

# EMC Data Compression

## *A Detailed Review*

---

### **Abstract**

EMC® Data Compression is a storage efficiency feature for CLARiiON® that allows information to be stored in a smaller footprint than it normally would be. It allows users to maximize their total cost of ownership (TCO) benefits for relatively inactive data that requires the five 9s of protection offered by CLARiiON.

September 2010

---

---

Copyright © 2010 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](http://EMC.com)

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

Part Number h8045.1

---

## Table of Contents

<b>Executive summary .....</b>	<b>4</b>
<b>Introduction .....</b>	<b>4</b>
Audience .....	4
<b>Introducing EMC Data Compression.....</b>	<b>4</b>
<b>Data reduction technologies.....</b>	<b>6</b>
Best use of Data Compression .....	6
<b>How compression works.....</b>	<b>7</b>
Compress operations.....	8
Decompress operations.....	10
<b>Managing compression.....</b>	<b>11</b>
<b>Limits and interoperability .....</b>	<b>14</b>
Limits.....	14
Interoperability .....	15
Special cases for LUN Migration and SAN Copy.....	16
<b>Resource consumption and performance .....</b>	<b>16</b>
<b>Conclusion .....</b>	<b>19</b>
<b>References .....</b>	<b>19</b>
<b>Appendix A: NTFS file system conditioning with SDelete .....</b>	<b>20</b>
<b>Appendix B: Compression states.....</b>	<b>21</b>

---

## Executive summary

EMC Data Compression is an EMC® CLARiiON® efficiency feature that allows users to store information using as little as half the storage capacity that they use today. All compression and decompression processes are handled by CLARiiON, so no server cycles are consumed in the process, and no additional server software is required. With average compression rates as high as 2:1, users can reap significant cost reductions in initial capital expenditure and ongoing operational costs.

There are many data types that can benefit from compression, such as typical end-user office documents; home directories; archives for virtual-machine files; and static copies of databases or messaging-system components. Deploying compression at the storage-system level provides a consistent deployment across several data types.

EMC Data Compression is available for CLARiiON customers running release 30 of the FLARE® Operating Environment.

## Introduction

This white paper discusses the implementation, management, and best practices of the Data Compression feature for CLARiiON.

## Audience

This white paper is intended for IT planners, storage architects, administrators, and others involved in evaluating, managing, operating, or designing CLARiiON storage systems. Readers should be familiar with the concepts and terminology of CLARiiON Virtual Provisioning™ as described in the *EMC CLARiiON Virtual Provisioning* white paper on EMC Powerlink®.

## Introducing EMC Data Compression

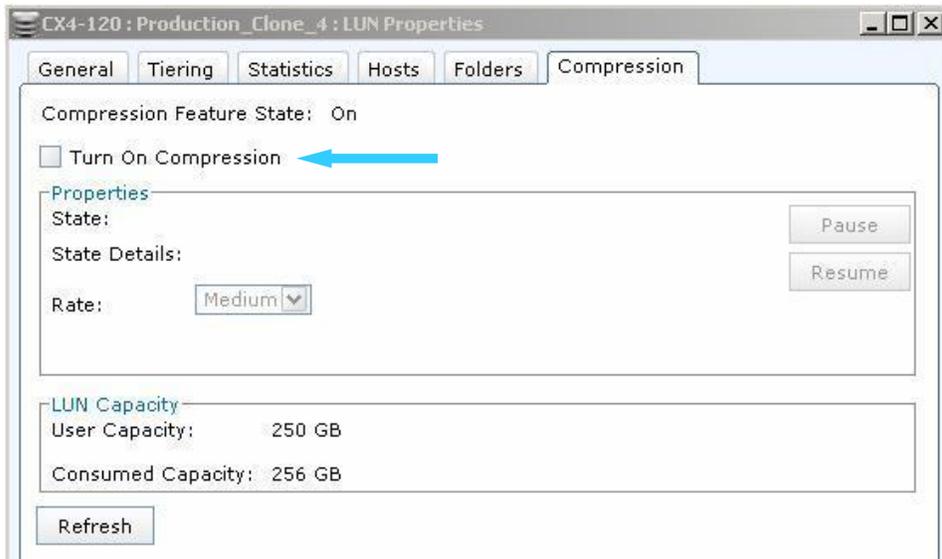
CLARiiON has a long history of storage efficiency innovations. CLARiiON was the first midrange storage system to offer ATA drive technology, which offered increased storage density compared to Fibre Channel disk drives. CLARiiON's disk-drive Spin Down feature, available for RAID groups on SATA drives, minimizes power consumption for inactive drives. CLARiiON was the first midrange storage system to offer Flash drives, which, in addition to performance benefits, use the least power per I/O. Thin LUNs can be used to allocate physical capacity on demand in very small increments. Sub-LUN Fully Automated Storage Tiering (FAST) can automatically move your least active data to less-expensive SATA drives, while leaving the most active elements on Fibre Channel and Flash drives. CLARiiON's Data Compression feature is yet another option to further optimize storage capacity allocation.

EMC Data Compression is a licensed software feature that is tightly integrated with CLARiiON Virtual Provisioning technology. CLARiiON Virtual Provisioning provides ease-of-use features such as pools, thick LUNs, thin LUNs, and FAST sub-LUN tiering. These features greatly reduce storage system management complexity by employing EMC best practices and automating many mundane tasks.

Deciding which service level is appropriate for a dataset is typically a tradeoff between storage costs and performance. Data compression can offer significant reductions in storage costs and is ideal for datasets that have low performance requirements, but still require the protection of CLARiiON five 9s availability. It is also flexible; since compression is implemented at the LUN level, it can be applied to any data type that resides on the array.

It is very easy to manage the Data Compression feature. As you can see in Figure 1 and Figure 2, compression is a LUN attribute that you simply enable or disable on each LUN. When enabled, data on disk is compressed as a background task. When the host writes new data to a LUN, it is stored on the LUN until system-defined thresholds are met, which automatically triggers compression of new data asynchronously. When the host reads compressed data, the data is decompressed in memory, but remains

compressed on disk. Therefore, once enabled, no user intervention is required; the system automatically manages processing of new data in the background.



**Figure 1. LUN Properties dialog box – Compression tab**

Name	ID	State	Thin	Compression	Storage Pool	Tiering Policy
Dastore_1	14	Ready	Off	Off	Pool_1_FC_SATA	Auto-Tier
Development_1	4	Ready	On	Off	Pool_2_FC	Auto-Tier
Production_1	0	Ready	Off	Off	RAID Group 0	
Production_Clone_1	7	Ready	On	On	Pool_1_FC_SATA	Auto-Tier
Test_Data_1	12	Ready	On	On	Pool_1_FC_SATA	Auto-Tier

**Figure 2. Compression LUN attribute in a LUN table**

Compression can be enabled on any LUN type, including RAID group-based LUNs (traditional LUNs and metaLUNs) and pool-based LUNs (thick LUNs and thin LUNs). When a LUN is compressed, if it is not already a thin LUN it becomes a thin LUN. Any RAID group LUN is automatically migrated to the user’s pool of choice as part of the initial compression. Pool-based LUNs remain in the pool in which they currently reside when they are compressed. Thin LUNs are a key element of Data Compression, because as data is compressed, thin-LUN blocks are freed and eventually returned to the pool for use by other LUNs.

Compression is compatible with all CLARiiON replication features including LUN Migration, SnapView™, MirrorView™/A and MirrorView/S, and SAN Copy™. As a result, there are many use cases where these products can create a compressed copy of a dataset. There are also many use cases where compressed LUNs can be rapidly repurposed with these tools.

Compression is complementary to other advanced features such as sub-LUN FAST and FAST Cache. Sub-LUN FAST helps facilitate placement of compressed LUNs on lower-tier drives. LUNs with lower I/O activity are usually placed on lower-tier drives because of FAST’s auto-tiering policy. Users can also explicitly mandate that LUNs be placed on the lowest-tier drives. Under normal conditions, FAST Cache may not provide much benefit to compressed LUNs because of their low I/O profile. However, if compressed data becomes active, FAST Cache may help bridge the gap while data is being decompressed

---

or migrated. Using these features with compression is discussed further in the “Interoperability” section of this paper.

## Data reduction technologies

Compression is part of a larger group of data reduction technologies that also include block-level deduplication and file-level deduplication. Certain data reduction technologies are extremely effective when applied to particular data types, while others provide modest benefits for almost any data type. When choosing which data reduction to use, the most important considerations are the type of data and the requirements for accessing the data. Following are overviews of how these technologies are applied in various EMC products.

Data deduplication products, such the Data Domain<sup>®</sup> and Avamar<sup>®</sup> families, can process massive amounts of backup and archive datasets, and over time reduce capacity requirements by as much as 10x to 30x. For example, datasets that contain several full backups have a lot of duplicate data. A single file may be backed up several times but only have a few bits change between versions. These systems examine an entire dataset at the bit level, and store unique blocks of data only once. A single file may have version 1.0, 1.1, 1.2, and so on. If there are small incremental changes in each version, the combined capacity to store all versions may not be much more than the capacity required for version 1.0 alone.

File-level deduplication, or file-level single instancing, can reduce capacity requirements by 3x and sometimes more. Content Addressed Storage (CAS) systems, like EMC Centera<sup>®</sup>, perform this type of deduplication along with other optimizations. In these cases, each file is given a unique identifier and stored only once. When datasets such as email archives are stored, tremendous benefits can be realized when attachments that are part of several email messages are only stored once on the system. However, version 1.0 and version 1.1 of a file may be stored as two unique objects. EMC Centera conducts its deduplication across all of the data in the Centera cluster. EMC Centera is designed primarily for static content, because another feature of this technology is data authenticity for compliance purposes. This technology can be used to prove that data has not changed over long periods of time. As a result, Centera is ideal for data types that may be involved in litigation, such as archived email or medical records.

EMC Celerra<sup>®</sup> implements a combination of technologies. It performs single-instance, file-level deduplication, so that multiple instances of the same file are only stored once. It also employs compression of those files. This can reduce data footprints for typical file share data as much as 2x. Celerra Data Deduplication is enabled and managed at the file-system level. It is designed to service a dynamic file sharing environment by offering flexible deduplication policies based on file type, size, and time of access.

Data Compression for CLARiiON processes all data residing in a LUN. Data is processed in fixed 64 KB increments. It encodes any recurring strings or patterns within the data currently being processed and then moves on to the next 64 KB increment. There is no notion of single instancing either at the file level (like EMC Centera or Celerra) or across blocks (like the Data Domain and Avamar). Each 64 KB increment of data is processed independently of other data on the LUN. This process can reduce data footprints by as much as 2x, while using modest CPU and memory resources.

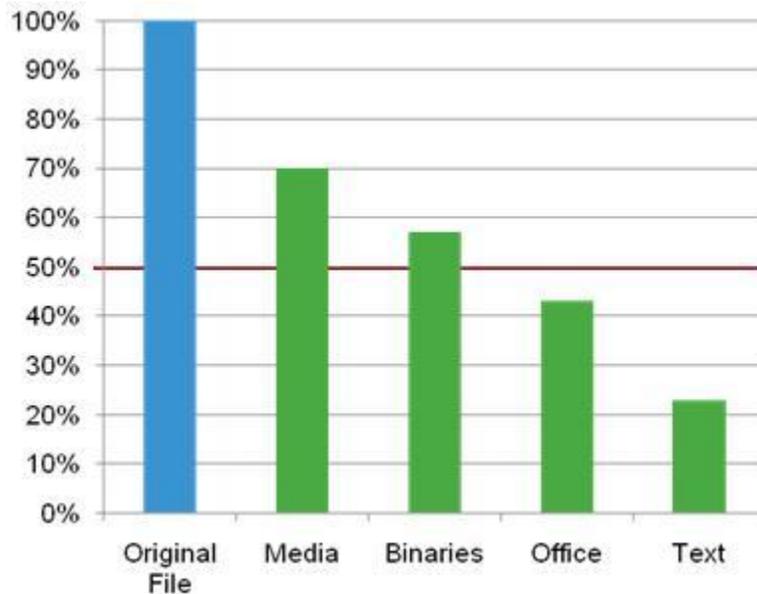
EMC Data Compression uses the same compression libraries that RecoverPoint and Celerra Data Deduplication use. Through these products, this library has been successfully deployed over a large variety of datasets.

### ***Best use of Data Compression***

There are many places where compression can be implemented in an environment. Some applications and databases offer native compression. There are also many host-side compression utilities that can be used for archive and compression of individual files and directories, such as tar, Winzip, gzip, and bzip2. These utilities typically require some form of manual execution or scripting and offer little to no control over the amount of resources dedicated to the task. If a user’s compression granularity requirement is at the LUN level, Data Compression offers an easy-to-implement solution that can be managed consistently across any data type. For file system data, Celerra Data Deduplication offers similar autonomy while also allowing users a high level of control at the individual file and directory level.

---

Typical results for initial compression of common data types are shown in Figure 3. Initial compression benefits include both reducing the data footprint and removing any unused space. This chart assumes 25 percent of space in the volume was unused before compression. The red line shows a rough estimate of savings across all data types of 50 percent (2:1 reduction).



**Figure 3. Initial compression rates of common file types**

Compression works best on data that contains whitespace and/or recurring patterns. Media files, like .pdf, .jpeg, and .mpeg, often include some form of compression. Applying any subsequent compression algorithm to these file types will not save additional space. Binary files, like executables, shared libraries, and font files as well as common office files, are moderately compressible. Text files, which can include HTML and XML documents, often have a higher degree of compressibility. Any data type that is encrypted is very random in nature by design and does not benefit from compression.

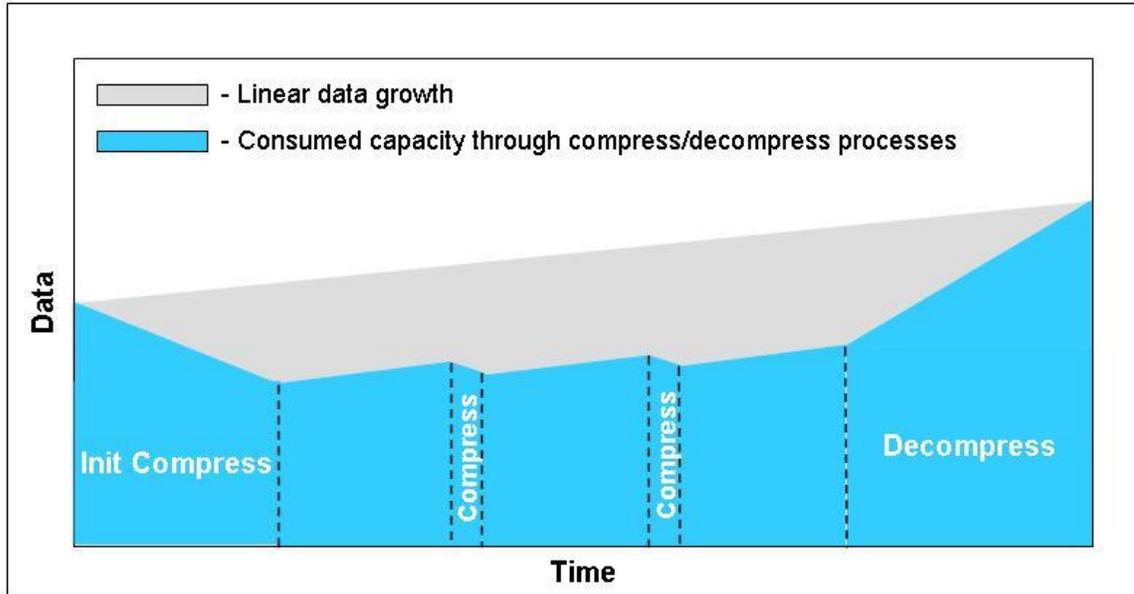
When compressing LUNs that host NTFS file systems, users should consider using the SDelete utility to zero blocks occupied by deleted files. These blocks are not zeroed by default, but doing so with the SDelete utility can greatly improve capacity savings from Data Compression. The SDelete utility is discussed in greater detail in “Appendix A: NTFS file system conditioning with SDelete.”

## How compression works

There are three phases of compression for LUNs:

1. Initial compression — This occurs when compression is first enabled for a LUN. The entire LUN is processed during this phase.
2. Compression of new data — As new data is written to the LUN, it is written uncompressed and then compressed asynchronously. This is an ongoing cycle for as long as compression is enabled on a LUN.
3. Decompression when compression is disabled — When compression is disabled on a LUN, the data on the LUN is decompressed.

The entire LUN is processed when the Data Compression feature is enabled or disabled on a LUN. Figure 4 shows linear data growth through two different scenarios. The gray (upper) area represents simple linear data growth over a period of time without compression. The blue (lower) area represents storage system consumed capacity for the same data profile through initial compression, new data compression, decompression.



**Figure 4. Consumed capacity of a LUN through the compress/decompress processes**

The initial compression and decompression phases take the longest. When new data is written to a compressed LUN, the consumed capacity of the compressed LUN increases. When the system-defined threshold for new data is reached for the LUN, the system automatically starts the compression process for the new data. The system-defined threshold is equal to 10 percent of the user capacity of the LUN or 10 GB, whichever comes first. New data must increase the LUN's consumed capacity by at least 1 GB before it is eligible for the compression process to run. This cycle of consumed capacity increasing with new data, and then being reduced with compression, continues indefinitely until either the user disables compression for the LUN or the LUN is migrated to a LUN without compression enabled.

The compression process is resilient to LUN trespass operations. If a LUN is trespassed due to faults in the host path, compression operations continue on the SP that takes ownership of the LUN. Compression activities of a trespassed LUN count toward the concurrent operation limit of the new SP. These limits are provided in the "Limits and interoperability" section of this paper.

Initial compression of thin LUNs, compression of new data, and decompression operations are checkpointed every few minutes. Checkpointing involves saving the current status of the operation to disk. Therefore, in the event of a system interruption, the processes that were running prior to the interruption can effectively restart at the correct place. Compression operations that utilize LUN Migration restart from the beginning if there is a system interruption.

Compression is tightly integrated with Virtual Provisioning and does not require private LUNs to work. Compression uses the existing metadata structure that is in place for thin LUNs. It has a modest memory budget, used for management purposes, that is shared with other features. However, because of the large number of objects supported, the combined memory budget is increased on the CX4-480 and CX4-960 when either compression or FAST is enabled. This may affect SP cache allocations set by the user. After installation, users should verify that SP cache memory is allocated to read and write cache as desired. If FAST is already enabled, no memory allocation changes are made when Data Compression is installed. The same is true if the feature is installed before FAST.

## **Compress operations**

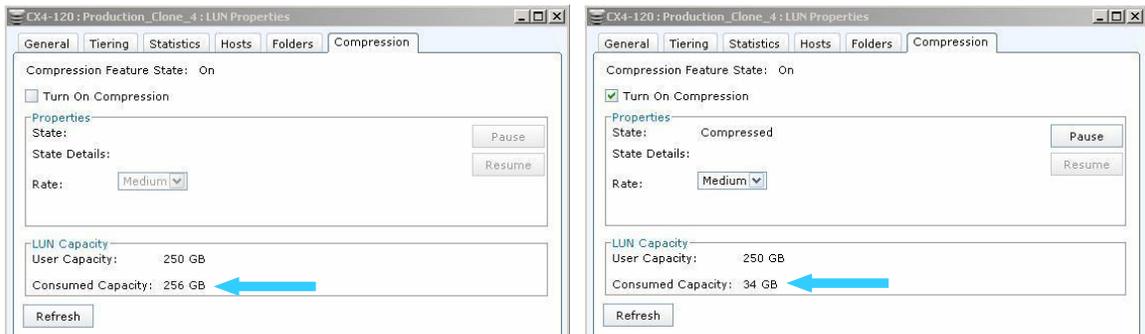
Although compression can be enabled on any LUN type, when the initial compression process completes, the LUN will be thin, because thin LUN technology allows the compression process to return space to the pool. For thin LUNs, 8 KB is the smallest unit of capacity allocation. As data is processed in 64 KB

increments, compressed data is written to the LUN if at least 8 KB of the 64 KB can be saved. If the resulting savings from compression is less than 8 KB, the data is written to the LUN uncompressed.

Capacity is allocated from a pool to a thin LUN in 1 GB slices. Over time, if enough 8 KB blocks are freed by compression, 1 GB slices can be returned to the pool for use by any LUN in the pool. This process starts once enough capacity has been saved to free a slice and may continue after the compression process has completed.

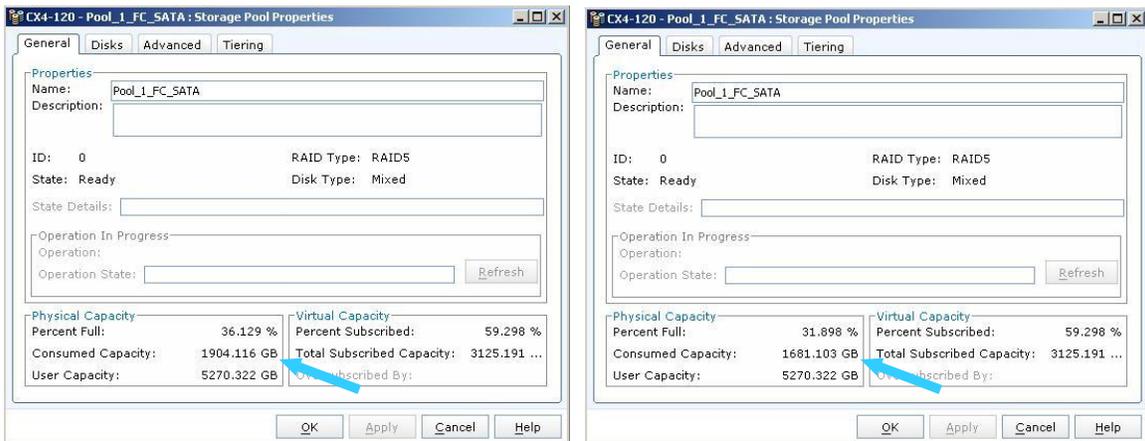
After compressing a LUN, users will see the consumed capacity of both the LUN and pool decrease. Consumed capacity in this case is reduced due to savings from both thin LUNs and compression. Examples of this can be seen in Figure 5 and Figure 6. Users will not see any change in reported capacity usage at the host level.

Figure 5 shows the compression of a thick LUN. The user capacity of the LUN is 250 GB, but 256 GB are consumed in the pool since thick LUNs consume capacity equal to the user capacity plus metadata. After compression, the LUN only consumes 34 GB of space in the pool, for a savings of just over 220 GB. This savings represents the benefits of both moving the LUN from a thick LUN to a thin LUN and compressing the data. Though there are savings in the pool, there will be no change in capacity usage reported by the server operating system.



**Figure 5. LUN Properties – Compression tab before and after compression**

The change in consumed capacity and resulting change in Percent Full can be seen in Figure 6. The capacity saved from compressing the LUN is returned to the pool for use by any other LUN in the pool. Server-visible metrics like User Capacity, Subscribed Capacity, and Percent Subscribed remain unchanged.



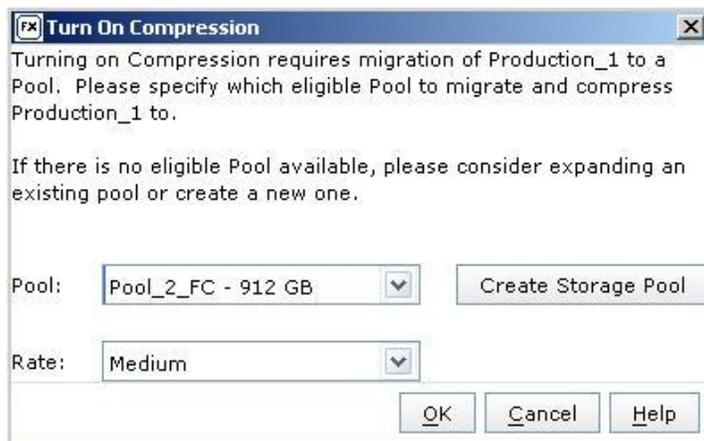
**Figure 6. Pool properties before and after compression**

There are user settings for the compression rate of High, Medium, and Low in the LUN Properties dialog box under the Compression tab as shown in Figure 5. The compression rate setting determines how fast

---

compression is performed, not the level of data reduction. The compression rate setting applies to initial compression, subsequent compression of new data, and decompression operations. The default rate setting is Medium, but it can be changed at any time.

An online migration of RAID group LUNs and thick LUNs to thin LUNs is performed when compression is enabled. All aspects of the migration are handled for the user with no interruption to server I/O. Thick LUNs are migrated within the pool where they reside. Since RAID group LUNs do not already reside in a pool, the user is prompted to select a destination pool for the compressed LUN. Figure 7 shows the popup Turn On Compression dialog box that allows the user to select an existing pool or launch the Create Pool dialog box to create a new destination pool. Users can also set the rate of the migration in this dialog box.



**Figure 7. Destination pool selection for RAID group LUN compression**

As shown in Figure 7, all currently eligible pools are listed in the Pool pull-down menu. Pools are only shown if they have enough free capacity to accommodate the user capacity of the RAID group LUN. Capacity equal to the user capacity of the RAID group LUN is reserved in the pool to ensure the process can complete. After the migration is complete, the original RAID group LUN is unbound and its capacity is available to bind new LUNs in the RAID group.

## ***Decompress operations***

The decompression process restores compressed data to its original size. When compression is disabled on a compressed LUN, the entire LUN is processed in the background. When the decompression process completes, the LUN remains a thin LUN and in the same pool. Capacity allocation of the thin LUN after decompression depends on the original pre-compression LUN type. If the LUN was originally a RAID group LUN or a thick LUN, the thin LUN is fully allocated after decompression. During the decompression of these LUNs, any unallocated capacity is filled with zeros. The LUN, while remaining a thin LUN, is fully allocated to mimic the consumed capacity of the original LUN type. If a LUN being decompressed was originally a thin LUN, the decompressed LUN will remain thin.

There are no capacity reservations made before the decompression process starts. However, if the pool becomes about 90 percent full, the compression state of the LUN will become **system paused**. Host I/O can continue to the LUN uninterrupted while compression is in this state. This behavior is implemented as a safeguard to ensure there is pool capacity available for new host I/O. If the user decides they would like to continue decompression without adding capacity to the pool, the user can manually override this safeguard and resume decompression. However, the compression state will again become **system paused** if the pool reaches 98 percent full. This safeguard cannot be overridden.

LUN Migration offers a compelling alternative to decompression. This is an attractive option when a LUN is more active than expected or in anticipation of a known period of high I/O activity. In these cases, in addition to decompressing data, a change in LUN type may be warranted and/or different physical spindles

may be used. All of these changes can be addressed with a single online LUN Migration operation. Users can migrate a compressed LUN to any LUN of their choosing as long as the target LUN's user capacity is the same or larger than the compressed LUN's user capacity. The target LUN can be any LUN type, compressed or not compressed, and reside in any pool or RAID group. As long as the target pool has enough capacity, there are no restrictions.

## Managing compression

Data Compression is a licensed feature that requires that a compression enabler be installed on the storage system. The compression enabler also requires that the thin provisioning enabler be installed if it is not already. The thin provisioning enabler is shipped with the compression enabler. Once installed, the compression options are visible within Unisphere™. Only the user roles of Manager and Administrator are allowed to manage Data Compression. The Monitor role and replication roles do not have the privileges to manage Data Compression. This section discusses the management dialog boxes available for compression and the options within them.

Figure 8 shows the LUN table in Unisphere, where users can easily see which LUNs have compression enabled. The dialog boxes that manage compression can be launched from this page, which include the LUN Properties and Compressed LUN Summary.

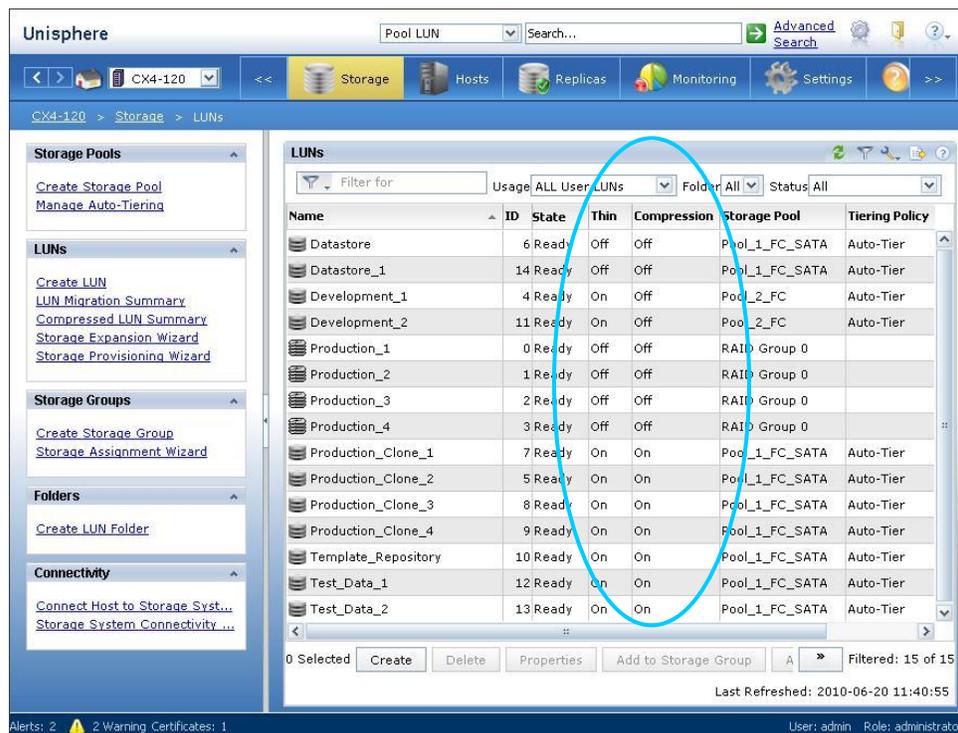


Figure 8. LUNs table in Unisphere

The Thin and Compression columns are not shown in the LUN table by default, but are helpful for understanding the status of the Data Compression feature. To display these columns, as shown in the figure, right-click the column header and choose **Performance View**, or go to the **Tools > Choose Columns** option (in the upper right-hand corner of the table) and select the **Thin** and **Compression** checkboxes.

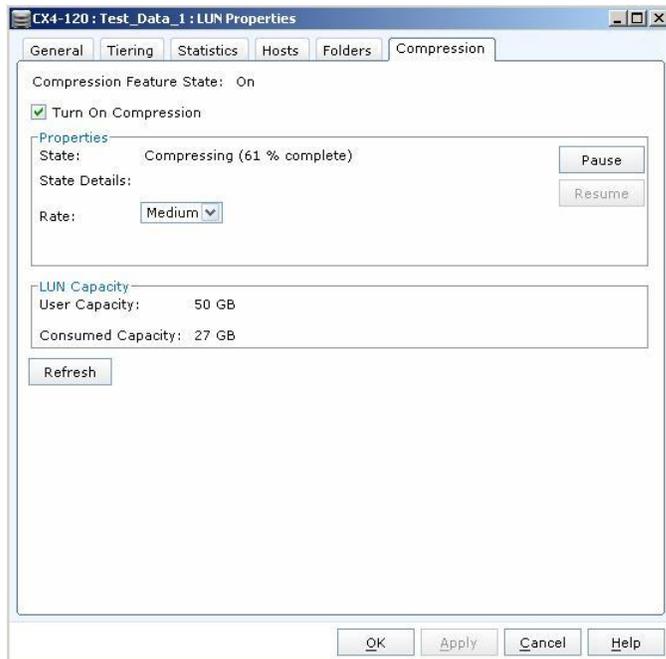
The Compression column displays **On** for LUNs that have compression enabled, and **Off** for LUNs that do not have compression enabled. The Thin column displays **On** for any LUN created as a thin LUN or any non-thin LUN that has completed its initial compression process. Since RAID group LUNs and thick LUNs

---

are migrated to thin LUNs during the initial compression, the Thin column remains **Off** until that process is complete.

The LUN Properties dialog box can be launched by double-clicking the LUN in the table; right-clicking the LUN and selecting **Properties**; or selecting the LUN in the table and clicking the **Properties** button at the bottom of the table.

Virtually all LUN-level controls for compression can be selected in the **Compression** tab of the LUN Properties dialog box. In Figure 9 a LUN is shown in the **Compressing** state.



**Figure 9. LUN Properties dialog box – Compression tab**

The Compression Feature State indicates the system-level state of Data Compression. Its values can be **On** or **Paused**. When On, all compression options are available to the user. When Paused, LUN-level compression controls are visible but are grayed out. Also, any in-process compression operations are paused with the exception of initial compression operations for RAID group LUNs and thick LUNs. If already started, those operations will continue even when the system paused.

The Data Compression feature can be enabled or disabled at any time by selecting and clearing the checkbox and clicking **OK** or **Apply**. If compression is disabled during an initial compression, the original LUN remains in place. This is most noteworthy for RAID group and thick LUNs. If compression is disabled during the initial compression of a RAID group LUN or thick LUN, the automatic migration is aborted and the user can continue to use the original LUN.

Users can pause compression and decompression at the LUN level via the LUN Properties dialog box. Any in-process compression or decompression will be paused. If the LUN is in between compression operations, no new compression processes will start. The only exception is that initial compression operations for RAID group LUNs and thick LUNs cannot be paused. They can be aborted as described above, but they cannot be paused. After the initial compression completes for these LUNs, any subsequent compression of new data or decompression can be paused.

The Compression state provides information on current compression operations and their progress (% Complete). A complete list of compression states is provided in “Appendix B.” Common states are:

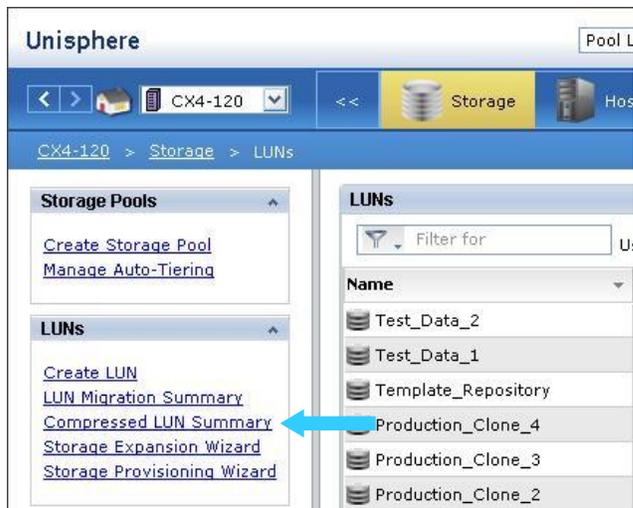
- **Compressing** — The compression process is running. This includes compression of new data as triggered by the system-defined threshold as well as initial compression of thick and thin LUNs.
- **Compressed** — The LUN is compressed and there are no active compression operations.
- **Compression Queued** — The maximum number of compression operations has been reached, so this operation is placed in the compression queue and will be automatically serviced as the system completes in process operations.
- **Decompressing** — The LUN is in the process of being decompressed.
- **Migrating** — This is the state of RAID group LUNs during the initial compression.

The Status field can provide additional information as to why compression is in a particular state. Under normal conditions this field is empty. It can provide useful information, like error codes and descriptions, if an error occurs.

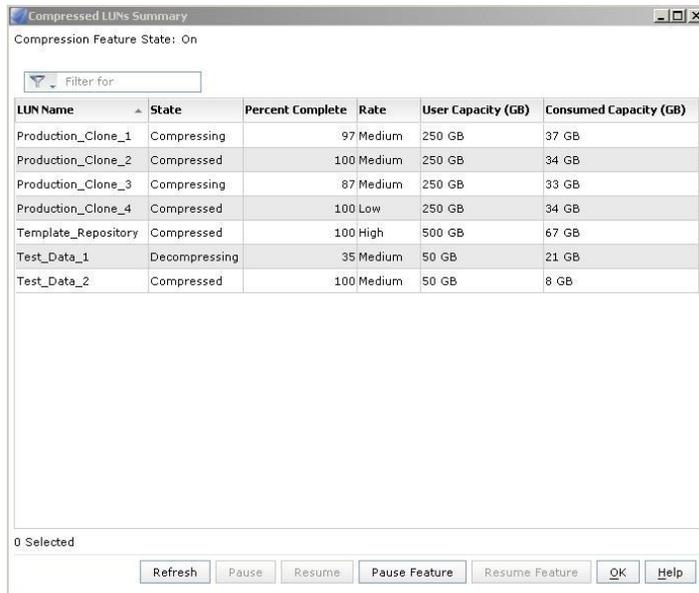
The Rate field shows the rate for initial compression, compression of new data, and decompression operations. It can be changed at any time, including during or between compression operations.

User Capacity is the LUN capacity as seen by the host. Consumed capacity represents the physical capacity used in the pool. For compressed LUNs, consumed capacity is the end result of thin LUN space savings and compression.

The Compressed LUNs Summary dialog box can be launched under the LUNs menu in the left menu pane of Unisphere as seen in Figure 10. This dialog box is used to monitor the status of all compression operations on the storage system and provides the user system-level pause and resume capabilities.



**Figure 10.** The blue arrow shows where to launch the Compressed LUNs Summary dialog box



**Figure 11. Compressed LUNs Summary dialog box**

Users can filter the LUN list to include only LUNs of interest. The user can click **Refresh** to update the status of any LUNs displayed in this dialog box, and individual LUN compression operations can be paused and resumed. **Pause Feature** and **Resume Feature** provide the user with system-level control of the Data Compression feature. This toggles the Storage System Compression Setting field between **On** and **Paused**. When paused, compression cannot be enabled on any LUN. All compression and decompression operations are paused with the exception of initial compression operations of RAID group LUNs and thick LUNs. This is most useful for halting and resuming compression activities before and after known periods of high-resource utilization.

A Compression Summary Report is available via the Unisphere Reports wizard. You can launch this by clicking the **Monitor > Reports > Reports Wizard** option. The report includes the System Compression State, the total number of compression operations, and significant attributes for each compressed LUN. LUN attributes include LUN Name, State, Percent Complete, Rate, User Capacity, and Consumed Capacity.

Lastly, all compression controls described in this section are available via the Navisphere® Secure CLI command, `compression`. The `compression` command has all the subcommands necessary to perform any of the above actions. Details can be found in the *Navisphere Command Line Interface (CLI) Reference* document on Powerlink.

## Limits and interoperability

### Limits

The following table details the limits for the Data Compression feature.

---

**Table 1. Limits for the Data Compression feature**

	<b>CX4-120</b>	<b>CX4-240</b>	<b>CX4-480</b>	<b>CX4-960</b>
<b>Total compressed LUNs per pool or system</b>	512	1024	2048	2048
<b>Concurrent compression operations per SP</b>	5	5	8	10
<b>Concurrent migration operations per array</b>	8	8	12	12

Notes:

- Concurrent compression operations include initial compression of data on thin LUNs, compression of new data, and decompression.
- Concurrent migrations include any initial compression of non-thin LUNs (RAID group LUNs and thick LUNs) where the LUN being compressed must be migrated to a thin LUN.
- Compression and migration limits are not dependent on each other. For example, SPA of a CX4-120 can have five compression operations running, and the sixth will be queued. At the same time eight thick LUN initial compressions could be running.

The following cannot be compressed:

- Private LUNs (including write intent logs, clone private LUNs, reserved LUNs, metaLUNs, and component LUNs)
- Snapshot LUNs
- Celerra iSCSI or file system LUNs
- A LUN that is already being migrated
- A LUN that is expanding or shrinking
- A mirrored LUN replicating to a storage system running pre-29 FLARE code

## **Interoperability**

In general, a compressed LUN has the same interoperability as a regular thin LUN. Compressed LUNs can be a source or destination LUN for CLARiiON replication applications. A LUN's compression status has no bearing on the compression status of any other LUN in a replication operation. For example, if a non-compressed LUN is the source of a SnapView clone group, the LUNs in the clone group can be compressed and/or non-compressed. Furthermore, if a compressed LUN is the source of a clone group the clones can be compressed and/or non-compressed. Compression can be enabled or disabled while in use by a replication application. Using the same example, compression can be enabled or disabled on a clone without affecting the operation of the clone group.

When a replication application such as SAN Copy writes to a compressed LUN, the I/O is treated as new data. The I/O written by the replication application can exceed the uncompressed data threshold, in the same manner host I/O would, and trigger the compression process. Some common replication writes that may trigger recompression are initial synchronizations, MirrorView/S syncs, MirrorView/A updates, Incremental SAN Copy updates, and clone synchronizations.

When creating a new replica that is intended to be compressed, users may want to enable compression *after* the initial synchronization. This way, compression is not triggered by the initial sync, which is typically a high bandwidth operation. After the initial sync completes, compression can be enabled to process all of the data at once, and compress new data as warranted by the system defined threshold. This is not a requirement, but merely a suggestion that may make initial synchronizations faster. To do this, you must have the capacity to accommodate the uncompressed data.

Sub-LUN FAST and FAST Cache are also compatible with compressed LUNs. Sub-LUN FAST can help facilitate the lowest capacity for compressed LUNs. Since compressed LUNs are likely to have low

---

performance profiles, FAST would probably place these LUNs on the lowest storage tier. Users can also explicitly set the tiering policy to **Lowest Tier**, thereby ensuring compressed LUNs remain on the lowest cost drives. To learn more about FAST, see the *EMC FAST for CLARiiON* white paper on Powerlink.

FAST Cache may help performance in cases where there are intermittent bursts of activity to a compressed LUN. However, it is not expected that it will bring compressed LUN performance up to par with non-compressed LUN performance. LUNs that become more active should either be decompressed or migrated to a non-compressed LUN. To learn more about FAST Cache see the *EMC CLARiiON and Celerra Unified FAST Cache — A Detailed Review* white paper on Powerlink.

## Special cases for LUN Migration and SAN Copy

When using LUN Migration to migrate data to a compressed LUN, always create a new thin LUN target and enable compression before starting the migration. LUN Migration is tightly integrated with compression; it is used in the initial compression of RAID group LUNs and thick LUNs. User-initiated LUN migrations to compressed LUNs can reap the same “compress on-the-fly” benefit as system-managed migration/compression operations. The destination LUN in this case should be a newly created thin LUN with no data and with compression enabled. After creating the new thin destination LUN and enabling compression, a user-initiated LUN migration compresses data inline with the migration I/O. At the end of the migration, the resulting LUN consumes the minimum amount of capacity possible. Specifying a destination LUN that has data on it is less efficient. The data is overwritten during the migration, and space reclamation is conducted as a separate process.

If using SAN Copy to replicate from a traditional LUN to a compressed destination LUN, do *not* enable compression on the destination LUN. Release 30 adds a space reclamation enhancement to CLARiiON Virtual Provisioning. Users can take advantage of this enhancement with SAN Copy when performing remote SAN Copy pull operations, or SAN Copy push or pull operations, within the same storage system.

This enhancement allows users to copy data from CLARiiON, Symmetrix<sup>®</sup>, or third-party storage systems to thin destination LUNs in a more efficient manner. Whitespace or strings of zeros in the data are detected by the system, and physical capacity is not allocated to the thin destination LUN for those areas. This provides immediate capacity savings. In prior releases, this would have resulted in a fully consumed thin LUN. This enhancement is for remote SAN Copy pull sessions and SAN Copy push and pull sessions within the same storage system. Remote SAN Copy push sessions result in a fully consumed thin LUN unless the source of the migration is a thin LUN on a system running release 29 or later.

If compression is enabled on the thin destination LUN prior to running the SAN Copy session, the space reclamation enhancement is not employed. In this case, the thin LUN is written to as if it were being fully consumed. At some point, the data written by SAN Copy triggers the compression process. After compression processes the data, the end result would likely be better than the space reclamation case alone. The compression algorithm would probably remove the same whitespace and zero strings, and compress the non-zero data wherever possible. Overall, it is more efficient to allow the SAN Copy pull session to complete without compression. This way, space reclamation is realized at the outset and the amount of capacity used is minimized throughout the process. After the copy is complete, compression can be enabled. Additional capacity savings can be realized from compression and the compression feature will not have to process as much data.

## Resource consumption and performance

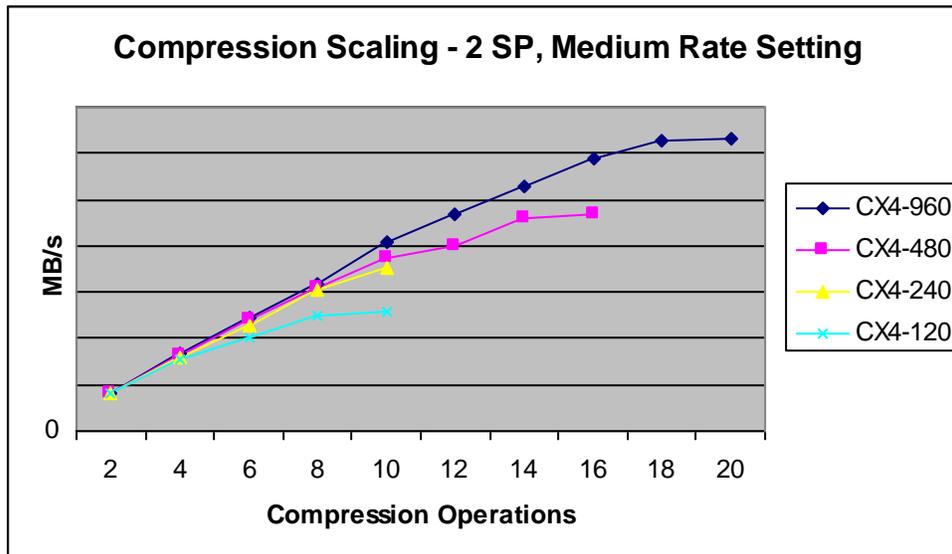
Moving datasets from traditional LUNs or thick LUNs to thin LUNs provides the benefits of recapturing and repurposing unutilized capacity and provides a “pay as you go” capacity consumption model. Users make this tradeoff when they do not have stringent IOPs and response time requirements. Compression takes this a step further, providing even greater capacity savings at the cost of additional system resources and LUN performance. Compression is best suited for data that is largely inactive but requires five 9s availability.

The compression rate settings offer control over the amount of resources that are dedicated to the compression process. For the standard compression process, which includes initial thin LUN compression and recompression of new data, the Low setting uses minimal ( $\leq 15$  percent) CPU resources for compression. At the Medium rate, CPU usage can range from 30-50 percent if the maximum number of concurrent compression operations is running. At High, CPU utilization can approach 60-80 percent if the maximum number concurrent compression operations are running. If only running half of the allowable concurrent operations at Medium or High, the range of expected CPU utilization is approximately half of the utilization ranges given above. Similar CPU utilization ranges can be expected for compression and decompression.

EMC recommends that you pause compression at the system level when response-time critical applications are running on the storage system. The process of returning capacity to the pool that has been freed by compression can contribute to CPU and write cache usage as data is consolidated onto as few slices as possible. This process can also continue after the compression process completes. Pausing the compression feature will ensure any background compression or associated space reclamation operations do not impact server I/O.

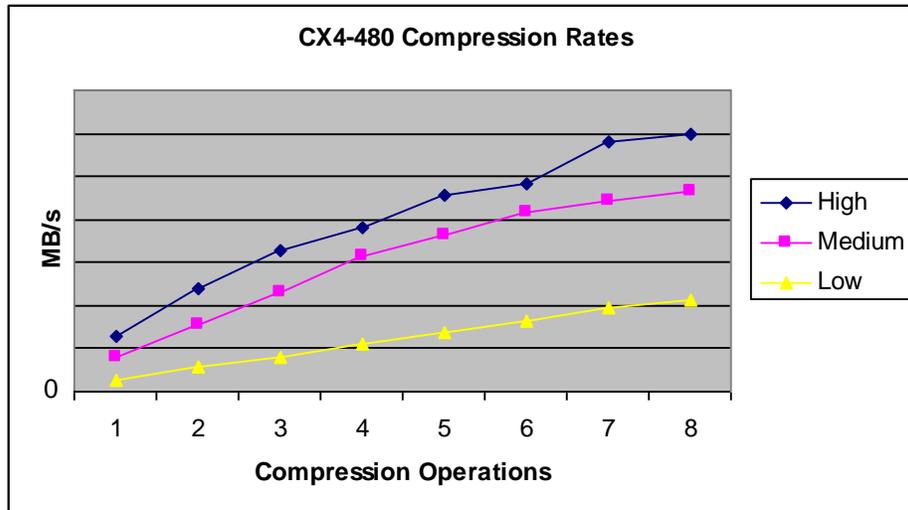
The initial compression of traditional RAID group-based LUNs and thick LUNs more closely mimics the behavior of standard LUN Migration to a thin LUN target. Compression is performed inline with the migration, so the overall rate may be 20-30 percent lower than migrating to non-compressed thin LUNs.

The relative compression throughput at the Medium rate for different storage system models is shown in Figure 12. This represents how much data is being processed by compression operations. Data used to generate the charts was compressed at a ratio of 1.5:1, or 65 percent. Compression operations were scaled two at a time, one for each SP, up to the maximum number of concurrent compression operations allowed.



**Figure 12. Compression scaling**

The difference in throughput with different compression rate settings for the CX4-480 is shown in Figure 13. There is a much larger difference between the Low and Medium rates than between Medium and High. This behavior is consistent across all models.



**Figure 13. Compression rates**

The compressibility of the data itself can have some bearing on compression throughput. When data is highly compressible (8:1), compression throughput may be 10-20 percent lower than is the case when compressing moderately compressible data (1.5:1). The differential is highest when the compression rate is set to High. The two cases are about equal when the rate setting is Low.

When data is highly compressible, compression and decompression throughput rates are essentially equal. Data that is moderately compressible, like the sample data used in the charts, results in lower decompression rates. For the sample data used that has a compression rate of 1.5:1, decompression throughput is 50-60 percent of compression throughput. Since the decompress operation only looks at compressed 64 KB areas, it is possible it will have to process less data than the compression operation, which has to process all new data. Assuming most areas are compressed, the decompression process is likely to take longer than initial compression. LUN Migration is an alternative that also allows for a change in physical drive location, LUN type, and RAID type.

Impact on server I/O can be moderate to high when compared to the performance of non-compressed thin LUNs. In some cases, like simulated file sharing environments, 25-50 percent lower throughput has been observed. In other cases, like large block, high bandwidth I/O, the impact can be much higher. The inline operations, inherent to reads and writes of compressed data, affect the performance of individual I/O threads; therefore, we do not recommend this for I/O-intensive or response-time-sensitive applications.

Compression's strength is improved capacity utilization. Therefore, compression is not recommended for active database or messaging systems, but it can successfully be applied to more static datasets like archives, clones of database, or messaging-system volumes. The best use cases are those where data needs to be stored most efficiently and with a high degree of availability.

---

## Conclusion

EMC Data Compression offers compelling capacity savings for many data types that have lower-performance requirements, but still require CLARiiON five 9s availability. Management of the Data Compression feature is very easy. It's simply enabled on a per-LUN basis and the system automatically monitors and manages compression operations. Compressed LUNs are compatible with virtually all CLARiiON replication and optimization products and features including LUN Migration, SnapView, FAST, and FAST Cache.

## References

- Unisphere Online Help
- *EMC CLARiiON Virtual Provisioning* white paper
- *EMC FAST for CLARiiON* white paper
- *EMC CLARiiON and Celerra Unified FAST Cache* white paper
- *Achieving Storage Efficiency through EMC Celerra Data Deduplication* white paper

---

## Appendix A: NTFS file system conditioning with SDelete

Many file systems do not efficiently reuse the space associated with deleted files. When files are deleted from NTFS file systems, the deleted files' data continues to be stored in the file system until it is overwritten by new data. When files are frequently deleted, the free space in the file system may gradually become filled with deleted file data that is no longer accessible by the file system. Deleted file data reduces the effectiveness of EMC Data Compression. The retention of deleted file data is a characteristic of the file system and is relevant regardless if LUNs are presented directly to a Windows server or the Windows server resides in a VM.

Data Compression processes blocks associated with deleted files the same way that data is processed for valid files. Unused (never-used) blocks in the NTFS file system compress down to zero consumed capacity, but deleted file blocks are only as compressible as the deleted file data that continues to be stored on them.

The SDelete utility from Microsoft will replace deleted file blocks with zeros when it is invoked with the `-c` option. Note that deleted file blocks cannot be removed by defragmentation or reformatting the file system. The blocks zeroed by SDelete do not consume any capacity once they are processed by Data Compression. This can have a profound impact on compression results.

For example, let's assume a 100 GB file system has files that are 1.5:1 compressible (33 percent space savings). This file system resides on a 100 GB traditional LUN. Also assume that the file system reports used space to be 60 GB, but since the file system has been in use for some time, there are actually 90 GB of used blocks in the file system. This means that only 10 GB of the file system capacity is unallocated to data.

Without SDelete, the resulting compressed LUN will consume roughly 60 GB in the pool, since all 90 GB worth of blocks is compressible at a 1.5:1 ratio. Thin LUN metadata may add another 3-4 GB of consumed capacity. The 10 GB of unallocated data does not consume any capacity once it is compressed. This results in overall capacity savings of roughly 40 GB.

If SDelete is run on the file system prior to being compressed, Data Compression is more effective. There will still be 60 GB of data, but the extra 30 GB of deleted file blocks and the 10 GB of unallocated capacity are overwritten by zeros by SDelete. In this case, the compressed LUN only consumes 40 GB in the pool for a total savings of 60 GB. The data itself compressed at a ratio of 1.5:1, but the zeroed capacity does not consume any space in the compressed LUN. In this example, using SDelete yielded an additional 50 percent of space savings.

SDelete writes to all space not consumed by a valid file data, therefore if the file system resides on a thin LUN, SDelete will cause the thin LUN to become fully consumed. Therefore, users must be sure they have adequate pool capacity if they choose to run SDelete on a thin LUN. On traditional LUNs or thick LUNs hosting NTFS file systems, SDelete should be run prior to enabling compression to maximize space savings. Additional information, including the download of the utility can be found at the following link: <http://technet.microsoft.com/en-us/sysinternals/bb897443.aspx>.

## Appendix B: Compression states

Table 2. LUN compression states

State	Description	Valid for LUN type
<b>Initializing</b>	Compression is first enabled, and background setup operations are performed by the system.	All
<b>Compressing</b>	The thin LUN or thick LUN is being compressed.	Pool LUNs
<b>Compressed</b>	An initial or subsequent compression has completed.	Pool LUNs
<b>Compression Queued</b>	The state of compression prior to moving to Compressing. It is in this state for a short duration unless system limits for concurrent operations are exceeded.	Pool LUNs
<b>Compression Paused</b>	The compression operation has been paused. I/O continues to the LUN. New data is not compressed; compressed data remains compressed.	Pool LUNs
<b>Compression Faulted</b>	The compression operation did not complete successfully. More information may be available in State Details.	Pool LUNs
<b>Decompressing</b>	The LUN is being decompressed.	Thin LUNs
<b>Decompression Queued</b>	The LUN is waiting to be decompressed. It is usually in this state for a short duration unless system limits for concurrent operations are exceeded.	Thin LUNs
<b>Decompression Paused</b>	The decompression operation has been paused. I/O continues to the LUN. New data is not compressed; compressed data remains compressed.	Thin LUNs
<b>Decompression Faulted</b>	The decompression operation did not complete successfully. More information may be available in State Details.	Thin LUNs
<b>System Paused</b>	The system has paused a decompression operation due to lack of capacity in the pool. The user can either add capacity to the pool or migrate LUNs out of the pool to free up capacity.	Thin LUNs
<b>Migrating</b>	This is the state of a RAID group LUN during the initial compression. RAID group LUNs are migrated to a thin LUN, and compressed inline with the migration. This operation cannot be paused. It can be canceled.	RAID group LUNs
<b>Migration Queued</b>	This is the state of a LUN while it is waiting to be migrated. It is usually in this state for a short duration unless system limits for concurrent operations are exceeded.	RAID group LUNs
<b>Migration Faulted</b>	The migration operation did not complete successfully. More information may be available in State Details.	RAID group LUNs