



Implementing IBM Easy Tier with IBM Real-time Compression

IBM Redbooks Solution Guide

Overview

IBM® Easy Tier® is a performance function that automatically and non-disruptively migrates frequently accessed data from magnetic media to solid-state drives (SSDs). In that way, the most frequently accessed data is stored on the fastest storage tier, and the overall performance is improved.

How does it work

Every volume is split into logical units called *extents*. Easy Tier is based on algorithms that are developed by IBM Research, which evaluates the access frequency of each extent. Each extent is rated according to the number of I/Os going to that extent. Extents with a high rating, receiving the most I/Os, are marked as “hot” extents and become candidates for migration to SSDs in the same storage pool. Periodically, but no greater than 24 hours, a migration plan is created according to the “heat” of the extents and the data is migrated to the SSD MDisk. When the SSD becomes full, and there is a hotter extent to move onto the SSD, the “cooled” extents are migrated back to the lower-tiered MDisk.

Migrations are typically minimal, and add up to a maximum of two terabytes of data per day. The number of host read and write operations to a specific extent determines the rating of the extents. Only I/Os smaller than 64 KB are considered when determining ‘heat’ to prevent sequential I/O patterns from filling up the SSDs with data that is not likely to be accessed again frequently.

For more information about Easy Tier, see Chapter 7, “Easy Tier”, in *Implementing the IBM System Storage SAN Volume Controller V6.3*, SG24-7933, found at <http://www.redbooks.ibm.com/redbooks/pdfs/sg247933.pdf>.

Easy Tier with compressed volumes

IBM Real-time Compression™ software is embedded in IBM Storwize® V7000 and IBM SAN Volume Controller systems. Compressed volumes have a unique write pattern to the MDisks. When a host writes data to a certain offset in a compressed volume, the system compresses this data, which is then written to another offset of the underlying volume as it is represented in the storage pool. Such a change in offsets triggers unnecessary migrations of data into SSDs because repetitive writes to the same logical offset end up written in various locations instead. A new Easy Tier algorithm is therefore required to support compression.

What is new in Storwize V7000 and SAN Volume Controller V7.1

Starting with Version 7.1, Easy Tier supports compressed volumes. A new algorithm is implemented to monitor *read* operations on compressed volumes instead of reads and writes. The extents with the highest number of read operations that are smaller than 64 KB are migrated to SSD MDisks. As a result, frequently read areas of the compressed volumes are serviced from SSDs. Easy Tier on non-compressed volumes operates as before and it is based on read and write operations smaller than 64 KB.

Performance results

The performance improvement that is achieved with Easy Tier and compression has an up to 3x faster application response time by having 5% of SSDs in the configuration. Throughput (maximum IOPS) depends on compression processor usage; therefore, in most cases, throughput remains the same.

Figure 1 shows the test results of a Transaction Processing Performance Council benchmark C (TPC-C) on a compressed volume with Easy Tier enabled and disabled. The TPC-C was used with an Oracle database and represents a realistic Online Transaction Processing (OLTP) workload. (For more information about TPC-C, go to <http://www.tpc.org/tpcc/default.asp>.)

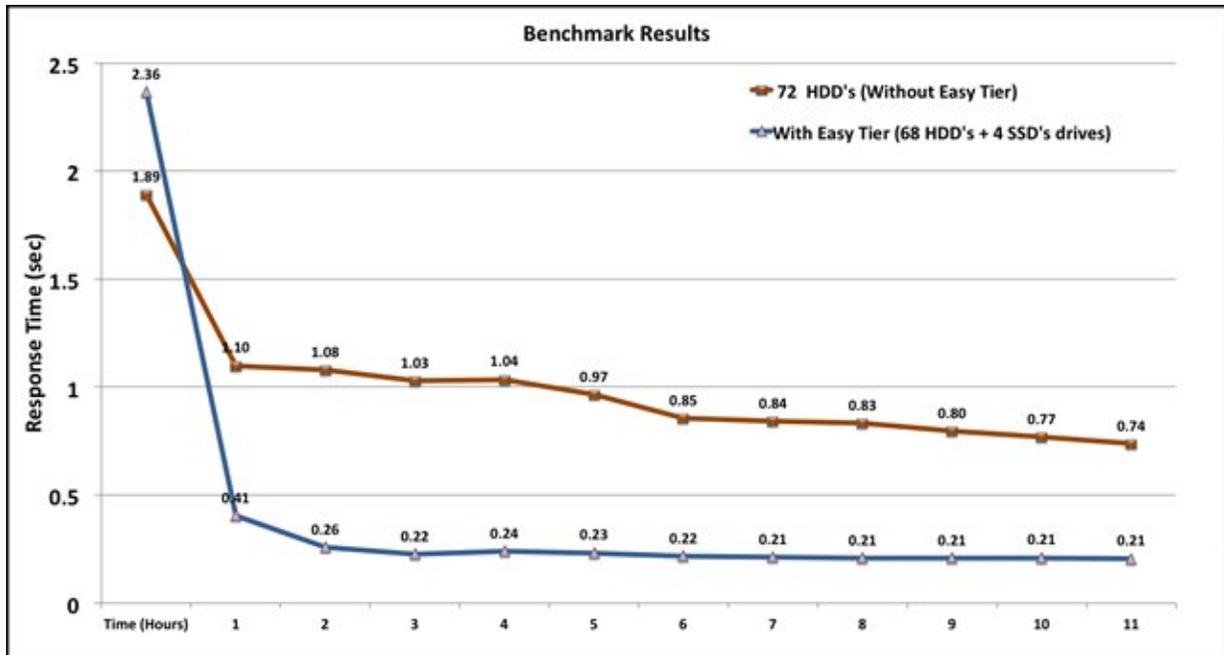


Figure 1. Benchmark results

Test results show that the application response time became faster by more than 3x when the configuration used SSDs with Easy Tier, compared to a similar configuration *without* SSDs and Easy Tier.

The Storwize V7000 system that was used in the benchmark was running software Version 7.1.0.1 and was using the following disk configuration:

Without Easy Tier:

- 72 x 300 GB SAS HDDs

With Easy Tier:

- 68 x 300 GB SAS HDDs
- 4 x 300 GB SAS SSDs

IBM Storage Tier Advisor Tool (STAT): A tool to monitor Easy Tier

The IBM Storage Tier Advisor Tool (STAT) is a Windows console application that analyzes heat data files that are produced by Easy Tier and produces a graphical display of the amount of "hot" data per volume (with predictions about how additional SSD capacity could benefit the performance for the system) and per storage pool.

The tool is available at no additional cost and can be found at the following website:

<http://www-01.ibm.com/support/docview.wss?uid=ssg1S4000935>

To use the tool, you should use the `dpa_heat` file as a source file. The tool provides a report of volume heat distribution and recommendations.

To download the file, from the IBM Storwize V7000 GUI, navigate to the Settings icon in the left pane and click **Support**, as shown in Figure 2.



Figure 2. Support option

Click **Show full log listing...**, as shown in Figure 3.



Figure 3. Show full log listing option

Download the dpa_heat file from the list of files that is displayed, as shown in Figure 4.

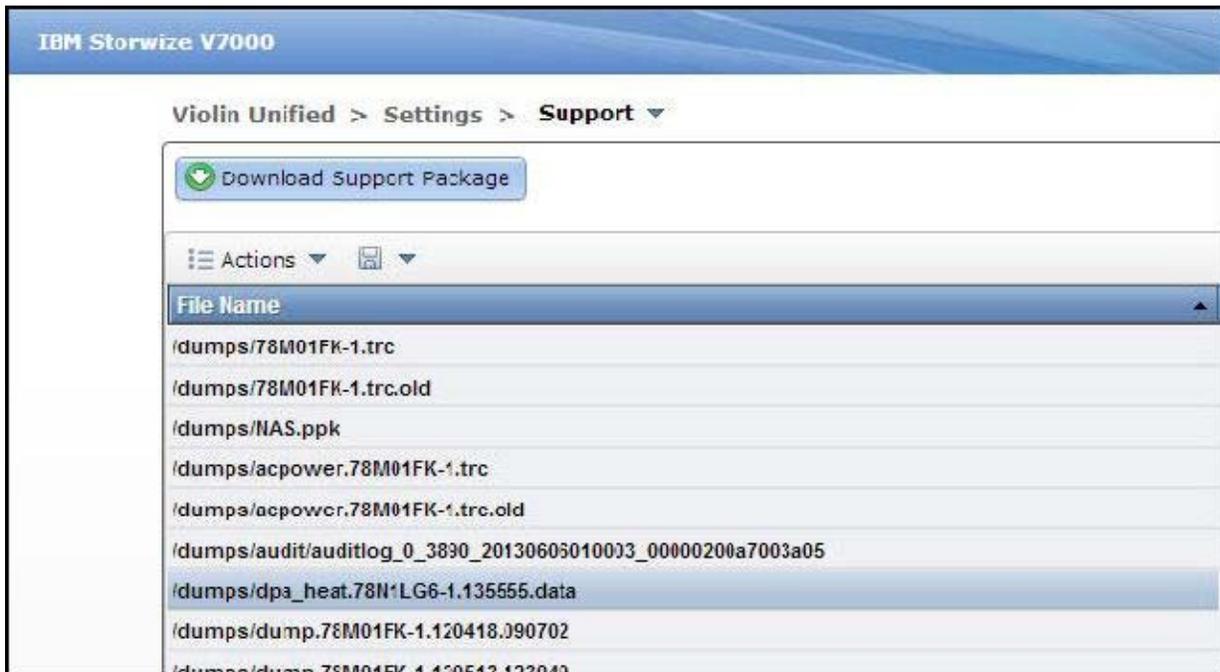


Figure 4. File selection

The dpa_heat file is also in the full support package.

Understanding the results

This section describes how to interpret the results.

Volume heat distribution

The *Volume Heat Distribution* report is useful for understanding the amount of capacity that is migrated to the SSD when Easy Tier is enabled. The heat areas of compressed volumes are reported based on read operations only. The non-compressed volume is based on reads and writes. The “hot” part of the volume is marked in red, as shown in Figure 5.

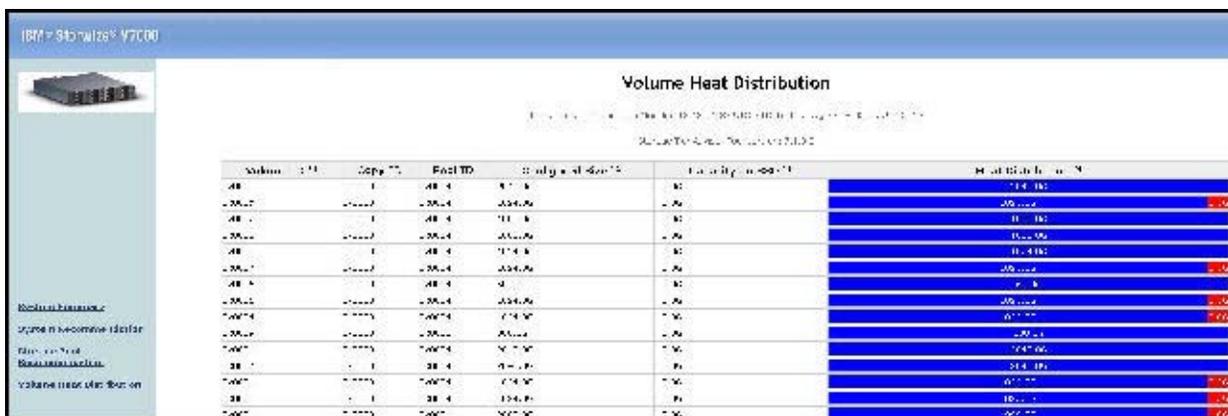


Figure 5. Volume Heat Distribution report

Note: The tool's recommendations are based on the state of the volume. Recommendations about generic volumes are based on both reads and writes, but compressed volumes are based only on read operations. Therefore, if you consider enabling Easy Tier on a compressed volume, first compress the volume and then use the STAT utility. Otherwise, the STAT tool's recommendations will be different from the actual results.

Performance improvement

The *system recommendation* and the *system pool recommendation* reports show the potential performance improvement in percentages according to the number of SSDs that are added.

Note: When compressed and non-compressed volumes are in the same storage pool, they might affect the predicted performance improvement results of the entire pool.

Compressed volumes are not directly supported by the STAT tool and therefore its recommendations will be inaccurate for compressed volumes.

Use the results there were obtained to estimate the performance improvement. Figure 6 shows the Storage Pool Recommendation.

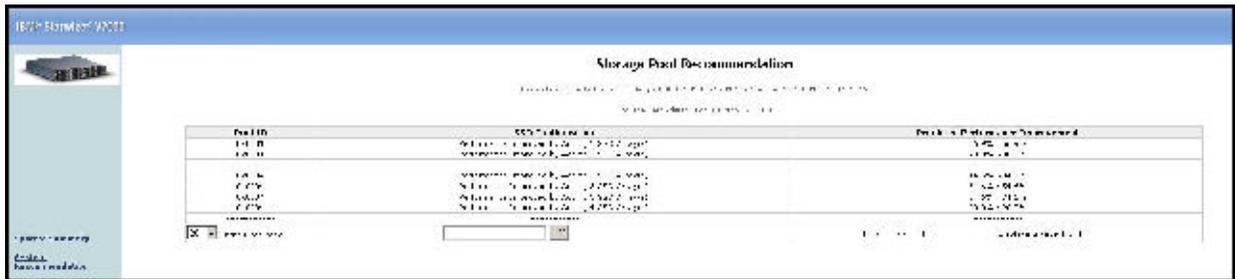


Figure 6. Storage Pool Recommendation

Configuration

Easy Tier is defined at a storage pool level and the algorithm runs on all the volumes in the pool. If Easy Tier must be disabled for a certain volume, you can disable it by running the following command-line interface (CLI) command:

```
svctask chvdisk -easytier off volume name
```

To configure Easy Tier, complete the following steps:

1. Create a storage pool with HDD MDisk.
2. Add an MDisk with SSD to the same pool.

Easy Tier is automatically turned on for pools with both SSD MDisks and HDD MDisks, so all the volumes in the pool have Easy Tier enabled. Figure 7 shows Easy Tier activated.



Figure 7. Easy Tier activated

Conclusion

As shown, Easy Tier with Real-time Compression can greatly improve read I/O activity response time. Therefore, you should enable Easy Tier with compression on volumes with a high read workload.

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. 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 August 22, 2013.

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:
redbook@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/tips1072.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:

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

The following terms are trademarks of other companies:

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Windows, 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.