

# IT 认证电子书



质 量 更 高 服 务 更 好

半年免费升级服务

<http://www.itrenzheng.com>

**Exam** : **E10-001**

**Title** : Information Storage and  
Management Exam Version  
2

**Version** : DEMO

1.Which cache management algorithm is based on the assumption that data will not be requested by the host when it has not been accessed for a while?

- A. LRU
- B. HWM
- C. LWM
- D. MRU

**Answer: A**

Explanation:

Cache Management: Algorithms Cache is a finite and expensive resource that needs proper management. Even though modern intelligent storage systems come with a large amount of cache, when all cache pages are filled, some pages have to be freed up to accommodate new data and avoid performance degradation. Various cache management algorithms are implemented in intelligent storage systems to proactively maintain a set of free pages and a list of pages that can be potentially freed up whenever required.

The most commonly used algorithms are discussed in the following list:

Least Recently Used (LRU): An algorithm that continuously monitors data access in cache and identifies the cache pages that have not been accessed for a long time. LRU either frees up these pages or marks them for reuse. This algorithm is based on the assumption that data that has not been accessed for a while will not be requested by the host.

However, if a page contains write data that has not yet been committed to disk, the data is first written to disk before the page is reused.

Most Recently Used (MRU): This algorithm is the opposite of LRU, where the pages that have been accessed most recently are freed up or marked for reuse. This algorithm is based on the assumption that recently accessed data may not be required for a while.

EMC E10-001 Student Resource Guide. Module 4: Intelligent Storage System

2.What does the area ID of the FC address identify?

- A. Group of ports within a switch
- B. An individual port within a fabric
- C. Location of the name server within the fabric
- D. Unique number provided to each switch in the fabric

**Answer: A**

Explanation:

FC Addressing in Switched Fabric An FC address is dynamically assigned when a node port logs on to the fabric. The FC address has a distinct format, as shown in the slide.

The first field of the FC address contains the domain ID of the switch. A Domain ID is a unique number provided to each switch in the fabric.

Although this is an 8-bit field, there are only 239 available addresses for domain ID because some addresses are deemed special and reserved for fabric management services. For example, FFFFFC is reserved for the name server, and FFFFFE is reserved for the fabric login service. The area ID is used to identify a group of switch ports used for connecting nodes. An example of a group of ports with common area ID is a port card on the switch. The last field, the port ID, identifies the port within the group. Therefore, the maximum possible number of node ports in a switched fabric is calculated as: 239 domains X 256 areas X 256 ports = 15,663,104

- FC Address is assigned to nodes during fabric login
  - ▶ Used for communication between nodes within FC SAN
- Address format



- Domain ID is a unique number provided to each switch in the fabric
  - ▶ 239 addresses are available for domain ID
- Maximum possible number of node ports in a switched fabric:
  - ▶ 239 domains X 256 areas X 256 ports = 15,663,104

EMC E10-001 Student Resource Guide. Module 5: Fibre Channel Storage Area Network (FC SAN)

3.An organization performs copy on first access (CoFA) replication to create a local replica of application data. To perform a successful restore, what should be considered?

- Source devices must be healthy
- Save location size must be larger than the size of all source devices
- Save location size must be equal to the size of all source devices
- All changes to the source and replica must be discarded before the restore starts

**Answer:** A

Explanation:

Replication: Restore & Restart Considerations Local replicas are used to restore data to production devices. Alternatively, applications can be restarted using the consistent point-in-time replicas. Replicas are used to restore data to the production devices if logical corruption of data on production devices occurs—that is, the devices are available but the data on them is invalid.

Examples of logical corruption include accidental deletion of data (tables or entries in a database), incorrect data entry, and incorrect data updates. Restore operations from a replica are incremental and provide a small RTO. In some instances, the applications can be resumed on the production devices prior to the completion of the data copy. Prior to the restore operation, access to production and replica devices should be stopped.

Production devices might also become unavailable due to physical failures, such as production server or physical drive failure. In this case, applications can be restarted using the data on the latest replica. As a protection against further failures, a “Gold Copy” (another copy of replica device) of the replica device should be created to preserve a copy of data in the event of failure or corruption of the replica devices. After the issue has been resolved, the data from the replica devices can be restored back to the production devices.

Full-volume replicas (both full-volume mirrors and pointer-based in Full Copy mode) can be restored to the original source devices or to a new set of source devices. Restores to the original source devices can be incremental, but restores to a new set of devices are full volume copy operations.

In pointer-based virtual and pointer-based full-volume replication in CoFA mode, access to data on the replica is dependent on the health and accessibility of the source volumes. If the source volume is inaccessible for any reason, these replicas cannot be used for a restore or a restart operation. EMC

E10-001 Student Resource Guide. Module 11: Local Replication

4. Which host component eliminates the need to deploy separate adapters for FC and Ethernet communications?

- A. Converged network adapter
- B. TCP Offload Engine NIC
- C. FCIP bridging adapter
- D. iSCSI host bus adapter

**Answer: A**

Explanation:

Converged Network Adaptor (CNA)

- Provides functionality of both – a standard NIC and an FC HBA
  - ▶ Eliminates the need to deploy separate adapters and cables for FC and Ethernet communications
- Contains separate modules for 10 Gigabit Ethernet, FC, and FCoE ASICs
  - ▶ FCoE ASIC encapsulates FC frames into Ethernet frames



A CNA provides the functionality of both a standard NIC and an FC HBA in a single adapter and consolidates both types of traffic. CNA eliminates the need to deploy separate adapters and cables for FC and Ethernet communications, thereby reducing the required number of server slots and switch ports. CNA offloads the FCoE protocol processing task from the server, thereby freeing the server CPU resources for application processing. A CNA contains separate modules for 10 Gigabit Ethernet, Fibre Channel, and FCoE Application Specific Integrated Circuits (ASICs). The FCoE ASIC encapsulate FC frames into Ethernet frames. One end of this ASIC is connected to 10GbE and FC ASICs for server connectivity, while the other end provides a 10GbE interface to connect to an FCoE switch.

EMC E10-001 Student Resource Guide. Module 6: IP SAN and FCoE

5. What is a function of unified management software in cloud computing?

- A. Defining cloud service attributes
- B. Consolidating infrastructure resources scattered across one or more data centers
- C. Metering based on usage of resources by the consumer
- D. Providing an interface to consumers to request cloud services

**Answer: B**

Explanation:

Cloud Management and Service Creation Tools The cloud management and service creation tools layer includes three types of software:

This classification is based on the different functions performed by these software. These software interact with each other to automate provisioning of cloud services.

The physical and virtual infrastructure management software is offered by the vendors of various infrastructure resources and third-party organizations. For example, a storage array has its own management software. Similarly, network and physical servers are managed independently using network and compute management software respectively. These software provide interfaces to construct a virtual infrastructure from the underlying physical infrastructure.

Unified management software interacts with all standalone physical and virtual infrastructure management software. It collects information on the existing physical and virtual infrastructure configurations, connectivity, and utilization. Unified management software compiles this information and provides a consolidated view of infrastructure resources scattered across one or more data centers. It allows an administrator to monitor performance, capacity, and availability of physical and virtual resources centrally. Unified management software also provides a single management interface to configure physical and virtual infrastructure and integrate the compute (both CPU and memory), network, and storage pools. The integration allows a group of compute pools to use the storage and network pools for storing and transferring data respectively.

The unified management software passes configuration commands to respective physical and virtual infrastructure management software, which executes the instructions. This eliminates the administration of compute, storage, and network resources separately using native management software.

The key function of the unified management software is to automate the creation of cloud services. It enables administrators to define service attributes such as CPU power, memory, network bandwidth, storage capacity, name and description of applications and platform software, resource location, and backup policy. When the unified management software receives consumer requests for cloud services, it creates the service based on predefined service attributes.

The user-access management software provides a web-based user interface to consumers. Consumers can use the interface to browse the service catalogue and request cloud services. The user-access management software authenticates users before forwarding their request to the unified management software. It also monitors allocation or usage of resources associated to the cloud service instances. Based on the allocation or usage of resources, it generates a chargeback report. The chargeback report is visible to consumers and provides transparency between consumers and providers. EMC E10-001 Student Resource Guide. Module 13: Cloud Computing