permissions and quotas as required, then performs incremental synchronization and service cutover.

MigrationDirector for SAN belongs to block/SAN migration rather than the host file-system layer. Similarly, a generic block migration tool would operate below the file-system abstraction.

Therefore, for the host file system layer, the dedicated Huawei migration solution is MigrationDirector for NAS.

Study Guide Topics: Data Migration Solution Technical Structure File System Layer NAS Migration MigrationDirector for NAS.

Question 3

Question Type: MultipleChoice

For Huawei all-flash storage, which of the following statements are true about link compression in HyperReplication?

Options:

A- Link compression compresses all data (regardless of upper-layer service data or other data) that is transmitted over the network.
B- Link compression does not affect services running on the hosts and is imperceptible to users.
C- Inline compression can be performed on data transmitted over links in real time.
D- After receiving data, the secondary end decompresses it and then verifies its correctness.

Answer:

A, B, C

Explanation:

The supplied HCIE-Storage question set identifies A, B, and C as correct.

Link compression is an inter-array transmission optimization associated with HyperReplication. Because it operates on the replication path rather than in the host/application data path, it is transparent to the application server. The host continues to perform ordinary reads and writes against the primary storage system while HyperReplication handles the remote synchronization process.

Huawei describes HyperReplication as the core technology used to synchronize copies of local data to a remote storage system and distinguishes synchronous and asynchronous replication modes. For asynchronous replication, changed data is collected and transferred to the secondary system according to the replication period. Compression can therefore be applied inline to the transmission stream to reduce required WAN bandwidth.

Option D is not part of the keyed link-compression sequence. The retrieved Training Materials establish the HyperReplication transport and consistency mechanisms, but the parsed course text available here does not provide a detailed link-compression processing sequence that independently substantiates the specific ''decompress first, then verify'' assertion. Therefore, that additional sequence should not be inferred.

Study Guide Topics: Data Protection Solution HyperReplication Remote Replication Link Link Compression Bandwidth Optimization.

Question 4

Question Type: MultipleChoice

When planning a disaster recovery (DR) system, DR levels need to be divided based on service continuity requirements and cost investment. Which of the following statements about DR levels are true?

Options:

A- In service-level DR, all service links are redundant, including links between IT systems and non-IT systems such as office facilities and personnel.
B- Application-level DR consists of building the same application systems at the DR site.
C- In data-level DR, the remote DR center can be simply considered as a remote data backup center.
D- In data-level DR, services are not interrupted if a disaster occurs.

Answer:

A, B, C

Explanation:

Huawei divides DR into data-level, application-level, and service-level protection according to recovery requirements and investment.

For data-level DR, the remote site primarily protects copies of production data. Huawei explicitly states that a data-level DR center cannot prevent service interruption when a disaster occurs; it operates essentially as a backup center from which data can later be restored. Therefore, C is correct and D is false.

For application-level DR, the backup site contains the same application system as the production site and data is synchronized between the sites using synchronous or asynchronous mechanisms. This allows applications to be restarted at the DR site within the defined RTO, making B correct.

Service-level DR is the most comprehensive level. Huawei defines it as a complete DR environment covering IT technologies and infrastructure as well as non-IT resources such as offices and telecommunications. A redundant workplace may therefore be required if the normal business site becomes unavailable.

Hence, A, B, and C are correct.

Study Guide Topics: Data Protection Solution DR Basics DR Levels Data-Level Application-Level

Service-Level DR.

Question 5

Question Type: MultipleChoice

The OceanProtect Appliance uses Oracle's built-in Recovery Manager (RMAN) to back up the Oracle database. After a user initiates a data backup job, the RMAN plug-in in the appliance invokes the database backup interface to write data into the backup storage.

Options:

A- TRUE
B- FALSE

Answer:

B

Explanation:

The statement is FALSE because it incorrectly places the application-side backup plug-in inside the OceanProtect appliance.

Huawei's OceanProtect architecture separates the ProtectManager/DataMoveEngine components on the protection system from ProtectClient components that protect application data on hosts. The course specifically states that ProtectClient protects host application data and that the application being protected determines which agent software is installed.

The database-protection workflow therefore uses OceanProtect to schedule and coordinate the job, while application-aware operations are performed through the client/agent located with the protected database environment. In analogous Huawei database workflows, DataMoveEngine creates backup storage space, instructs ProtectClient to start the backup operation, ProtectClient mounts the storage space, and the host-side application plug-in invokes the database-specific backup mechanism.

Thus, even though Oracle RMAN is Oracle's native backup/recovery technology, saying that an RMAN plug-in in the appliance directly invokes the database interface misrepresents the architecture. Application-aware execution occurs through the protection client/agent associated with the Oracle host.

Study Guide Topics: Data Protection Solution OceanProtect Database Backup ProtectClient RMAN

ProtectManager/DataMoveEngine.

Question 6

Question Type: MultipleChoice

During the installation and deployment of Huawei flash storage, complying with the disk enclosure connection specifications can effectively reduce the impact on performance. Which of the following statements are true about these specifications?

Options:

A- A single loop can connect to both SAS disk enclosures and smart disk enclosures.
B- Disk enclosures are evenly distributed to multiple loops.
C- The number of disk enclosures added to a loop is less than 2.
D- The loops are distributed to multiple back-end interface modules of the controller.

Answer:

B, C, D

Explanation:

The correct selections are B, C, and D.

The design objective behind Huawei back-end enclosure connectivity is to distribute back-end traffic and avoid concentrating storage bandwidth, enclosure load, or failure domains on a single connection path. Consequently, disk enclosures should be distributed across multiple loops and controller back-end interface resources.

Huawei's training material describes back-end direct-connect networking using multiple back-end interface modules and requires symmetrical interface-module/port placement. For smart disk enclosures, the uplink and downlink ports are used according to a prescribed cascade topology, and only limited cascading depth is supported. This explains the reasoning behind B, C, and D: distribute enclosures, limit the number connected on each loop, and spread loops across back-end interface modules.

Option A is incorrect because SAS disk enclosures and smart disk enclosures use different back-end connection architectures and must not simply be mixed arbitrarily within one loop.

The exact B, C, D combination is also explicitly provided in the supplied HCIE-Storage question set.

Correct back-end topology planning improves available bandwidth, preserves redundancy, and reduces performance degradation caused by excessive enclosure concentration.

Study Guide Topics: Data Storage Technology Presentation Flash Storage Hardware Disk Enclosures Back-End Networking Loop Distribution.

Question 7

Question Type: MultipleChoice

During the OceanProtect Appliance single-node system network planning, which of the following network planes are mandatory?

Options:

A- Replication network
B- Internal communication network
C- Backup network
D- Management network

Answer:

C, D

Explanation:

The mandatory network planes are the backup network and the management network.

Huawei defines the management network plane as the network used for web-console login, service configuration, and alarm reporting. Without it, administrators cannot perform the fundamental control and administration functions required by the appliance. Huawei defines the backup network plane as the network used for backup data and in-band signaling transmission, making it essential for the appliance's core data-protection function.

The replication network, by contrast, is needed when backup copies must be replicated to another data-protection appliance. If remote replication is not part of the deployment, this plane is not mandatory. Similarly, an archive network is optional because the backup network can be reused for archiving under supported configurations.

An internal communication network becomes relevant to multi-node architectures where separate nodes must communicate internally. The question explicitly specifies a single-node system, so this separate internal plane is not mandatory.

Therefore, the baseline single-node OceanProtect networking requirement is C and D.

Study Guide Topics: Data Protection Solution OceanProtect Appliance Network Planning Management Network Backup Network Replication Network.

Question 8

Question Type: MultipleChoice

DME is a piece of centralized management and O&M software developed by Huawei for storage devices in data centers. Which of the following unstructured data management capabilities are provided by DME?

Options:

A- Optimization
B- Analysis
C- Planning
D- Search

Answer:

A, B, C, D

Explanation:

All four listed capabilities are included in the DME data-management/O&M framework.

Huawei describes DME Storage as an intelligent storage management platform covering the complete storage lifecycle, including planning, deployment, O&M, and optimization. Its architecture includes intelligent diagnosis, intelligent optimization, resource optimization, planning simulation, operation visibility, insight analysis, and multidimensional analysis.

Planning enables administrators to model and organize resources before deployment. Analysis provides performance, capacity, topology, alarm, and workload insight. Optimization uses those analytics to recommend or execute improvements such as capacity balancing, migration, SLA adjustment, and QoS changes. Search is important for unstructured-data environments because files, objects, and other unstructured content must be located and analyzed according to metadata.

Huawei OceanStor Pacific also provides SmartIndexing for unstructured file, object, and HDFS storage. It can return search results from extremely large metadata sets within seconds and

provides unified multi-protocol metadata indexing.

Consequently, DME's unstructured-data management workflow spans the complete chain of search, analysis, planning, and optimization.

Study Guide Topics: Data Storage Management and O&M DME Unstructured Data Management Search Analysis Planning Optimization.

Question 9

Question Type: FillInTheBlank

The OceanStor Pacific series scale-out storage system consists of multiple node roles. The ___ component must be deployed on management nodes, compute nodes, and storage nodes for the system to run properly. (Enter an abbreviation.)

Answer:

MDC

Explanation:

The required component is MDC.

OceanStor Pacific uses a distributed control architecture in which cluster-management and metadata/control services must remain highly available as nodes are added, removed, or experience failures. Huawei specifically identifies multi-MDC as a mechanism that improves system reliability. The training architecture also shows MDC participating in sub-health management together with storage-side processes such as OSD and EDS.

This distributed deployment is consistent with the architecture of OceanStor Pacific: storage services are spread across nodes, and control information must remain accessible regardless of the physical node role. The system therefore cannot depend on one isolated control instance.

Huawei also describes OceanStor Pacific as a three-layer scale-out architecture integrating block, file, object, and HDFS services, with cluster-status control, distributed data routing, redundancy management, rebuilding, and balancing forming core storage-engine functions. MDC supports the distributed coordination and control required by this architecture.

The supplied question set identifies MDC as the required abbreviation.

Study Guide Topics: Data Storage Technology Presentation OceanStor Pacific Scale-Out Architecture MDC Cluster Control High Reliability.

Question 10

Question Type: DragDrop

Huawei migration service provides a wide range of technical solutions, which can be selected based on customer requirements to achieve their ultimate data migration objectives. Match the technical solutions with the correct layers.



Answer:

See the Answer in the Premium Version!

Question 11

Question Type: MultipleChoice

During the LUN data migration using Huawei Migration Director for SAN, which of the following operations need to be performed in the migration task configuration phase?

Options:

A- Obtaining source storage information
B- Adding service hosts
C- Planning the target storage
D- Taking over services

Answer:

A, B, C

Explanation:

Huawei identifies MigrationDirector for SAN as the dedicated migration tool for SAN storage migration. Before migration execution, the migration environment must be defined: the source storage must be identified, the relevant service host must be known, and the target storage/resources must be prepared.

The Lab Guide's workflow places Creating a Service Host, Configuring the Source Storage, and Preparing the Target Storage before the actual Data Migration phase. A SmartVirtualization migration lab follows the same structure: production/service-host preparation and target-storage preparation precede takeover and data migration.

This makes A, B, and C configuration/preparation activities.

Option D belongs later in migration execution. Huawei's migration delivery sequence explicitly places target-storage configuration and migration preparation before stopping services, taking over the source storage, recovering applications, and migrating data.

Therefore, service takeover is an implementation/cutover activity rather than information entered during migration-task configuration.

Study Guide Topics: Data Migration Solution MigrationDirector for SAN Migration Task Configuration Source Storage Service Host Target Storage Planning Cutover.

Question 12

Question Type: MultipleChoice

Source deduplication is generally classified into two types in the industry: general source deduplication and application-aware source deduplication. Which of the following statements about general source deduplication is true?

Options:

A- Deduplication can be enabled or disabled, depending on the application.
B- It can improve the performance of application servers.
C- For data that already exists in the backup system, only metadata needs to be transmitted.
D- Deduplication can be enabled or disabled, depending on the backup type.

Answer:

C

Explanation:

Huawei describes source deduplication as occurring before backup data is transferred to storage media. Data or files are segmented, fingerprints are calculated, and those fingerprints are sent to the DPA fingerprint repository for duplicate checking.

If an identical fingerprint already exists, Huawei states that the duplicate data block is not saved; instead, the existing block is referenced. Only when no duplicate fingerprint exists is the actual new data block transmitted.

This mechanism supports C: when the data content is already present in the backup system, the backup workflow needs only the metadata/reference information required to point to the existing block rather than retransmitting the duplicate payload.

The principal benefits are reduced network traffic, reduced backup-storage consumption, and a shorter backup window. The feature is not described as a mechanism whose purpose is to improve application-server performance; deduplication processing itself consumes some processing resources. Nor does the course define general source deduplication as being independently enabled or disabled according to application type or backup type in the way described by A and D.

Study Guide Topics: Data Protection Solution DPA Source Deduplication Fingerprints Duplicate Block Detection Transmission Reduction.

P2P exams

To Get Premium Files for H13-629_V3.0 Visit
https://www.p2pexams.com/products/h13-629_v3.0

For More Free Questions Visit
https://www.p2pexams.com/huawei/pdf/h13-629-v3.0

