

# EMC Virtual LUN Technology

## *A Detailed Review*

---

### **Abstract**

This white paper describes the Virtual LUN capabilities available with EMC® CLARiON® storage systems beginning with FLARE® Release 16. It provides details on the terminology, features, benefits, use cases, concepts, storage system operations, and management of Virtual LUNs.

September 2006

---

---

Copyright © 2006 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Part Number H1529.1

---

## Table of Contents

<b>Executive summary .....</b>	<b>4</b>
<b>Introduction .....</b>	<b>4</b>
Audience .....	5
Terminology .....	5
Virtual LUN use cases .....	6
Reacting to changing information value (ILM).....	6
Load balancing/improving performance .....	6
Nondisruptive system consolidation/reconfiguration .....	6
<b>Virtual LUN concepts.....</b>	<b>6</b>
Migration components.....	7
Migration implementation.....	7
Initiating the Virtual LUN migration process.....	7
Selecting the destination LUN.....	8
Start migration .....	9
<b>Virtual LUN rules, facts, limitations.....</b>	<b>10</b>
Migration rules .....	10
Migration facts.....	10
Virtual LUN limitations.....	11
<b>Virtual LUN restrictions .....</b>	<b>11</b>
Source LUN restrictions .....	11
LUN or MetaLUN restrictions .....	11
Layered application restrictions.....	11
Destination LUN restrictions .....	12
<b>Virtual LUN failure/recovery.....</b>	<b>12</b>
Trespass scenario.....	13
SP panic scenario .....	13
SP failure, power loss, and double-fault scenarios.....	13
<b>Virtual LUN management .....</b>	<b>13</b>
Navisphere Manager.....	13
Virtual LUN properties .....	13
Virtual LUN summary .....	14
Navisphere CLI .....	15
<b>Conclusion .....</b>	<b>15</b>

---

## Executive summary

Virtual LUN technology is a feature that allows a user to change the characteristics of an existing LUN (RAID type, size, and so on) without disruption to host applications. Within a single storage system, Virtual LUNs move data from a source LUN to a destination LUN of the same or larger size with the changed characteristics that the user selected.

Virtual LUNs can bring value to your business in a number of ways:

- Enabling Information Lifecycle Management (ILM) within the array
- Improving performance
- Nondisruptive reconfiguration

Virtual LUNs enable users to easily move data within a storage array without downtime, allowing the creation and execution of an effective ILM strategy. Users can quickly migrate data to different storage tiers within the array as its value changes over time, enabling the more efficient use of hardware resources and maximizing their return on investment. Virtual LUNs also have the capability to improve array performance by migrating data from overutilized to underutilized disk spindles. Many other migration techniques do not allow the characteristics of the LUN to be changed; however, Virtual LUNs allow those characteristics to be changed to help improve performance. Virtual LUNs also enable easier system reconfiguration; for example, data can be moved away from disk-array enclosures as the system is upgraded without disruption. This can all be accomplished while maintaining service levels and reducing host overhead.

Virtual LUN technology is available starting with EMC® CLARiiON® Release 16. Virtual LUNs leverage FLARE®—CLARiiON's existing operating system—for data integrity and RAID protection features. The CLARiiON layered architecture allows advanced replication applications—such as EMC SnapView™, MirrorView™, and SAN Copy™—to work with Virtual LUNs. The new functions are integrated into EMC Navisphere® Manager GUI and CLI. The driver that facilitates Virtual LUN operations is packaged with the FLARE Operating Environment, so existing customers can leverage the new feature after a nondisruptive software upgrade.

## Introduction

The Virtual LUN feature may be applied on a per-LUN basis to migrate any public LUN (any LUN that is host-accessible) to any other public LUN or metaLUN. The LUN or metaLUN to which the source LUN or metaLUN has been migrated can have different LUN characteristics or even be of different disk types. The changeable characteristics of the source LUN are as follows:


- **RAID geometry**
  - RAID type (RAID 0, 1, 3, 5, 1/0, individual disk)
  - Number of disk spindles
  - Alignment offset
  - Stripe element size
  - LUN to metaLUN
  - MetaLUN to LUN
  - LUN size
- **Drive type**
  - Migrate from Fibre Channel to ATA or low-cost Fibre Channel
  - Migrate ATA or low-cost Fibre Channel to Fibre Channel

In any case, Virtual LUNs can enhance performance or increase disk utilization by allowing the user to change any of the aforementioned characteristics while their production volume remains online.

---

Leveraging FLARE offers customers peace of mind with respect to data integrity. CLARiiON has over 12 years of R&D investment in the FLARE Operating Environment, which provides industry-leading performance and data-integrity features. The architecture as a whole has well-defined software interfaces that allow layered application drivers to be added nondisruptively to provide additional value-added services, such as local and remote copies to end users. Virtual LUNs follow the same layered application paradigm. Virtual LUN functions are implemented as an additional driver that leverages the already-proven platform.

Figure 1 shows where the migration driver fits into the driver stack. Although it is a discrete driver, the Virtual LUN driver is part of FLARE so customers do not have to buy any additional licenses to take advantage of the Virtual LUN feature.



**Figure 1. CLARiiON driver stack with CLARiiON Virtual LUN driver**

The distinct advantages to implementing features in this manner are:

- Adding the feature is nondisruptive to the customer's hardware and configuration, as well as to any ongoing host I/O. You add the Virtual LUN feature by simply upgrading to FLARE 16 or higher.
- FLARE, a proven operating environment, handles fault tolerance and data integrity. Every LUN that is part of a Virtual LUN session is managed independently by FLARE.
- The Virtual LUN driver sits below the layered application drivers so, given certain restrictions, the layered applications can still be utilized during the migration process.

## **Audience**

This white paper is intended for EMC prospects, EMC customers, system engineers, EMC partners, members of the EMC and partners professional services community, and anyone interested in learning more about CLARiiON Virtual LUNs.

## **Terminology**

The following is a summary of terms that are used throughout this white paper. Both old and new terms and concepts are included. The definitions of new terms are further defined as they are used in context throughout the paper.

**Checkpoint** — A value that is used to keep track of the progress of synchronization. The checkpoint is periodically updated and persistently stored. In case of failures during synchronization, it can be used to continue the migration operation without having to restart a full synchronization from the beginning.

**Controlling storage processor (SP)** — The storage processor that owns both source and destination LUNs during the migration process.

---

**Destination LUN** — The LUN that is the target of the Virtual LUN migration process.

**LUN** — A logical partition of a RAID group. The LUN is the basic logical unit managed by FLARE and serves as the building block for Virtual LUN components.

**LUN ID** — The local identifier for any LUN.

**Virtual LUN** — The process of migrating a public LUN/metaLUN to another public LUN/metaLUN that has different characteristics.

**MetaLUN** — A composite LUN made up of multiple LUNs.

**Nick name** — The simple plain-text naming convention for any LUN.

**Source LUN** — The original LUN/metaLUN that the Virtual LUN driver acts upon. The source LUN remains online during the entire migration process.

**Synchronize (sync)** — The process of updating data contents of the destination LUN with those of the source LUN.

**WWN** — The World Wide Name, which is the global identifier for any LUN.

## ***Virtual LUN use cases***

The following section discusses the scenarios as to where Virtual LUNs can be utilized.

### **Reacting to changing information value (ILM)**

Virtual LUNs allow users to move data to cheaper or better performing arrays as the value of that data changes over time. For example, historical data within a database, which is not accessed very often, can be migrated from Fibre Channel to ATA drives. This “archiving” within the array frees up better performing FC drives for the information that is most valuable. This functionality makes CLARiiON a perfect fit for data warehousing.<sup>1</sup>

### **Load balancing/improving performance**

If a particular source LUN on a given set of disk spindles becomes hot with I/O activity, Virtual LUNs allow that LUN to be moved to faster or less-utilized spindles. Virtual LUNs also allow the characteristics of the LUN to be changed, thereby enhancing the performance of the source LUN. Similarly, if a particular source LUN on a set of high-performance disk spindles becomes cool with I/O activity, Virtual LUNs allow that LUN to be moved to more cost-effective spindles, freeing up the higher performance spindles for a more productive use.

### **Nondisruptive system consolidation/reconfiguration**

In many configurations, multiple LUNs are spread across a varying number of disk-array enclosures (DAEs). Virtual LUNs allow for the consolidation of those LUNs to a particular DAE, allowing multiple DAEs to be balanced across back-end loops or simply freeing up a particular disk enclosure to be used elsewhere. Virtual LUNs can also be used to migrate data while a piece of the array is being reconfigured, to ensure minimal disruption to the host.

## **Virtual LUN concepts**

The following sections discuss the major concepts associated with Virtual LUN. The “Migration components” section discusses exactly what is necessary in order for a migration to occur. Once an

---

<sup>1</sup> See EMC white paper *SQL Server Data Warehouse Deployments with EMC CLARiiON Storage Systems* on EMC Powerlink

understanding of the components is complete, the question of how a Virtual LUN session is implemented is looked upon in detail.

## Migration components

The Virtual LUN process consists of a source LUN as well as a destination LUN. The source LUN can be any public LUN or metaLUN that is currently not in a state of expanding. Once the source LUN has been selected, the user is presented with a list of all the public LUNs that are available on the storage system. The public LUNs available to be a destination LUN can belong to either SP within the storage system. Figure 2 shows how the Navisphere UI depicts the list of available LUNs that can be chosen to be the destination LUN for the Virtual LUN process. In this example, the source LUN is a public LUN belonging to SP A; the chosen destination LUN is a public metaLUN belonging to SP B.

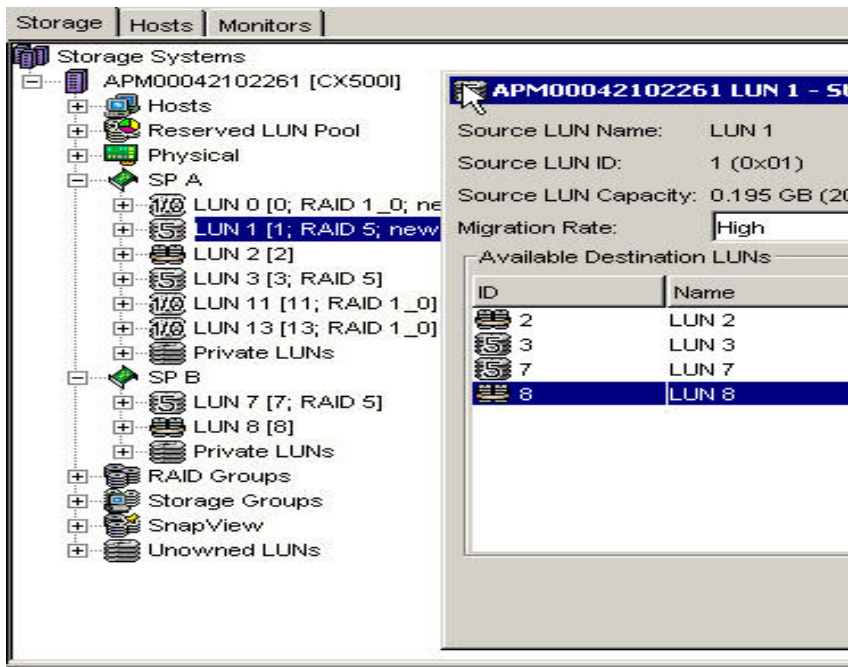


Figure 2. CLARiiON Virtual LUN process as represented in the Navisphere UI

## Migration implementation

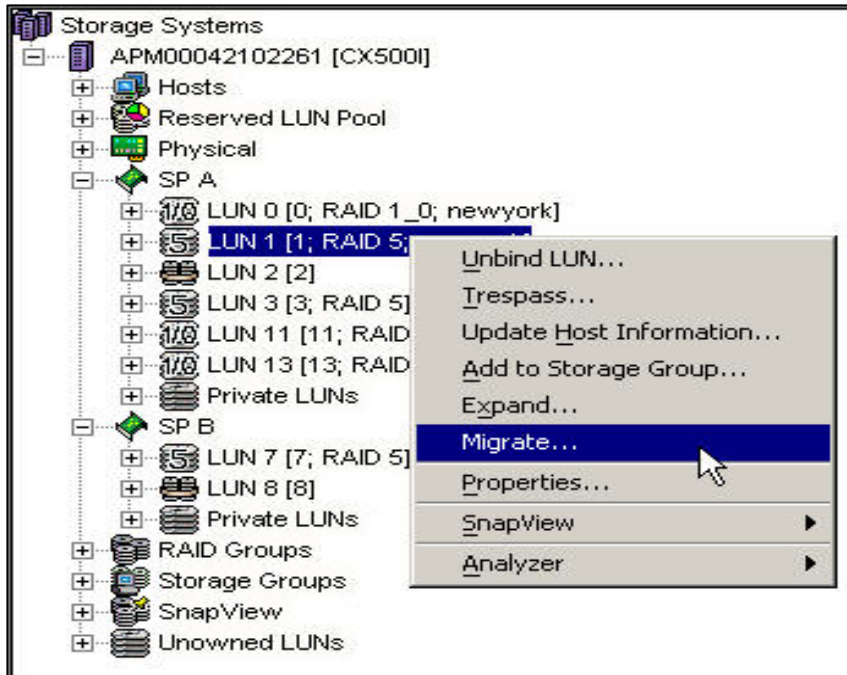
Until now, SnapView clones and SAN Copy were the only storage-system-based methods for intra-storage-system LUN migration. To take advantage of these capabilities, the user would need to have licenses for these layered applications. The major advantage to the addition of the Virtual LUN driver is that it not only is available with FLARE 16, but that it also allows the source LUN to remain online while the migration is taking place without any interruption to host I/O. It eliminates the need to unmount the original source and then mount the destination LUN. Once the migration completes, the destination LUN assumes the source LUN's *nickname*, WWN, and LUN ID.

The following sections describe how to use Virtual LUNs as well as how the Virtual LUN driver handles the migrating processes.

## Initiating the Virtual LUN migration process

Figure 3 depicts the selection of the source LUN—the public LUN that the user would like to migrate—using the object-driven menus of the Navisphere UI. In this case, the user is selecting a public LUN

(LUN 1) to be migrated to any other public LUN within the storage system.

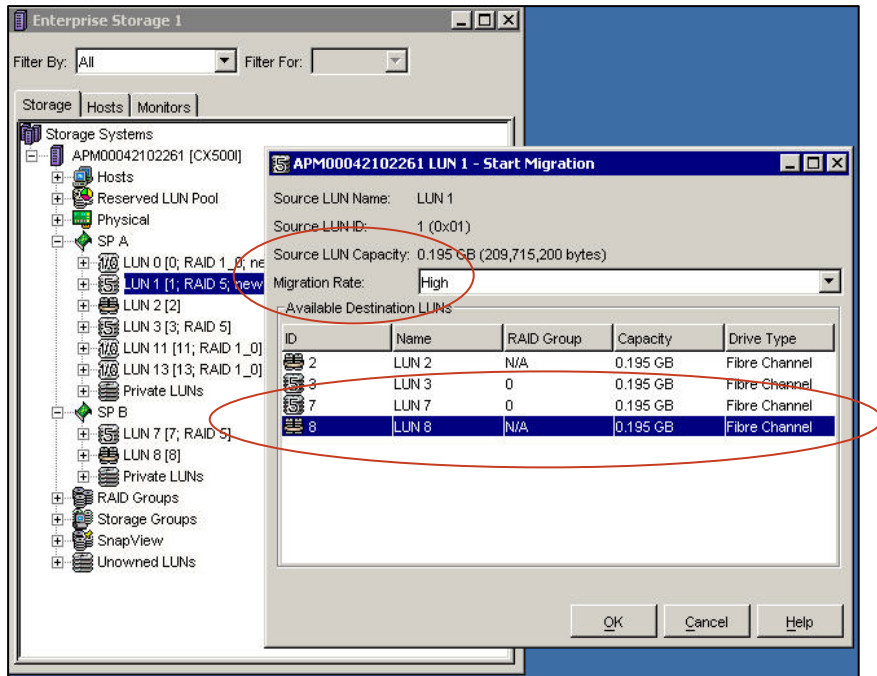


**Figure 3. Initiating the Virtual LUN process**

The source LUN remains online during the entire migration process so that the host application can continue sending I/O traffic without interruption. Once the user selects **Migrate**, the process flow begins and the selected source LUN eventually becomes a public LUN of any type, with any number of varying characteristics.

### Selecting the destination LUN

Figure 4 depicts the object menu where the user selects the targeted destination LUN from the **Available Destination LUN** field, as well as the migration rate, which is the rate at which the migration takes place. The **Migration Rate** menu offers four levels of service: **ASAP**, **High**, **Medium**, and **Low**. The user-selected migration rate depends on the amount of SP utilization that the user would like to have allocated for the migration.



**Figure 4. Selecting the destination LUN and migration rate**

After selecting the migration rate, the user selects the destination LUN, which is the eventual target for migration. In this case, the user is targeting a metaLUN—LUN 8—and, as shown, the LUN’s SP ownership makes no difference since all public LUNs within the storage system are available.

### Start migration

Once the destination LUN has been selected, the Virtual LUN migration should successfully start, as shown in Figure 5. When the session begins, ownership of the destination LUN is changed to that of the source LUN so that the migration process is managed by a controlling SP.

During the migration, the source LUN remains online to receive host I/O. The Virtual LUN driver handles the synchronization of source and destination LUNs, as well as marking the checkpoint file in a set periodic increment. Synchronization of the source and destination is a block-for-block copy, so if a block on the source LUN has been marked as bad, then that particular block will also be marked as bad on the destination.

If a host write is received by the source LUN to a region that has been synchronized, then that write is automatically replicated to the destination LUN. However, if a host write is received by the source LUN and is written to a block of the LUN that has yet to be synchronized, then that write is simply made to the source LUN as that region of the LUN waits for synchronization to occur. This mechanism of the migration driver allows the source LUN to remain protected should an error occur or the session fail.

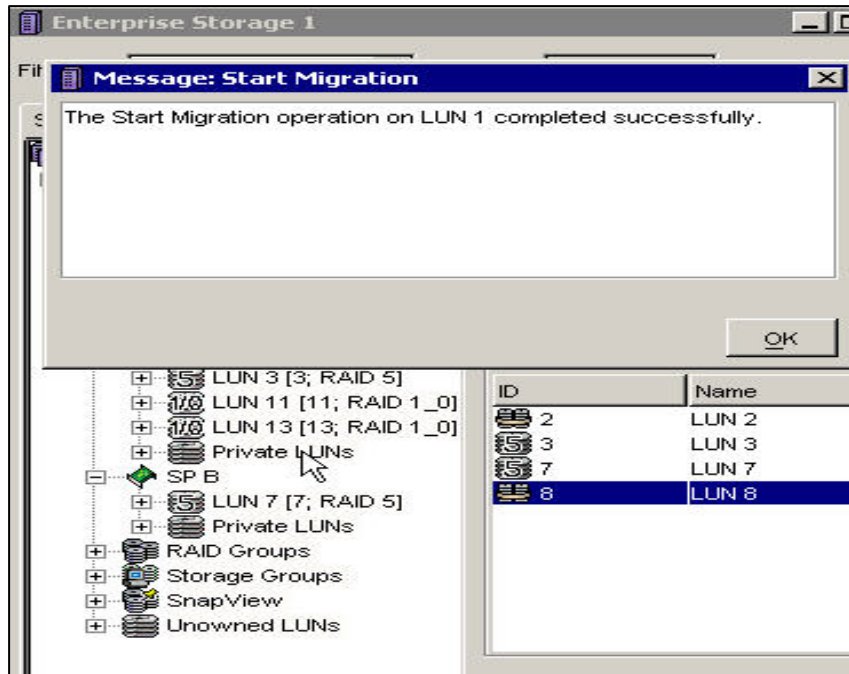


Figure 5. Successful start of the Virtual LUN process

## Virtual LUN rules, facts, limitations

The following sections discuss all the necessary rules, facts, and limitations associated with Virtual LUNs. Each section provides insight into what can be done with Virtual LUNs.

### Migration rules

- Source and destination LUNs must:
  - Reside on the same storage system
  - Have the same controlling SP during migration
- Source and destination LUN can be in the same or different RAID groups.
- Any public LUN or metaLUN of size X can migrate to any LUN or metaLUN of size X or larger.
- Any RAID type can migrate to any RAID type.
- Any disk type (Fibre Channel or ATA) can migrate to any disk type (Fibre Channel or ATA).

### Migration facts

- Source LUN can be in a user-created storage group.
- Source LUN is unbound upon migration completion.
- Source and destination LUN SP ownership can be different.
  - Destination LUN will trespass to source LUN SP upon start of migration.
- Each migration has a persistent checkpoint file to keep track of the migration synchronization process.
- Destination LUN, upon completion, assumes the source LUN's name, WWN, and LUN ID.

---

## Virtual LUN limitations

The number of active and queued jobs that the Virtual LUN driver can handle depends on the CLARiiON CX Series storage-system model.

Table 1 shows the limits for each of the CX Series storage systems.

**Table 1. CX Series storage system limits**

Storage system	Active jobs	Queued jobs
CX700	12	1012
CX500	8	504
CX300	4	252
CX600	4	508
CX400	4	252
CX200	4	124

## Virtual LUN restrictions

The following sections discuss both the source and destination LUN restrictions for Virtual LUN migration. The first section takes a look at what cannot be done from the perspective of the source LUN by considering not only the LUN restrictions, but also the restrictions with regards to layered applications. The last section takes a look at what cannot be done from the perspective of the destination LUN.

### Source LUN restrictions

The restrictions for Virtual LUNs with regard to the source LUN must be considered from both the LUN perspective as well as the layered application perspective. These restrictions are as follows.

#### LUN or MetaLUN restrictions

For a LUN to be an eligible source LUN, the following restrictions apply:

- The LUN must be a public LUN or metaLUN.
- The LUN must *not* be in a transitioning state.
- The LUN must *not* be expanding.

#### Layered application restrictions

Each layered application has certain restrictions before the LUN being controlled by the particular layered application can be migrated.

##### SnapView

- A snapshot LUN cannot be migrated.
- A reserved LUN cannot be migrated.
- If migrating to the same-size LUN, the source LUN may be migrated at any time.
- If migrating to a larger-size LUN, the user must stop all sessions and destroy any snapshots for this LUN.

##### BCV/clones

- A clone private LUN may not be migrated.
- If migrating to the same-size LUN, the source or clone may be migrated at any time.

- 
- If migrating to a larger-size LUN:
    - If the LUN is a clone, then remove the clone from clone group.
    - If the LUN is the clone's source, then remove all clones of that LUN from its clone group and destroy the clone group.

#### **MirrorView/S**

- A write-intent LUN may not be migrated.
- If migrating to the same-size LUN, the primary or secondary LUN may be migrated as long as the mirror is not in the synchronizing state.
- If migrating to a larger-size LUN:
  - If the LUN is the secondary image, then remove it from the mirrored relationship.
  - If the LUN is the primary image, then remove all secondary images from the mirrored relationship and destroy the mirror.

#### **MirrorView/A**

- If migrating to the same-size LUN, the primary or secondary LUN may be migrated at any time.
- If migrating to a larger-size LUN, the primary or secondary LUN may be migrated once the MirrorView/A session has been *stopped*.

#### **SAN Copy**

- If migrating to the same-size LUN:
  - The local LUN may be migrated unless it is participating in an active SAN Copy session.
  - The remote LUN may be migrated at any time.
- If migrating to a larger-size LUN:
  - All SAN Copy sessions must be stopped and all the copy descriptors for the LUN must be destroyed.
  - The remote LUN may be migrated at any time.

#### **Incremental SAN Copy**

- If migrating to the same-size LUN, the local or remote LUNs may be migrated at any time.
- If migrating to a larger-size LUN, the local or remote LUNs may be migrated once the incremental SAN Copy session has been stopped.

### ***Destination LUN restrictions***

For a LUN to be an eligible destination LUN, the following restrictions apply:

- The LUN must be a public LUN or metaLUN.
- The LUN must be the same size or larger.
- The LUN must *not* be in a storage group.
- The LUN must *not* be in use by any of the layered drivers within its driver stack.
- The LUN must *not* be expanding.

### **Virtual LUN failure/recovery**

While the source LUN is undergoing Virtual LUN migration, a checkpoint file is periodically written to in order to keep track of how far along the source LUN is in the migration process. The checkpoint file is persistently stored within the vault area of the first five drives in the storage system and under certain circumstances allows the migration to restart at the last known good checkpoint. The following describes the potential failure scenarios that could occur and the resulting migration recovery that would take place.

---

## ***Trespass scenario***

If the source LUN is undergoing Virtual LUN migration and a trespass occurs, the Virtual LUN driver writes to the checkpoint file indicating a point of clean recovery for the migration. During the trespass, both the source and destination LUNs—which are already being controlled by the controlling SP—trespass together. Once the trespass is complete, the migration restarts at the last checkpoint.

## ***SP panic scenario***

If the source LUN is undergoing a Virtual LUN migration and an SP panic occurs, the Virtual LUN driver is unable to update the checkpoint file with a point of clean recovery for the migration. Once the SP is back online, the migration then restarts from the beginning because the checkpoint file has been marked dirty.

## ***SP failure, power loss, and double-fault scenarios***

If the source LUN is undergoing Virtual LUN migration and the SP fails, a power loss or a double-fault occurs. The migration then restarts from the beginning once the issue is resolved.

## **Virtual LUN management**

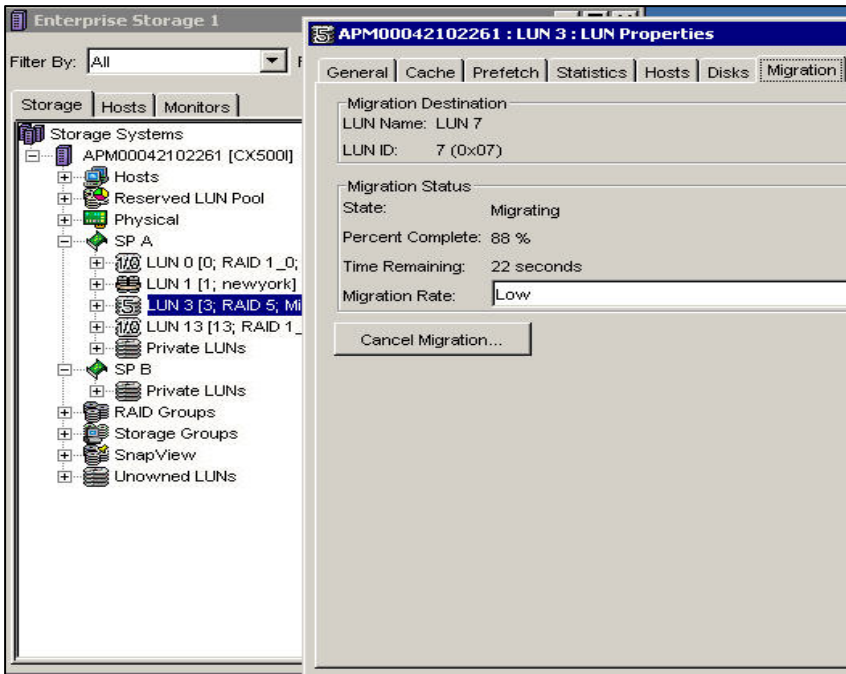
The user can easily manage Virtual LUNs through either the Navisphere Manager UI or the Navisphere CLI.

### ***Navisphere Manager***

The user interface within Navisphere Manager has several additional features due to the addition of Virtual LUNs. These features are as described next.

#### **Virtual LUN properties**

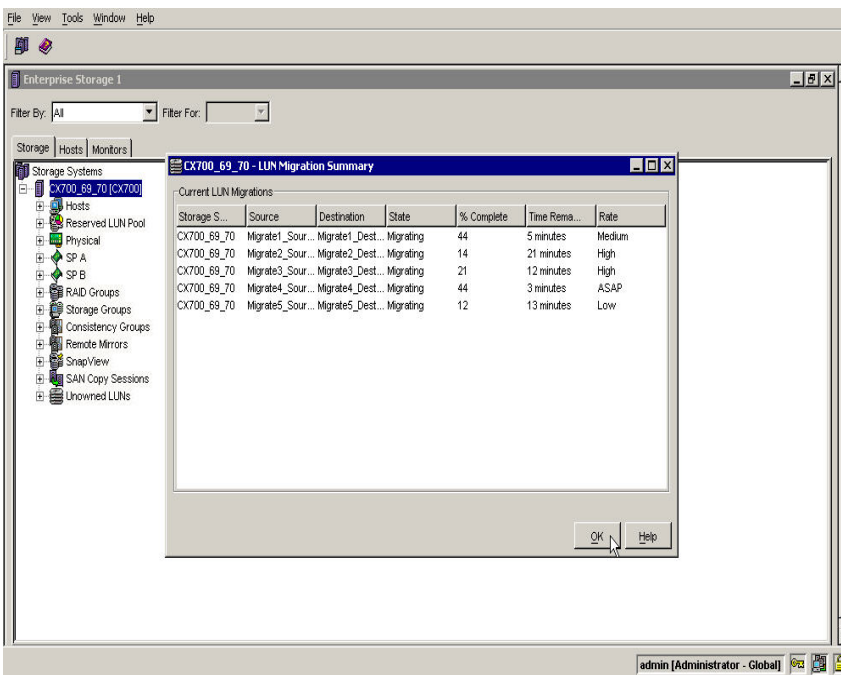
When a source LUN is undergoing the Virtual LUN migration process, an additional tab is added to the **LUN Properties** dialog box. As Figure 6 shows, the user can change the **Migration Rate** or click **Cancel Migration** to cancel the process altogether. If the Virtual LUN session is canceled, then the destination LUN is destroyed in order to maintain data security.



**Figure 6. Properties of a Virtual LUN**

### Virtual LUN summary

Figure 7 shows the **Virtual LUN Migration Summary** dialog box, which displays all the migration currently in process as well as in queue. From this summary, you can see how much longer a particular session may have left to complete. You can also use the summary to determine if it is necessary to change a session's migration rate.



**Figure 7. Virtual LUN session summary**

---

## Navisphere CLI

A set of Java-based CLI commands has been added to Navisphere CLI to manage Virtual LUN operations. The commands are executed by the `navicli.jar` file that resides in the Navisphere CLI directory. Since the commands use Java, the Java Runtime Environment must be present on the host executing the commands. Using the security mechanisms of Navisphere 6, the `navicli.jar` file communicates directly with the Management Server package. The root command **migrate** is used to control the Virtual LUN operations. There are several flags for starting, canceling, modifying, and information gathering. The commands are:

```
migrate
  -start
    -source LUN ID/WWN -dest LUN ID/WWN
    -rate low|medium|high|asap

  -cancel
    -source LUN ID/WWN [-o]

  -modify
    -source LUN ID/WWN -rate low|medium|high|asap [-o]

  -list
    [-source LUN ID/WWN] [-destination] [-rate] [-state]
    [-percentcomplete] [-timeremaining]
```

## Conclusion

Virtual LUN technology enables users to change the characteristics of their information without any disruption to host applications. Virtual LUNs add value by migrating information freely throughout the array, empowering users to get a better return on their storage investment while still maintaining service levels on their most critical applications. Virtual LUN technology is a feature unique to EMC CLARiiON and requires no additional software licenses.

Virtual LUN migration is performed with the source LUN remaining online, and with no interruption to host I/O. The ability to migrate data anywhere within the storage system allows the user to realize better performance by migrating data away from overutilized drives to faster, less-utilized disk drives. Virtual LUNs can be a method of supporting an ILM strategy because data can be migrated to less-expensive hardware, effectively “archiving” within the array. Additionally, Virtual LUNs can change the characteristics of the source LUN during migration, which can also improve performance. Furthermore, the technology can be used when performing system reconfiguration or consolidation, by allowing the migration of data spread across the array to one DAE, either as a consolidation effort or to minimize disruption to the host while you add or change capacity.

Virtual LUN technology is a feature that is managed through Navisphere and continues to leverage FLARE, the proven CLARiiON operating system, to provide online data migration capabilities in CX Series storage systems (existing users can get the new functionality by simply performing a nondisruptive upgrade to FLARE 16 or higher). No additional hardware or appliance is necessary to utilize Virtual LUNs. In addition, all of the CLARiiON layered applications for local and remote replication work with Virtual LUNs, and taking advantage of the feature will require little or no change to your existing storage infrastructure.