



IBM Real-time Compression Evaluation User Guide for IBM Storwize V7000 and SAN Volume Controller Version 6.4

Version 2.2
June 20, 2012

TABLE OF CONTENTS

TABLE OF CONTENTS 2

Real-time Compression Free Evaluation for 45 Days Program 3

Overview 4

Introduction..... 4

Data Compressibility 5

Additional Considerations..... 5

Compression and Various Data Types..... 7

Volume Types 8

A Note about Synthetic Benchmark Tools..... 9

Step-by-step Instructions to Set Up for 45 day Evaluation..... 10

Step 1 – Evaluate expected compression rates for existing data..... 10

Step 2 –Assess the Suitability of I/O Group for Compression 11

Step 3 – Enable Real-time Compression evaluation license 12

Step 4 - Compress existing data..... 12

Step 5 – Confirm Compression savings..... 14

Step 6 – Evaluate performance 15

Step 7 – Remove Original Copy..... 16

Rollback to Non-Compressed Volume Instructions..... 16

Additional Resources 17

Real-time Compression Free Evaluation for 45 Days Program

IBM authorizes existing Storwize V7000 and SAN Volume Controller customers to evaluate the new Real-time Compression capability at no cost. Customers who wish to take advantage of this opportunity are encouraged to download the Storwize V7000 Version 6.4 software and enable the Real-time Compression function to help evaluate the potential benefits they can obtain based on their own specific environment and application workloads. The evaluation period is for up to 45 days and it begins when a customer enables the compression function. At the end of the evaluation period, the customer must either purchase the required licenses for Real-time Compression or disable the function.

Customers taking advantage of this opportunity must work with their local IBM Business Partner or their IBM sales representative through the evaluation process to ensure that the version 6.4 software is downloaded from a designated IBM site, the evaluation period allowance is maintained, and an appropriate decision is made afterwards. This User Guide provides step by step instructions on how to setup IBM Real-time Compression for evaluation on an installed Storwize V7000 or San Volume Controller system. Note that IBM recommends that Real-time Compression Evaluation Software not be used in a production environment. If you have questions about how to set up IBM Real-time Compression your system, please contact your IBM or IBM Business Partner Sales or Technical Sales representative for assistance. If you encounter problems with your system after implementing IBM Real-time Compression Evaluation software, it is recommended that you disable Real-time Compression following the roll back instructions in this document prior to calling IBM Support for assistance. If the problem persists with Compression disabled, follow the normal process for contacting IBM Technical Support for assistance. The base Storwize V7000 or San Volume Controller system and related software on that system for which you have current SW Maintenance in place remain fully supported as entitled from IBM support during the Real-time Compression evaluation period. It is the customer's responsibility to ensure that their system is properly licensed for the Real-time Compression function should they choose to use it after the evaluation period is over.

IBM recommends reading through this document before using compression in order to achieve optimal results with implementing compression in the customer's environment. In addition, this document includes a detailed step-by-step walkthrough that provides a guide for first time use of compression. The guide covers the planning, execution, and evaluation of compression efficiency in the system environment, in a completely non-disruptive way using Storwize V7000 and SAN Volume Controller advanced data management features.

Overview

Software version 6.4 of Storwize V7000 and SAN Volume Controller products introduces a new capability of Real-time Compression. This capability is made possible through the integration of the Random Access Compression Engine (RACE) technology. IBM's Real-time Compression delivers high-performance compression that supports active primary workloads. This innovative, easy-to-use compression is now fully integrated into the Storwize V7000 and SAN Volume Controller products.

Storwize V7000 Real-time Compression typically delivers 50% or better compression for data that is not already compressed, and helps reduce:

- Storage purchase costs
- Rack space
- Cooling and energy
- Software costs for additional functions

IBM Real-time Compression increases the capacity of existing storage infrastructure, helping enterprises meet the demands of rapid data growth while enhancing storage efficiency. IBM's Real-time Compression uses IBM's patented real-time data compression techniques to primary storage, delivering optimization and savings in Storwize V7000 and SAN Volume Controller, for both internal as well as external/virtualized storage. The result is exceptional cost savings, ROI, operational and environmental efficiencies.

Introduction

Release 6.4 introduces a new 'Compressed' volume preset. Compressed volumes are a new type of thinly provisioned volumes. Compressed volumes use IBM's Random Access Compression Engine (RACE) technology to deliver Real-time Compression capabilities in the SAN Volume Controller and Storwize V7000 products. Using release 6.4, users can implement compression and evaluate its value and benefits in their Storwize V7000 and SAN Volume Controller environments. Evaluation of this unique IBM storage efficiency offering provides users with a good understanding of how this technology can help improve their storage systems efficiency as well as operational efficiencies. Specifically, evaluation users can assess key areas such as: compression efficiency, transparency in the environment, and application performance.

This document discusses the attributes of IBM's Real-time Compression in Storwize V7000 and SAN Volume Controller products. The document also provides technical information and guidelines on what should be considered when deploying compression in the evaluation environment.

Data Compressibility

Choosing which workloads are good candidates for compression is determined by the nature of the data that is targeted to be stored in a compressed volume and the compression savings expected for this data. As a general rule of thumb, IBM recommends to compress workloads that get more than 45% compression savings. Use the following threshold for volume compressibility in order to decide whether to compress a volume or not to:

Data Compression Rate	Recommendation
Higher than 45% compression savings	Use Compression
25-45% compression savings	Evaluate workload with compression
Lower than 25% compression savings	Avoid Compression

IBM provides a utility called **Comprestimator** that analyzes the expected compression savings for an existing volume containing data considered for compression. Comprestimator is a command line host-based utility that can be used to estimate expected compression rates for existing block devices in their native environment. The utility uses advanced mathematical and statistical algorithms to perform the device sampling and analysis process efficiently and in a very short time period. The utility displays its accuracy level by showing the maximum error range of the expected compression results achieved, based on the formulas it uses. The utility runs on a host that has access to the devices that will be analyzed, and performs only read operations so it has no effect whatsoever on the data stored on the device. It is important to emphasize that since Comprestimator is a host-based utility that analyzes any block device accessible from the host it is running on, it can be used on any device, regardless of the storage system it is stored on. Please note that it is highly recommended to use Comprestimator to analyze volumes that contain as much data as possible rather than volumes that are mostly empty of data. This increases accuracy and reduces the risk of analyzing old data that was already deleted, but may still have traces on the device.

Depending on the environment configuration, in some cases Comprestimator will be used on more than one host, in order to analyze additional data types that are stored in volumes accessible from other hosts. A copy of Comprestimator can be downloaded from IBM's technical support website 'Fix Central' – [StorageDisk-2076-Comprestimator_Utility_v1.14](#) .

Additional Considerations

In addition to compression rate achieved for the compressed workload, there are other factors and metrics to be taken into consideration when choosing what data to compress on which system. These considerations include the following:

- IBM Real-time Compression is supported only on the following systems:

- IBM Storwize V7000
- IBM SAN Volume Controller (SAN Volume Controller)
 - CF8 models
 - CG8 models
- Real-time Compression is supported for both internal storage and externally virtualized storage. Using the external virtualization capability of Storwize V7000 and SAN Volume Controller, compression can be used with all storage systems supported for external virtualization with release 6.4.
- Compression is supported for a maximum of 200 compressed volumes per I/O group in this release – this limit applies only to compressed volumes; there is no restriction for the number of non-compressed volumes in the I/O group. A cluster with 4 I/O groups can support as many as 800 compressed volumes.
- System resources for using compression – before using compression on existing Storwize V7000 and SAN Volume Controller systems it is highly recommended to measure the current system utilization, in order to ensure it has the CPU resources required for compression.
 - On V7000 systems, SAN Volume Controller CF8 nodes, and early SAN Volume Controller CG8 nodes (4-core systems)
 - If CPU utilization on the nodes in the I/O group is below 25%, then this I/O group is suitable for using compression.
 - If the CPU utilization of a node is sustained above 25% most of the time, then this I/O group may not be suitable for compression, as it is too busy. If your system is currently supported for the Non Disruptive Volume Move feature, then it is possible to move this volume to another I/O group that has the resources required for compression.
 - On latest SAN Volume Controller CG8 nodes (6-core systems) –
 - If CPU utilization on the nodes in the I/O group is below 50%, then this I/O groups is suitable for using compression.
 - If the CPU utilization of a node is sustained above 50% most of the time, then this I/O group may not be suitable for compression, as it is too busy. If your system is currently supported for the Non Disruptive Volume Move feature, then it is possible to move this volume to another I/O group that has the resources required for compression.
 - Note than system resource allocation for compression is done as soon as the first compressed volume is created, even if it is not used yet.
- Compressed volumes cannot be used with Easy Tier. Specifically, Easy Tier and compression can both work on the same Storwize V7000 or SAN Volume Controller system, but Easy Tier will not handle compressed volumes.
- Application I/O patterns –
 - Real-time Compression and its RACE foundations were designed and purpose built for enterprise-level random-access application workloads. While exceling in such application workloads, IBM does not recommend using Real-time Compression for heavy sequential write oriented workloads. The most common use-case for such workload is using Real-time Compression to compress volumes that are used as disk-based backup targets.

- Workloads generating random-access I/O patterns are highly recommended for use with Real-time Compression.

Compression and Various Data Types

The best candidates for data compression are data types that are not compressed by nature. These data types range between many workloads and applications, such as databases, character/ASCII based data, server virtualization, Computer Aided Design (CAD) data, software development, engineering data, vector data, and other general-purpose enterprise applications and workloads.

In cases where the data is compressed by nature, little additional savings can be achieved by using storage system compression. It is recommended not to spend system resources on data that can't be compressed. Workloads that create compressed data should not be stored on compressed volumes. The following example represents workloads and data that is already compressed by nature:

- Compressed audio, video, and image file formats – files such as jpeg, png, mp3, medical imaging (DICOM), mpeg2, etc.
- Compressed user productivity file formats – MS-Office 2007 newer formats (pptx, docx, xlsx, etc.), pdf files, MS-Windows executable files (exe), etc.
- Compressed files formats – such as zip, gzip, rar, cab, tgz, etc.

In cases where data is encrypted by the client or application used, no savings can be achieved by using any compression method. Due to the high entropy found in encrypted data, compression algorithms can't find repetitions within them, and are completely ineffective in most cases. Workloads that create encrypted data should not store their data on compressed volumes. If, however, encryption in the storage system/environment operates at the back-end disk level and is performed post compression – the data can be compressed and will have the same encryption capabilities in the back-end.

Following is a high-level short list focused on workloads that are good candidates and list of workloads that should be avoided for compression:

- Common workloads suitable for compression
 - Databases (OLTP/DW) – DB2, Oracle, MS-SQL, etc.
 - Applications based on databases – SAP, Oracle Applications, etc
 - Server Virtualization – KVM, VMware, Hyper-V, etc.
 - Other compressible workloads – engineering, seismic, collaboration, etc.
- Common workloads NOT suitable for compression
 - Workloads using pre-compressed data types such as video, images, audio, etc.
 - Workloads using application/network-based encryption
 - Heavy sequential write oriented workloads
 - Other workloads using incompressible data or data with low compression rate

Volume Types

It is very important to understand the nature of the data that is contained within a volume. There are two types of volumes to consider – homogeneous and heterogeneous. As a rule of thumb, homogeneous volumes are typically better candidates for compression. The following sections explain the types of volumes and how to decide if a volume should be compressed.

Homogeneous volumes are volumes that contain data that was created by a single application, and store data of the same kind. The following examples represent application workloads that typically store their data on homogeneous volumes:

- Databases – such as DB2, Oracle, MS-SQL etc.
- Server virtualization – such as VMware, Hyper-V, KVM etc.

Homogenous volumes are good candidates for compression since the majority of the data stored in such volumes typically achieves similar compression across the volume. This is true as long as the homogeneous volume contains compressible data. If the volume contains incompressible data (already compressed by nature or encrypted) it will not compress well, even though the volume is homogeneous.

Heterogeneous volumes are volumes that contain diverse content (data that was created by several different workloads). Heterogeneous volumes contain data of different types. Since different data types populate such volumes, there are situations where compressed or encrypted data will be stored on such volumes. In such cases, system resources can be spent on data that can't be compressed.

Heterogeneous volumes are typically found in user productivity environments, such as user home-directories, where users typically store diverse types of files. In many situations compressed content can be found in such environments as mentioned in 'Data Types' section above. As a result of this it is recommended to avoid compressing heterogeneous volumes, unless the heterogeneous volumes contain only compressible, unencrypted data.

Review the following guidelines in order to achieve the highest compression savings in your environment:

- Avoid using any client, file-system, or application compression when using V7000/SAN Volume Controller compressed volumes. In most cases, data that is already compressed will not get significant additional savings from compressing it again.
- Avoid using any client, file-system, and application encryption when using V7000/SAN Volume Controller compressed volumes. Encrypted data is not compressible.
- Avoid compression of V7000/SAN Volume Controller volumes that are used to store homogeneous or heterogeneous data types that contain compressed data or compressed file types.
- Avoid compression of V7000/SAN Volume Controller volumes that are used to store homogeneous or heterogeneous data types that contain encrypted data or encrypted files.

V7000/SAN Volume Controller management interfaces can be used to get rich compression savings information that can be used to measure compression effectiveness in the environment.

Compression savings can be reviewed and analyzed in volume, pool, or system level. This information can be used to measure the savings achieved in your environment for workloads using these volumes, pools, and systems.

A Note about Synthetic Benchmark Tools

Using real-life workloads when testing a system's performance provides the most useful and valuable information. This is especially true when testing compression and many other storage efficiency technologies. Tested workload generates I/O patterns that characterize the real application behavior, and the data stored and accessed looks like the real data that the system is highly likely to handle when it is used in production environments later.

Another way to test performance is to use synthetic test tools. These tools can be used for testing the limits of traditional storage systems or volumes. When using such tools it is important to remember that these tools typically generate synthetic I/O patterns using synthetic data sets. Since compression and other storage efficiency technologies are influenced by the nature of the data being optimized and how it is being accessed, synthetic benchmark results can easily provide the wrong impression for expected performance. This can make such results inaccurate, and in some cases even practically useless when trying to measure real-life application environment behavior. If you plan to use such tools, please consider their following characteristics:

- The data used in the benchmark – benchmarks typically use patterns that vary between being “all zeros” to being “completely random”. Production workloads generate content with varying data patterns that represent the information contained in the data. Using synthetic data will impact compression rates that can subsequently cause performance change in many cases.
- I/O access patterns generated – benchmarks typically generate a mix of synthetic I/O patterns that uses sequential/random reads/writes. Production workloads generate I/O patterns that are related to the data they store and access. In many cases the order in which data is read is identical or similar to the way it was written to storage. IBM's RACE technology uses time-based compression, which takes advantage of the similar read/write order. This creates both compression and performance efficiencies, that can't be tested when using synthetic I/O patterns. This cannot be achieved with application simulators as well, since they just generate synthetic I/Os that are completely unrelated to the data stored in the volumes they access.
- Hot spots – benchmarks sometimes generate workloads that are not realistically spread across application volumes. So a benchmark can create synthetic hotspots on small range of data. On the other hand, if there are no hot spots at all and the test accesses the entire range of offsets across a volume, it will not represent a real environment as well. Using production workloads create real hot spots rather than synthetic ones, allowing the system cache to be more efficient.

Step-by-step Instructions to Set Up for 45 day Evaluation

The following procedure guides you through the process of how to configure your SAN Volume Controller or Storwize V7000 system to evaluate compression in your storage environment. The procedure also includes steps on how to disable Real-time Compression if needed.

The procedure describes the configuring and evaluation of a single volume. The volume described is in I/O group X. I/O group X contains nodes or node canisters S and T. The same technique should be applied to multiple disks in a multiple I/O groups.

Step 1 – Evaluate expected compression rates for existing data

Use the Comprestimator tool to analyze the estimated compression rate for volumes you plan to compress. A copy of Comprestimator can be downloaded from IBM's technical support website 'Fix Central' – [StorageDisk-2076-Comprestimator Utility v1.14](#) .

Comprestimator is a command line host-based utility that can be used to estimate expected compression rate for block-devices. The utility uses advanced mathematical and statistical algorithms to perform the sampling and analysis process in a very short and efficient way. The utility runs on a host that has access to the devices that will be analyzed, and performs only read operations so it has no effect whatsoever on the data stored on the device.

After installing Comprestimator on a host that has access to the block-devices/LUNs/volumes you plan to compress. Following that, use Comprestimator to analyze the device for expected compression. It is highly recommended to use Comprestimator to analyze volumes that contain as much data as possible rather than volumes that are mostly empty of data. This increases accuracy and reduces risks to analyze old data that was already deleted, but may still have traces on the device.

The following example shows how to run the tool and analyze a device ('/dev/sdb'). Comprestimator can be used with the '-h' flag for all other available options.

```
comprestimator -d /dev/sdb
```

Review the output generated by Comprestimator for the volume. Depending on the volume type you are analyzing choose one of the following values from the output as expected compression rate for the volume:

- When analyzing an existing thinly provisioned volume/block-device, use the '*Compression Savings (%)*' value in Comprestimator output.
- When analyzing an existing fully allocated volume/block-device that was zeroed before it was used to store the current data on the volume, use the '*Compression Savings (%)*' value in Comprestimator output.
- When analyzing an existing fully allocated volume/block-device that was **not** zeroed before it was used to store the current data on the volume, use the '*Total Savings (%)*' value in Comprestimator output.

Verify the expected compression rates for this volume with the chart in Data Compressibility section to decide whether this volume is a good candidate for compression.

Step 2 – Assess the Suitability of I/O Group for Compression


Check that the CPU utilization in I/O group X is at an appropriate level.

1. Log in to the GUI.
2. Open Monitoring -> Performance.



Note: Monitoring Icon looks like this

3. Switch between selecting node S and node T from the dropdown box in the upper left corner of the GUI. This dropdown box will select “System Statistics” by default.



4. Monitor the CPU utilization over busy periods of the day or the week.
5. Only if you are using SAN Volume Controller 2145-CG8 nodes, you should first find out what is the number of CPU cores in your system. To determine if your system uses 4 or 6 CPU cores:
 - a. Select Monitoring > System to view VPD information related to the processor. The version entry for 2145-CG8 nodes contains one of the following two values:
 - i. Intel® Xeon® CPU E5630 - 4 cores
 - ii. Intel® Xeon® CPU E5645 - 6 cores



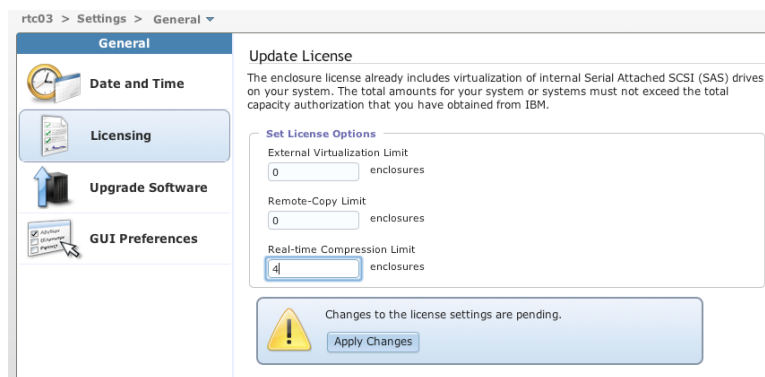
6. Verify if your system is suitable for using compression based on your system’s CPU utilization:

- On V7000 systems, SAN Volume Controller CF8, and early SAN Volume Controller CG8 nodes (4-core systems) –
 - If CPU utilization on the nodes in the I/O group is below 25%, then this I/O group is suitable for using compression.
 - If the CPU utilization of a node is sustained above 25% most of the time, then this I/O group may not be suitable for compression, as it is too busy. If your system is currently supported for the Non Disruptive Volume Move feature, then it is possible to move this volume to another I/O group that has the resources required for compression.
- On latest SAN Volume Controller CG8 nodes (6-core systems) –
 - If CPU utilization on the nodes in the I/O group is below 50%, then this I/O groups is suitable for using compression.
 - If the CPU utilization of a node is sustained above 50% most of the time, then this I/O group may not be suitable for compression, as it is too busy. If your system is currently supported for the Non Disruptive Volume Move feature, then it is possible to move this volume to another I/O group that has the resources required for compression.

Note: It is also possible to perform this analysis using performance metrics in TPC. Simply draw the CPU % graph in the StorageSubsystemPerformance ByNode view from TPC and observing the CPU utilization over a longer period of time.

Step 3 – Enable Real-time Compression evaluation license

In order to enable the 45 days evaluation license, enter '0' for total number of enclosures in your system. The system will flag that you do not have licenses for the compressed volumes, but you will be able to use the Real-time compression function for the 45 day evaluation period. Select “Apply Changes” to complete the process.



Step 4 - Compress existing data

Once you have confirmed that the I/O group is suitable for compression, perform the following steps to add a compressed copy to the volume:

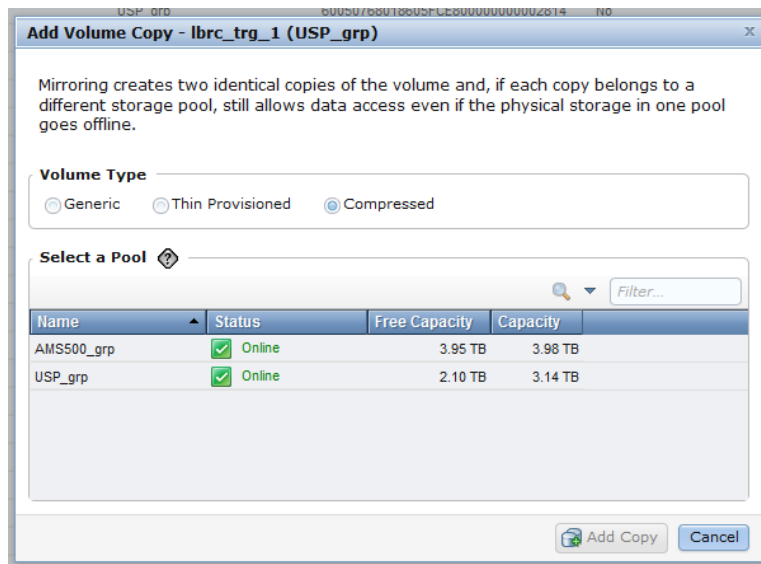
Important note: if the source volume you plan to compress is a generic (fully-allocated) volume, please make sure you thoroughly review section 4.4 of the “Real-time Compression in SAN Volume Controller and Storwize V7000” Redpaper and follow the recommended actions to get the best results before performing the following steps. The Redpaper is available at <http://www.redbooks.ibm.com/redpieces/pdfs/redp4859.pdf>

1. Open the one of the following Volume panels:
 - a. Volumes->Volumes
 - b. Volumes->Volumes by Pool
 - c. Volumes->Volumes by Host

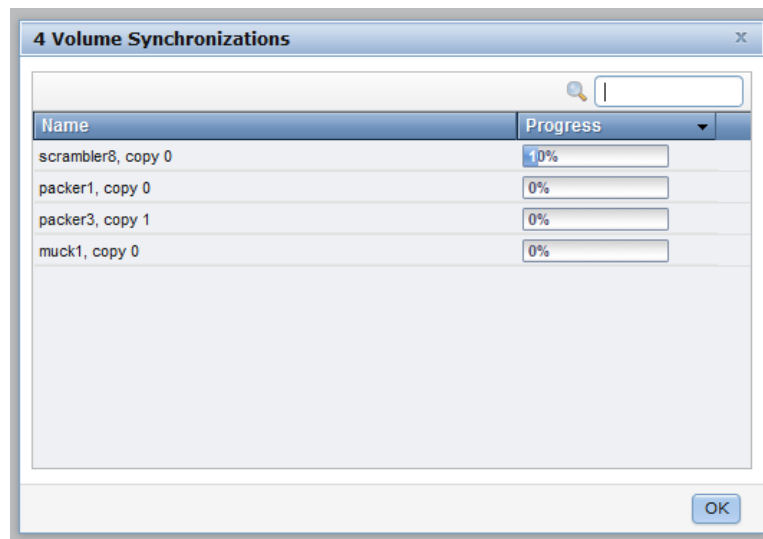
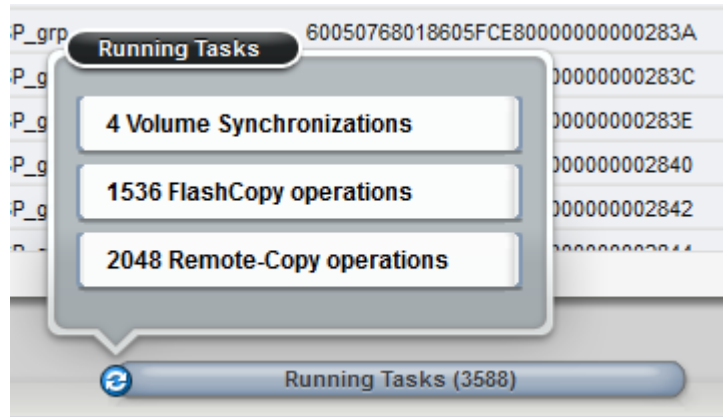


Note – the Volumes icon looks like this:

2. Select the volume that contains the data you want to compress.
3. Click on Actions at the top of the table and select Volume Copy Actions -> Add Mirrored Copy.
4. In the Add volume copy panel – select the Compressed radio button and choose the appropriate storage pool then press the Add Copy button.



5. Use the running tasks pod at the bottom of the GUI to monitor the Volume Synchronizations.



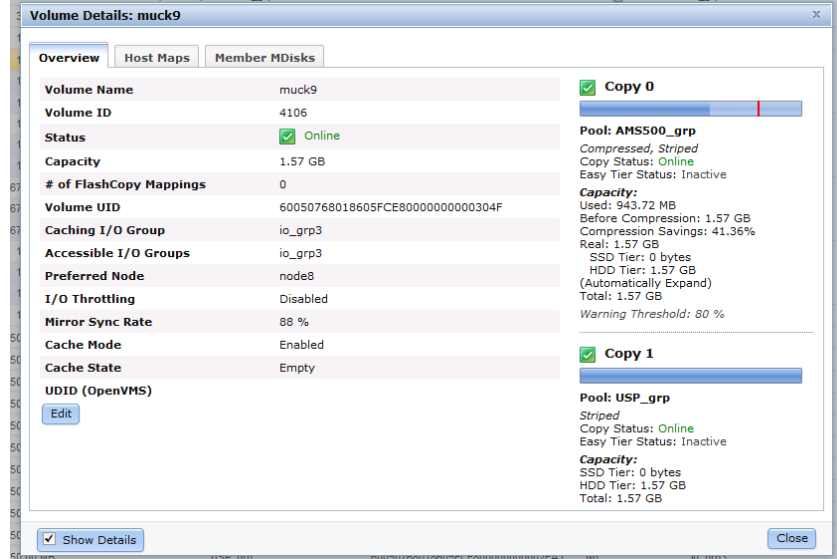
6. Once the synchronization has completed it will be removed from this list.

Step 5 – Confirm Compression savings

Use the GUI to check how much space was saved by using compression:

1. Right click on the volume and show properties.
2. Tick the “Show Details” checkbox.


3. The right hand side of the properties panel will show the disk space usage of each copy.



4. The following terms are used;
 - a. Used: How much disk space is currently being used by the compressed volume.
 - b. Before Compression: The amount of uncompressed data written to the compressed volume.
 - c. Compression Savings: This shows what is percentage of the capacity saved by using compression. This is the difference between the “Before Compression” and “Used.” Measures how well the data being written to the volume is compressing.
 - d. Real: The amount of capacity which has been assigned from the storage pool to this compressed disk. Note – this will always be greater than the used space.
5. Press ‘Close’ once you have finished reviewing the data. If you are satisfied with the compressions savings achieved then continue to Step 5.
6. If you no longer want to compress this volume, then delete the compressed volume copy by
 - a. Right click on the volume copy which is compressed.
 - b. Click “Delete this copy.”

Step 6 – Evaluate performance

To validate the performance realistically, it is recommended to switch the primary copy of the volume to the compressed copy. This causes all read I/O operations to be sent to the compressed volume.

1. The compressed copy of the volume is identified in the volume table with the following icon: 
2. The primary copy is identified by an asterisk (*) next to the Copy ID.
3. Right click on the Compressed copy and choose “Make Primary.”

At this point you should evaluate application performance of the application that is now using compressed volumes.

Note: The performance of the volume at this point will not be identical to the performance of the final compressed volume because any write operations will still be sent to both the compressed and non-compressed copy. Therefore, if both copies are still in the same storage pool, then the number of write operations to that storage will have been increased rather than reduced.

If you are satisfied with application performance when using compressed volumes, then continue to Step 6.

If you decide you no longer want to compress this volume, then delete the compressed volume copy by:

1. Right click on the volume copy which is compressed.
2. Click “Delete this copy.”

Step 7 – Remove Original Copy

If you are satisfied with the performance and space saving of the compressed volume, then contact your Business Partner or IBM to purchase a Real-time Compression license for your system. After purchasing a license you may remove the original copy of the volume to complete the migration process and add additional compressed volumes:

1. Right click on the volume copy which is not compressed.
2. Click “Delete this copy.”

We hope you are satisfied with the benefit that Real-time Compression can provide your Storage environment.

This completes the technical walkthrough section.

Rollback to Non-Compressed Volume Instructions

If for any reason during the process you wish to convert a compressed volume to a non-compressed volume, simply follow the volume mirroring process described in Step 3 above in reverse.

If at any time you wish to stop running compression in an entire I/O group, then all compressed volumes in that I/O group must either:

- Be moved to another I/O group using NVDM feature.
- Be converted back to a non-compressed volume by following this same procedure in reverse.

Additional Resources

IBM Redpaper Draft: [Real-time Compression in SAN Volume Controller and Storwize V7000](#).

[IBM Storwize V7000 Information Center](#) - Version 6.4 / Product Overview / What's New / Compressed Volumes.