

ISLEVER

E10-001

Information Storage and Management Exam

Version 2

DEMO

<https://www.islever.com/e10-001.html>

<https://www.islever.com/emc.html>

For the most up-to-date exam questions and materials, we recommend visiting our website, where you can access the latest content and resources.

Topic 1, Volume A

QUESTION NO: 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

QUESTION NO: 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:

$$239 \text{ domains} \times 256 \text{ areas} \times 256 \text{ ports} = 15,663,104$$

- 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

QUESTION NO: 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?

- A. Source devices must be healthy
- B. Save location size must be larger than the size of all source devices
- C. Save location size must be equal to the size of all source devices
- D. 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