HITACHI
DATA SYSTEMS

Enterprise Storage

Hitachi HiCommand® Tiered Storage Manager Software

Volume Migration Performance Report

A Performance Brief

By Rick Freedman

August 2006

Executive Summary

Hitachi HiCommand® Tiered Storage Manager data management software enables IT administrators to easily and interactively match application quality of service requirements to heterogeneous storage assets. Tiered Storage Manager software enables the nondisruptive movement of data volumes to match application-driven price, performance, and availability characteristics through an easy-to-use interface.

This performance brief is based on a series of tests run in a lab environment to understand the relative performance when migrating data from one storage tier to another. The observations may be useful for storage planning and to establish recommendations for data migration using Tiered Storage Manager software on the Hitachi TagmaStore® Universal Storage Platform model USP600.

These tests were limited and generated many more questions that will be addressed in additional tests and reports. However, the observations provided in this performance brief may provide some guidelines for consideration when using Tiered Storage Manager software in real-world environments.

Contents

Hitachi HiCommand® Tiered Storage Manager Software

Volume Migration Performance Report

A Performance Brief

By Rick Freedman

Overview

To help IT administrators to better serve business needs, Hitachi HiCommand® Tiered Storage Manager software makes it easy for them to match application quality of service requirements to heterogeneous storage assets—nondisruptively. But how does performance fare during the movement of data volumes? This performance brief considers results of a series of tests run on the Hitachi TagmaStore® Universal Storage Platform model USP600 in a lab environment, which assessed relative performance when migrating data from one storage tier to another.

Test Configuration

The test configuration for this series of tests included:

- One Universal Storage Platform
- Two Hitachi TagmaStore® Adaptable Modular Storage models AMS500 with both Fibre Channel and Serial-ATA (SATA) storage
- Multiple tiers that are configured based on
  - RAID configuration
  - Fibre Channel or SATA storage
  - Internal capacity on Universal Storage Platform or external capacity on AMS500

Objectives

The objectives of these tests, and the resulting performance paper, are to provide guidelines for planning tiered storage migrations as well as recommendations for optimizing migrations. In addition, these results will initialize a knowledge base for Hitachi Data Systems and its customers, which will provide information about factors that can help optimize data migration and minimize impact to applications.

Problems and questions to be addressed by this brief include:

- What sort of pre-planning do I need to do prior to migrating applications data?
- How long does it take to migrate a volume from A to B?
- How does a migration impact existing applications?

- What other factors affect migration time and system impact?
- What sort of recommendations/guidelines does Hitachi Data Systems have for the practical use of data migration on the Universal Storage Platform?

Those who may wish to look deeper into the performance of Tiered Storage Manager software may include: technical professionals seeking a greater understanding of its capabilities, such as storage administrators; organizations using the Universal Storage Platform or TagmaStore Network Storage Controller and considering Tiered Storage Manager software; and organizations considering the Universal Storage Platform or Network Storage Controller.

Testing Results

Volume Migration

Volume migration with Tiered Storage Manager software includes not only copying the data from the original source volume to the destination volume's location but also switching over the assignment of the LUN number so that I/O is directed to the volume's new location.

- As a general rule, the impact to application response time during a volume migration is comparable to that of maintaining a Hitachi ShadowImage™ In-System Replication software pair. The only additional impact is that after the volume copy operation has been completed, the pointers between the old and new LUNs are swapped.
- Because the physical location of the data has changed, once the LUN assignment is changed the previous cache entries are no longer valid, so the new data cache will be repopulated as I/O is processed. The greater the amount of cache used by the application, the longer the time it takes to fully populate the cache. Maximum I/O rates are often dependent upon cache hits; therefore, recovery to pre-migration I/O rate levels will be a function of cache repopulation and the responsiveness of the LUN's new storage location. In some test runs with large amounts of available cache, the I/O rates did not return to pre-migration levels for 15-20 minutes following the completion of the migration.

Volume Configuration and Layout

It was expected that the physical configuration of volumes would affect the time required to migrate volumes. A series of tests were run migrating to and from volumes of varying RAID configurations, as shown in Table 1 and Table 2. The first set of tests was to measure the time required to migrate data from various source volume types (or tiers) to the same target tier.

Table 1. Volume Migration to the Universal Storage Platform—146GB, 10K, RAID-6 (6D+2P)

Original Source Location			Migration Time (Minutes)	Relative Migration (Percent Longer)
System	**Disk Type**	**RAID Configuration**		
USP600	146GB 10K	RAID-5 – 7D+1P	29	—
USP600	146GB 10K	RAID-5 – 3D+1P	34	+17
USP600	146GB 10K	RAID-1+0 – 2D+2P	33	+14
AMS500	300GB 10K Fibre Channel	RAID-1+0 – 4D+4P	64	+121
AMS500	300GB 10K Fibre Channel	RAID-6 – 5D+2P	57	+97
AMS500	300GB 10K Fibre Channel	RAID-5 – 14D+1P	39	+34
AMS500	250GB 7.2K SATA	RAID-1+0 – 4D+4P	119	+310
AMS500	250GB 7.2K SATA	RAID-6 – 5D+2P	88	+203
AMS500	250GB 7.2K SATA	RAID-5 – 4D+1P	103	+255
AMS500	250GB 7.2K SATA	RAID-5 – 14D+1P	52	+79

(Total time to migrate three volumes, in parallel: 20GB, 50GB, and 100GB)

A second set of tests was run from the same volume to different target tiers.

Table 2. Volume Migration from Universal Storage Platform—146GB, 10K, RAID-6 (6D+2P)

Target Location			Migration Time (Minutes)	Relative Migration (Percent longer)
System	**Disk Type**	**RAID Configuration**		
USP600	146GB 10K	RAID-5 – 7D+1P	28	—
USP600	146GB 10K	RAID-5 – 3D+1P	48	+71
USP600	146GB 10K	RAID-1+0 – 2D+2P	63	+125
AMS500	300GB 10K Fibre Channel	RAID-1+0 – 4D+4P	53	+89
AMS500	300GB 10K Fibre Channel	RAID-6 – 5D+2P	44	+57
AMS500	300GB 10K Fibre Channel	RAID-5 – 14D+1P	53	+89
AMS500	250GB 7.2K SATA	RAID-1+0 – 4D+4P	209	+646
AMS500	250GB 7.2K SATA	RAID-6 – 5D+2P	164	+486
AMS500	250GB 7.2K SATA	RAID-5 – 4D+1P	189	+575
AMS500	250GB 7.2K SATA	RAI5 – 14D+1P	79	+182

(Total time to migrate three volumes, in parallel: 20GB, 50GB, and 100GB)

Note: The specific location of destination volumes on SATA parity groups can significantly affect migration times.

Consideration of volume layout can also help mitigate the impact (for example, by understanding the physical layout). If two destination volumes are on the same array group, migration will take much longer than if those volumes are located on different array groups.

It was observed that the greater the number of physical drives, the better the performance. Migration to RAID-5 groups with 7D+1P configuration ran up to 45 percent faster than migration to a RAID-5 group with 3D+1P (both on internal Universal Storage Platform drives). The number of copy tasks that are available to accomplish the migration may also affect this performance.

Migration to SATA Volumes

When testing SATA volumes on external storage, the difference in migration time between different RAID-5 configuration destinations was even greater. This is due to "verify-after-write" on SATA drives, which requires an additional read following each write, thereby compounding the difference in response times. Migrating to RAID-5 groups with 14D+1P was more than twice as fast as to a RAID-5 group with 4D+1P (both on external SATA drives). There is, however, a significant tradeoff in rebuild times in the event of a drive failure between the two configurations. Rebuild times will be much higher with a 14D+1P configuration than with 4D+1P. Note that these relative migration times are shown for comparison only; RAID-6 configurations are generally recommended over RAID-5 for SATA drives due to their extended recovery times during sparing.

Due to the performance characteristics, SATA drives are recommended for use with data that is accessed infrequently. Migrations to SATA drives should be planned when the time to complete the migration is not a concern. Review all available "best practices" concerning the use of SATA drives to fully design and deploy a SATA solution.[1]

Serial or Parallel Migration Tests

One significant feature of Tiered Storage Manager software is its capability to create Migration Groups and to initiate the migration of the entire group at the same time. A test series was set up to compare the time required to migrate multiple volumes serially and then in parallel. Four volumes were migrated individually (serially) and then again concurrently (parallel) as part of a Migration Group. In each test, the concurrent migration of volumes in a Migration Group completed in 25 percent to 35 percent less time than when migration was conducted serially, as shown in Table 3.

[1] You can find a solution "cookbook," "Hitachi TagmaStore® Universal Storage Platform: Configuring Externally Attached SATA Storage for Mainframe and Open Systems Archiving" at http://www.hds.com/tools/whitepapers.html?tres=usp

Table 3. Comparison of Sequential versus Concurrent Migration Times

Source Tier	→	Target Tier	Volume Size	Migration Time (minutes for each volume)	Total Migration Time
USP600: RAID-6 (6+2) 146GB, 10K	→	USP600: RAID-5 (3+1) 146GB, 10K	Migration Group		57:35
			20GB	8:21	
			50GB	13:42	
			100GB	19:18	
			200GB	33:35	
			Total Time for Migrations Performed Serially		1:14:56
USP600: RAID-5 (3+1) 146GB, 10K	→	USP600: RAID-6 (6+2) 146GB, 10K	Migration Group		42:46
			20GB	8:36	
			50GB	13:32	
			100GB	18:33	
			200GB	27:52	
			Total Time for Migrations Performed Serially		1:08:33

Impact of Application Load

Identical simulation-driven read-only application loads were run from two servers while the data volumes for each server were migrated (one at a time). This testing demonstrates the level of application impact by volume migration using Tiered Storage Manager software.

Impact of Application Load—with External Storage

The application load tests were performed migrating the data volumes between the Universal Storage Platform and externally attached storage. The data volumes were migrated in steps. First a migration from an internal Universal Storage Platform array group to an external volume on an AMS500 with Fibre Channel drives was performed. Then a migration to another volume on SATA drives in the same AMS500 was performed. The process was then repeated in the opposite order.

Application Load Summary

This suite of tests demonstrated that the migration of data volumes with Tiered Storage Manager software does affect the performance (or response time) of the application being migrated, and can also affect the performance of other applications that may be accessing data on either the source or the target array group.

While there are additional tests that these findings suggest, there are a number of observations worth noting:

- In each case there was some degradation of response time for the application while it was being migrated. In some cases there was a significant difference—application response times up to three to four times longer.

- The impact on application response time often showed a significant degradation at the final stage when the volume pointers are swapped and cache is invalidated. This degradation period has a short duration. Performance then improves as the cache starts to fill with ongoing data access.
- Volume migration can also affect performance of other applications accessing data on both the source and/or target array groups. The impact is generally more significant to applications sharing access to the migration target array group than those sharing access to the migration source array group.
- Application degradation appears to be magnified on slower tiers, particularly on external storage systems and very much so on external SATA volumes.

Conclusions and Recommendations

While HiCommand Tiered Storage Manager software supports nondisruptive migration of active data volumes, it is recommended that, whenever possible, migration of data volumes be planned during periods of lower data activity. This applies to activity on the volumes being migrated, as well as data activity on any of the source or target array groups affected.

Plan on an impact to I/O rates that will last longer than the actual migration. This is dependent on the amount of available cache as well as other factors.

Do not attempt to perform migrations when the write pending rate on the array is high. This will cause migrations to run for extended amounts of time. In this situation the I/O rate is minimally affected by the migration but it is already very low due to the elevated write pending rate. For example, as shown in Table 4, two tests illustrate differences in migration completion times when write pending is an issue.

Table 4. Migration Completion Times when Write Pending Is an Issue

Test 1: IOMeter 32K, 100% read, 0% write, 50% sequential, 50% random		
Start Time	**End Time**	**Elapsed Time**
10:19:43	10:30:02	10:19

Test 2: IOMeter 32K,100% read, 0% write, 50% sequential, 50% random (Test ran with second SQLIO load running to create write pending of 66%.)		
Start Time	**End Time**	**Elapsed Time**
16:57:36	20:26:30	3:28:54

Understanding the physical layout of volumes and array groups will aid in planning volume migrations and minimizing the impact to the storage systems and applications.

Important Note: Do not perform migrations concurrently targeting the same array group. This is especially critical on SATA drives.

All of the impacts of volume migration will be magnified when migrating to SATA volumes. Most of this can be attributed to the "verify-after-write" which requires multiple disk accesses for each write operation. Allow for longer migration times when moving to SATA drives.

As expected, the completion of these tests has led to initial conclusions and even more questions. Many of the observations have suggested areas of additional tests that will be further investigated in future performance briefs.

Appendix: Volume Migration Test Environment

The Testing Environment Consists of the Following Hardware

- One USP600 model
- Two AMS500 models
- Six Microsoft Windows 2003 Servers
  - Domain Controller
  - Management Server
  - File Server—for IOMeter testing
  - Exchange Server—for LoadSim testing
  - Two SQL Servers—for SQLIO testing and data mining testing
- Four Load Generators for LoadSim testing
- SAN connections were made through a Brocade® SilkWorm® 24000

USP600 Configuration Details

- 32GB cache
- 32 ports
- 128 146GB disks

AMS500 Configuration Details

AMS500 DF700M_75010286

- 4GB cache
- Four Fibre Channel ports
- 15 300GB Fibre Channel disks
- 30 250GB SATA disks

AMS500 DF700M_75010287

- 4GB cache
- Four Fibre Channel ports
- 15 300GB Fibre Channel disks
- 30 250GB SATA disks

Hitachi Data Systems Corporation

Corporate Headquarters
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
www.hds.com
info@hds.com

Asia Pacific and Americas
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
info@hds.com

Europe Headquarters
Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
Phone: + 44 (0)1753 618000
info.eu@hds.com

WHP-225-00 August 2006