

Implementing IBM Real-time Compression in SAN Volume Controller and IBM Storwize V7000

IBM Redbooks Solution Guide

The IBM® Redbooks® publication, *IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines*, SG21-7521-03, covers many aspects of implementing compression. This IBM Redbooks Solution Guide complements that publication, provides a detailed review of various workloads and use cases, and serves as a best practices guide for IT professionals who have the tasks of sizing, planning, and implementing compression in IBM SAN Volume Controller and IBM Storwize® V7000. Its target audience also includes customers, IBM, and IBM Business Partners.

Use the following process when you plan or implement a project with IBM Real-time Compression™. Step-by-step implementation instructions are beyond the scope of this solution guide; those details are covered in other IBM Redbooks publications and product documentation. The intention of this solution guide is to provide high-level guidelines and considerations of an implementation that can lead to successful implementation of the solution itself.

Figure 1 shows the IBM System Storage® SAN Volume Controller.



Figure 1. IBM SAN Volume Controller

Did you know?

IBM System Storage SAN Volume Controller, a member of the IBM Storwize family, is a storage virtualization system with a single point of control for storage resources. SAN Volume Controller improves business application availability and greater resource utilization so you can get the most from your storage resources, and achieve a simpler, more scalable, and cost efficient IT infrastructure. The newly enhanced SAN Volume Controller with IBM Real-time Compression offers up to three times the throughput for applications.

Business value highlights and guarantee

One of the highlights is the ability to double the effective storage capacity of your active data - guaranteed. IBM guarantees storage savings of at least 50 percent with Real-time Compression and the Storwize V7000.

Today, thousands of customers enjoy the extraordinary storage efficiency features of IBM Storwize V7000. And now that efficiency comes with a guarantee. That is correct; IBM now guarantees you a savings rate of at least 50% when deploying IBM Real-time Compression. Just turn on Real-time Compression on your new Storwize V7000 and, in the unlikely event you do not get 50% or greater compression savings, IBM will provide, at no cost, the hardware and authorizations to match the shortfall. That is a deal you cannot refuse!

Unlike other compression implementations, Real-time Compression is specifically designed to compress active primary workloads such as Oracle RDBMS, IBM DB2®, Microsoft SQL Server, and VMware. IBM is so confident that it is willing to guarantee storage savings for Storwize V7000 customers.

For more information about the guarantee, go to the following address:
<https://ibm.biz/BdDZE6>

Solution implementation overview and determining suitability for data compression

Before implementing this solution, use the Comprestimator utility. This host-based utility offers fast and accurate estimates of an expected compression rate for block devices. Use Comprestimator to estimate the compression ratio, and implement compression on volumes that have a compression ratio above 25%, and evaluate workload on volumes with low compression ratio (<25%).

The latest version of Comprestimator can be downloaded from the following location:
<http://www-304.ibm.com/support/customer/sasff/comprestimator/home.html>

As a general guideline, compress workloads that show more than 25% compression savings. Use the threshold for volume compressibility in Table 1 to help you decide whether to compress a volume.

Table 1 recommendations are based on the data compression rate.

Table 1. Recommendations

Data compression rate	Recommendation
Higher than 25% savings	Use compression
Below 25% compression savings	Evaluate workload with compression

After verifying the expected compression rate using Comprestimator, the workload for a volume should be reviewed, to decide whether it is suitable for compression.

The following common workloads are typically suitable for use with compression:

- Database applications: Oracle, IBM DB2, Microsoft SQL, SAP
- Server/Desktop virtualization: VMware, KVM, Hyper-V
- Messaging: IBM Notes, Microsoft Exchange, (as long as attachments are not compressed file types)
- Others: Engineering, collaboration, seismic

Solution architecture, ensuring the system has available CPU resource, and balancing the system

Version 7.3 supports the new hardware models - IBM Storwize V7000 Gen2 (2076-524) and SAN Volume Controller (2145-DH8). These models deliver outstanding performance for all kinds of workloads and are recommended when using Real-time Compression.

The following hardware configurations are highly recommended with Real-time Compression:

- Storwize V7000 Gen2 with 64GB and 2 x Compression Accelerator Cards (per canister)
- SAN Volume Controller (2145-DH8) with 1 or 2 Compression Accelerator Cards (per node)

*The basic configuration of Storwize V7000 Gen2 with 32GB (per canister) and a single onboard Compression Accelerator Card is also supported with Real-time Compression, but will not deliver the same performance levels as the recommended configuration with the extra RAM and Compression Acceleration Card.

If the configuration uses SAN Volume Controller, to preserve native SAN Volume Controller performance for non-compressed volumes and optimize Real-time Compression performance, order RPQ 8S1296 for each 2145-CG8 node in an I/O group running Real-time Compression or upgrade the hardware to the latest SAN Volume Controller 2145-DH8 nodes.

System resource allocation on Storwize V7000 and SAN Volume Controller changes upon creation of the first compressed volume in an I/O group. Be sure to review your existing system CPU resources before creating a compressed volume in an I/O group. While reviewing system resources, make sure the CPU utilization is lower for the respective model, as shown in Table 2.

Table 2. CPU resource consideration

Per-node	SAN Volume Controller CF8*	SAN Volume Controller CG8	SAN Volume Controller CG8 Dual-CPU (RPQ 8S1296)	Storwize V7000* (2076-112, 2076-124)	Storwize V7000 Gen2 (2076-524)	SAN Volume Controller (2145-DH8) Dual CPU
Processor close to, or above	25%	50%	No consideration necessary	25%	50%	No consideration necessary

* In most cases, CPU utilization is lower than 25% when compression is disabled.

Configure a balanced system

Do not configure only a few compressed volumes in an I/O group. This guideline is more applicable to systems with more than a single I/O group. For example, consider the following 4-node Storwize V7000 (two I/O groups) with 200 volumes:

- iogrp0: Nodes 1 and 2; 180 compressed volumes
- iogrp1: Nodes 3 and 4; 20 compressed volumes

This setup is not ideal because CPU and memory resources are dedicated for compression use in all four nodes; however, in nodes 3 and 4, this allocation is used only for servicing 20 volumes out of 200 compressed volumes.

The following two alternatives are expressed in Table 3 (X is "a number of"):

- Alternative 1: migrate all compressed volumes from iogrp1 to iogrp0.
- Alternative 2: migrate compressed volumes from iogrp0 to iogrp1 and load balance across the nodes.

Table 3. Alternative setups

	Node 1	Node 2	Node 3	Node 4
Original setup	90 compressed volumes X non-compressed volumes	90 compressed volumes X non-compressed volumes	10 compressed volumes X non-compressed volumes	10 compressed volumes X non-compressed volumes
Alternative 1	100 compressed volumes X non-compressed volumes	100 compressed volumes X non-compressed volumes	X non-compressed volumes	X non-compressed volumes
Alternative 2	50 compressed volumes X non-compressed volumes	50 compressed volumes X non-compressed volumes	50 compressed volumes X non-compressed volumes	50 compressed volumes X non-compressed volumes

Usage scenarios and best practices

Various scenarios and best practices are described next.

Converting non-compressed volumes to compressed volumes

Use volume mirroring to convert a volume from non-compressed to compressed (or vice versa).

- Use a maximum sync rate setting (100) if possible.
- Mirroring many volumes concurrently increases Compression CPU utilization. Mirroring to compressed volume is used as the method to convert from non-compressed to compressed. Do not do all conversions at once. Consider staging only a few mirrors at a time.

Evaluating performance for volumes targeted for compression

If the compression savings are in below 25%, evaluate the performance of the volume using the following steps:

1. Add a second compressed copy to the volume.
2. Mark the new compressed copy as primary.
3. Monitor the application performance and host performance to evaluate whether performance is acceptable. If the impact is unacceptable, delete the compressed copy to restore the original performance.

Converting an external MDisk to a compressed volume

An image mode volume provides a direct block-for-block translation from a managed disk (MDisk) to a volume with no virtualization. Image mode enables you to import external LUNs with existing data into the SAN Volume Controller or Storwize V7000 as-is. An image mode volume can then be compressed by adding a mirrored compressed volume copy in a different storage pool.

To import an external LUN to an image mode volume, use the following steps :

1. Map the external LUN to the SAN Volume Controller or Storwize V7000.
2. Create an empty MDisk Group:
`mkmdiskgrp -name import_mdg`

3. Create an image mode volume in the MDisk Group you created in step 2. In the command, specify the external LUN as the MDisk. The volume should be the same size of the original LUN:

```
mkvdisk -vtype image -mdisk 13 -mdiskgrp import_mdg -size 1 -unit tb -rsize auto -iogrp0
```
4. Now you can mirror this volume to a compressed volume:

```
addvdiskcopy -compressed -mdiskgrp striped -rsize 2% -autoexpand
```

FlashCopy

Note the following information:

- Grainsize: the IBM FlashCopy® default grainsize is 256 KB for non-compressed volumes, and 64 KB for compressed volumes (this is a new default from software versions 6.4.1.5 and 7.1.0.1). The preference is to use the default grainsize for compressed volumes.
- Consider using the background copy method: The two ways to use FlashCopy are with or without background copy. Without background copy, the host I/O is pending until the "split" event is finished. For example, if the host sends a 4 KB write, this I/O will wait until the corresponding grain (64 KB or 256 KB) is read and decompressed, and then it is written to the FlashCopy target copy. This adds latency to every I/O. With background copy, all the grains are copied to the FlashCopy target immediately after the FlashCopy mapping is created. This adds latency during the copy; however, it eliminates latency after the copy is complete.

Easy Tier

IBM Easy Tier® is supported with compressed volumes starting with software version 7.1.0.1. The performance improvement achieved with Easy Tier and compression is up to three times faster in terms of application response time.

Easy Tier with compressed volumes tracks only read operations and migrates hot-extents to SSDs. Therefore, a good approach is to enable Easy Tier on compressed volumes with high read workloads. If, for any reason, Easy Tier must be disabled on a certain volume, disable it by using the following command:

```
svctask chvdisk -easytier off volume_name
```

The performance improvement achieved with Easy Tier and compression is up to three times faster in terms of application response time by having 5% of SSDs in the configuration. Throughput (maximum IOPS) is dependent on compression CPU utilization; therefore, in most cases it remains the same.

Figure 2 shows the test results of a TPC-C benchmark on compressed volume with Easy Tier enabled and disabled. The TPC-C benchmark was used with an Oracle database and represents a realistic online transaction processing (OLTP) workload.

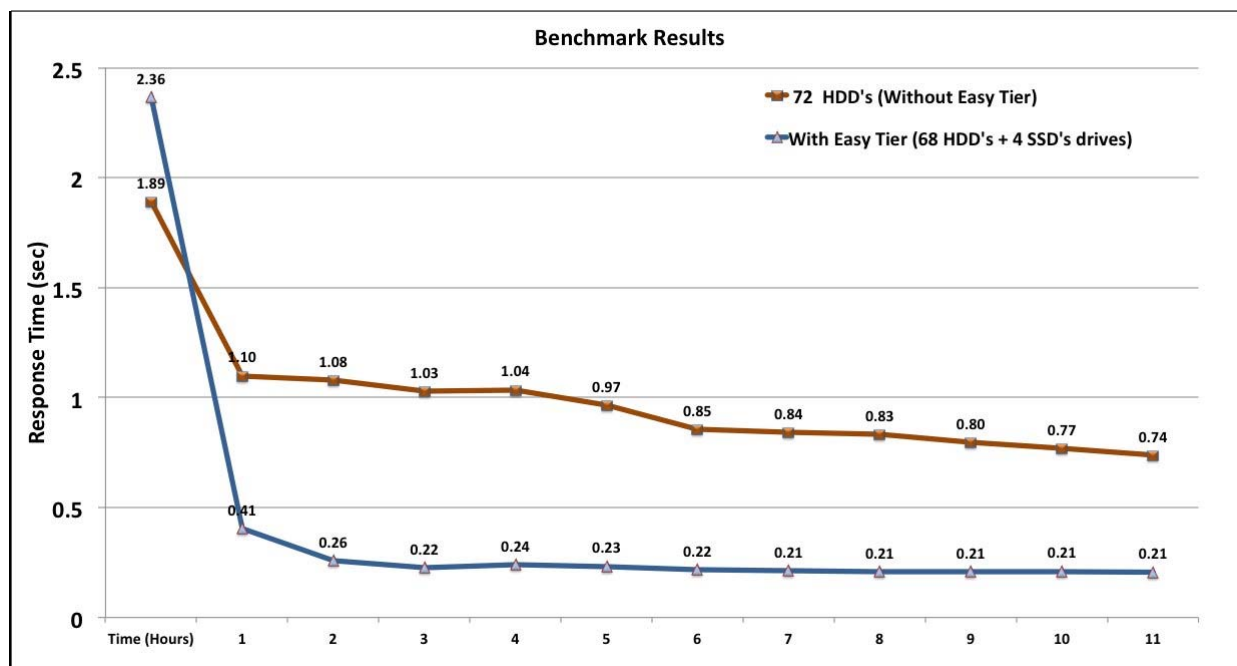


Figure 2. TPC-C benchmark results

The test results show that application response time became faster by more than three times when the configuration used SSDs with Easy Tier, compared to a similar configuration without SSDs and Easy Tier. Storwize V7000 system used in the benchmark was running software version 7.1.0.1 and was using the disk configuration without and with Easy Tier.

- Without Easy Tier:
 - 72 x 300 GB SAS HDDs
- With Easy Tier:
 - 68 x 300 GB SAS HDDs
 - 4 x 300 GB SAS SSDs

Supported software versions

IBM Real-time Compression in Storwize System family was first introduced in version 6.4.0. However, version 7.2 delivers significant performance improvements by software only. It is highly recommended to upgrade to the most recent PTF available in IBM Fix Central. Version 7.3 release contains support for the new hardware - Storwize V7000 Gen2 (2076-524) and SAN Volume Controller (2145-DH8) with Compression Accelerator Cards. For maximum performance, use the latest hardware models with software release Version 7.3.

Related information

For more information, see the following documents:

- *Real-time Compression in SAN Volume Controller and Storwize V7000*, REDP-4859
<http://www.redbooks.ibm.com/abstracts/redp4859.html?Open>
- *IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines*, SG24-7521-03 <http://www.redbooks.ibm.com/abstracts/sg247521.html?Open>

- IBM System Storage SAN Volume Controller
<http://www.ibm.com/systems/storage/software/virtualization/svc>
- IBM Offering Information page (announcement letters and sales manuals):
http://www.ibm.com/common/ssi/index.wss?request_locale=en

On this page, enter `Real-time Compression`, select the information type, and then click **Search**.
On the next page, narrow your search results by geography and language.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service. IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you. This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk. IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

© Copyright International Business Machines Corporation 2013-2015. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

This document was created or updated on May 5, 2015.

Send us your comments in one of the following ways:

- Use the online **Contact us** review form found at:
ibm.com/redbooks
- Send your comments in an e-mail to:
redbooks@us.ibm.com
- Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.

This document is available online at <http://www.ibm.com/redbooks/abstracts/tips1083.html> .

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>. The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

DB2®
Easy Tier®
FlashCopy®
IBM®
Real-time Compression™
Redbooks®
Redbooks (logo)®
Storwize®
System Storage®

The following terms are trademarks of other companies:

Microsoft, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.