

EMC Tiered Storage for Microsoft SQL Server 2008 Enabled by EMC Unified Storage and EMC Fully Automated Storage Tiering (FAST)

An Architectural Overview

EMC Information Infrastructure Solutions

Abstract

This white paper describes how EMC® Fully Automated Storage Tiering (FAST) technology can improve performance and resource utilization while lowering the overall costs for a Microsoft SQL Server 2008 database environment. EMC's FAST is an intelligent data rebalancing technology that monitors the activity of the most demanding applications and automatically rebalances them across two or more storage tiers that can utilize Enterprise Flash Drives, Fibre Channel drives, and SATA drives.

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

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

Part number: H7208.1

Table of Contents

Executive summary.....	5
Business case	5
Product overview.....	5
Key results.....	6
Introduction	7
Introduction to this white paper	7
Purpose.....	7
Scope	7
Not in scope	7
Audience	8
Terminology.....	8
Overview of components.....	10
Introduction to components.....	10
EMC CLARiiON CX4-960	10
EMC CLARiiON FAST	10
EMC Unisphere	11
Configuration.....	12
Overview of the physical components.....	12
Physical environment	12
Test configuration.....	12
Hardware resources	13
Software resources	14
Storage design considerations.....	15
Introduction to storage design considerations	15
Manual tiering.....	15
With EMC FAST	16
Tier preferences	17
Sub-LUN tiering.....	19
Implementation.....	21
Introduction to implementation	21
Creating a storage pool.....	21
Expanding a storage pool	23
Testing.....	26
Introduction to testing	26
Initial layout FC.....	26
FC baseline	28
Layout with Flash and FC	28
Flash and FC test.....	29

Post relocation.....	30
Test results.....	31
Comparison of manual against FAST	34
Conclusion	35
Summary.....	35
Findings.....	36
Next steps	36
References.....	36
White papers	36

Executive summary

Business case Performance is one of the top concerns for companies that rely on critical Microsoft applications and databases like Microsoft Exchange Server, SQL Server, and SharePoint.

As more companies move towards hyper-consolidated virtualized data centers and the private cloud, it is even more critical to look for intelligent infrastructure that can dynamically optimize performance while reducing overall costs.

Until now, the best practices for optimizing storage performance involved manual, resource-intensive processes. These processes can include tasks such as rebuilding RAID groups and LUNs or creating storage tiers—based on different storage types such as Flash drives and Serial Advanced Technology Attachment (SATA) drives—and then manually moving the data from one tier to another. Although typically these processes can take place while the application remains online, they are both manual and time-consuming.

The introduction of EMC's Fully Automated Storage Tiering (FAST) technology allows administrators to leverage an easy-to-use and potentially hands-off mechanism for optimizing the performance of their most demanding applications.

FAST is designed to **get the right data to the right place at the right time**. The introduction of FAST, with its ability to use ultra-performing Flash drives to fully optimize performance, is an exciting addition to EMC's midrange storage systems product line.

This white paper shows how FAST can improve performance and resource utilization while lowering the overall costs for a Microsoft SQL Server 2008 database environment. Automating the movement of data between storage tiers saves both time and resources. FAST eliminates the need to spend hours manually monitoring and analyzing data to determine a storage strategy, then maintaining, relocating, and migrating LUNs to the appropriate storage tiers.

Product overview

The purpose of this solution is to validate the use of EMC[®] FAST in a Microsoft SQL Server 2008 enterprise-class environment, supporting on-line transaction processing (OLTP) type workloads. It demonstrates the management capability of the EMC Unisphere[™] interface and shows the ease of use provided with the new storage management features available in the FLARE[®] software. This white paper also highlights how EMC FAST technology enables customers to **reduce costs while maintaining performance of demanding SQL databases**.

The solution demonstrates how an EMC CLARiiON[®] storage array with two tiers—Flash and Fibre Channel (FC)—can provide improved total cost of ownership (TCO) by validating the efficiency and granularity of FAST policies against the effort required for manual storage tiering, underlining the benefits of:

- Targeting data to Flash tiers to service high I/O workload requirements
- Enabling the movement of data through sub-LUN tiering, providing a level of reduced granularity not available with manual tiering strategies.

- Reducing administration effort through:
 - Automated analysis
 - Policy-based movement of data
 - Ability to relocate data through:
 - Manual activation
 - Scheduled relocations windows
- Providing a faster reaction to changes in application demand

This solution builds on the Microsoft Windows Server 2008 R2 Enterprise Edition environment hosting a Microsoft SQL Server 2008 physical cluster. For more information, refer to the EMC white paper: *EMC Tiered Storage for Microsoft SQL Server 2008 Enabled by EMC CLARiiON CX4 and Enterprise Flash Drives—A Detailed Review*.

Key results

This solution highlights the value of implementing FAST in a Microsoft SQL Server 2008 production environment. Administrators can create a storage pool on a CLARiiON CX™ series array and, after either migrating or relocating their database files to the pool LUN, can implement auto-tiering either as a manual or automated process to relocate data to the appropriate tiers at a sub-LUN block level.

It is no longer necessary to create separate LUNs on each tier; through pool LUNs FAST will spread data across multiple tiers, according to the statistics collected during the hourly polling cycles. Customers seeking database performance improvements due to hot data being in isolated portions of those SQL database files no longer need to use a large area of space on a faster tier of storage, but now can use **just** that area of space where they need to, **just** for their hot data portions.

When moving from a baseline of single-tier, 90 FC drives to a multi-tier that included 30 FC drives and four Flash drives, we reduced acquisition cost by 38 percent, reduced power/cooling cost by 45 percent, and saved 80 percent of management time. Furthermore, all of these savings were achieved while improving performance—tested IOPS were improved by 4.2 percent and tested TPS were improved by 2.4 percent. These results demonstrate that, by replacing 60 FC drives with four Flash drives, significant financial benefits can be achieved, as well as performance benefits provided.

<i>Configuration</i>	All Fibre Channel	Tiered Flash/FC
<i>Disks</i>	90 FC	30 FC / 4 Flash
<i>Tested TPS</i>	Baseline	2.4% Improvement
<i>Tested IOPs</i>	Baseline	4.2% Improvement
<i>Management</i>	Baseline	80% Less Time
<i>Acquisition Cost</i>	Baseline	38% Less Cost
<i>Power/Cooling</i>	Baseline	45% Less Cost

Introduction

Introduction to this white paper

This white paper introduces IT and database administrators to the new features of EMC CLARiiON FAST technology. It illustrates how, with the introduction of FAST into a SQL Server 2008 OLTP-type environment, the complexity of the storage environment can be greatly reduced, showing administrators how to gain improvements in storage array utilization and performance by adopting the data relocation recommendations that FAST suggests.

This white paper includes the following topics:

Topic	See Page
Overview of components	10
Configuration	12
Storage design considerations	15
Implementation	21
Testing	26
Conclusion	35
References	36

Purpose

The purpose of this white paper is to demonstrate the performance benefits of FAST in a given application workload, using Microsoft SQL Server 2008.

Scope

The scope of this white paper is to:

- Present an overview of the concepts and technologies in the solution.
 - Recommend how EMC CLARiiON FAST software should be used in conjunction with Microsoft SQL Server 2008.
 - Assess the storage performance impact of FAST on Microsoft SQL Server 2008.
 - Demonstrate the newest midrange storage management tool, EMC Unisphere.
 - Present the test results and consequent business benefits of the solution.
-

Not in scope

This white paper does not document the following:

- Supply and build of the physical environment
 - Installation or patching of Windows 2008 R2
 - Installation, patching, or Microsoft failover cluster configuration of Microsoft SQL Server 2008 SP1
-

The information in this paper is also not intended to replace existing, detailed product implementation guides or best practices. This solution does not provide a comprehensive guide to every aspect of a FAST-enabled solution.

Audience

This white paper is intended for:

- Field personnel who are tasked with deploying EMC CLARiiON CX4 as the storage platform
- Customers, including IT planners, storage architects, and SQL database administrators
- EMC staff and partners, for guidance and the development of proposals

It is assumed that the reader is familiar with:

- Microsoft SQL Server
 - EMC CLARiiON storage
-

Terminology

This section defines terms used in this document.

Term	Definition
Bandwidth	The amount of data a storage system can process over time, which is measured in megabytes per second.
Disk Transfers/sec	Disk Transfers/sec is the rate of read and write operations on the disk.
Enterprise Flash drives (Flash)	Also known as solid state drives (SSD), Flash drives contain no moving parts, which removes the storage latency associated with traditional magnetic disk drives.
FAST	FAST—Fully Automated Storage Tiering—represents the next generation of storage tiering. FAST automates the movement and placement of data across storage resources as needs change over time. FAST enables continuous optimization of your applications by eliminating trade-offs between capacity and performance, while lowering cost and delivering higher service levels at the same time.
FLARE	EMC FLARE software is a component of the FLARE Operating Environment that controls the operation of the CLARiiON storage system. FLARE release 30 supports features for Unisphere, FAST, pools, and compression for CLARiiON CX4.
Logical unit number (LUN)	A unique identifier used to identify logical storage objects in a storage system.

Pool LUNs	<p>A pool LUN consists of two types, a <i>thin</i> LUN and a <i>thick</i> (fully provisioned) LUN:</p> <ul style="list-style-type: none"> • A thin LUN consumes physical storage from the storage pool only when data is written to it, and competes with other thin LUNs in the pool for available storage. • A fully provisioned thick LUN is a logical unit of storage created within a storage pool that consumes physical storage of the LUN size up front from the storage pool. You can shrink or expand a fully provisioned LUN.
RAID 1/0	RAID method that provides data integrity by mirroring data onto another disk. This RAID type provides the greatest assurance of data integrity at the greatest cost in disk space, and the greatest performance and throughput.
RAID 5	RAID method where data is striped across disks in large stripes. Parity information is stored so data can be reconstructed, if necessary. One disk can fail without data loss. Performance is good for reads, but slower for writes.
Response time	The interval of time between submitting an I/O request and receiving a response.
SP	Storage processor on a CLARiiON storage system. On a CLARiiON storage system, a circuit board with memory modules and control logic that manages the storage-system I/O between the host's FC adapters and the disk modules.
Storage pools	<p>EMC storage pools comprise:</p> <ul style="list-style-type: none"> • Homogeneous pools that contain a single drive type, that is, only Flash, FC, or SATA. Auto-tiering is not available. • Heterogeneous pools that contain mixed drive types of Flash, FC, and SATA. Only two different drive types are needed for FAST to enable auto-tiering.
Throughput	The number of individual I/Os the storage system can process over time, which is measured in I/Os per second.
Unisphere	EMC Unisphere software provides the next generation of storage management and presents a single, integrated, and simple web interface for the CLARiiON storage system.

Overview of components

Introduction to components

This section identifies and briefly describes the components deployed in the solution environment. The components used are:

- EMC CLARiiON CX4-960
 - EMC CLARiiON FAST
 - EMC Unisphere
 - Microsoft Windows Server 2008 R2
 - Microsoft SQL Server
-

EMC CLARiiON CX4-960

The CLARiiON CX4-960 system is a high-end, enterprise storage array comprising a system bay that includes storage processor enclosures (SPEs), SPs, disk array enclosures (DAEs), and separate storage bays that can scale up to 960 disk drives. The CX4-960 arrays support multiple drive technologies, including Flash, FC, and SATA drives, and the full range of RAID types.

EMC CLARiiON FAST

The new CLARiiON FAST features include:

Sub-LUN tiering

Provides a more granular analysis and movement of data, delivering the optimal use of storage tier capacity. For example, instead of moving a 1 TB LUN to Flash, FAST sub-LUN tiering identifies the most critical active data within the LUN and automatically moves only those hot spots in 1 GB chunks. The result is a far better use of the available storage capacity through efficient storage utilization as customers increase the performance value of their investment in Flash technology.

Automated, policy-based movement

Automatic analysis and policy-based movement of data free up administrators' time by having the storage system do the work for them. FAST policy options include:

- Auto-tiering based on I/O patterns analysis
- Highest tier for optimal performance
- Lowest tier for greater TCO optimization

Unisphere management

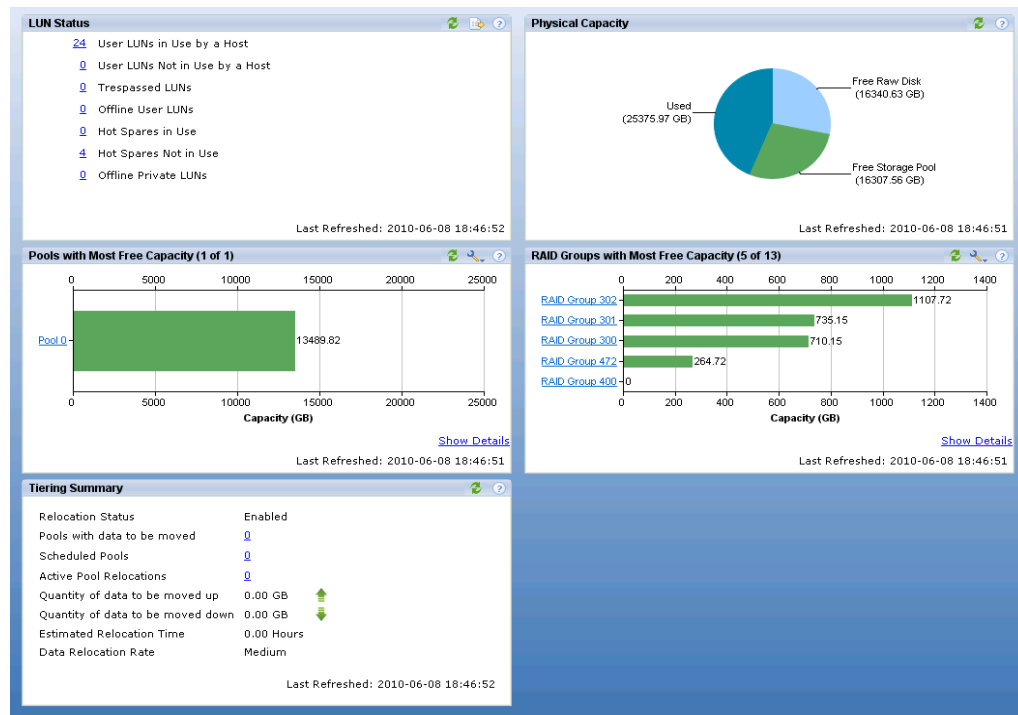
Integrated with the new Unisphere Manager, administrators can easily configure FAST on a per-LUN basis and schedule data movement to occur automatically as frequently or infrequently as needed.

EMC Unisphere EMC Unisphere provides a flexible, integrated experience for managing existing CLARiiON storage systems and next-generation EMC unified storage offerings in a single pane of glass. This new approach to midtier storage management fosters simplicity, flexibility, and automation. Unisphere's unprecedented ease of use is reflected in intuitive task-based controls, customizable dashboards, and single-click access to real-time support tools and online customer communities.

New Unisphere features include:

- **Task-based navigation and controls** that offer an intuitive, context-based approach to configuring storage, creating replicas, monitoring the environment, managing host connections, and accessing the Unisphere support ecosystem.
- **A self-service Unisphere support ecosystem**, accessible with one click from Unisphere, that provides users with quick access to real-time support tools, including live chat support, software downloads, product documentation, best practices, FAQs, online communities, ordering spares, and submitting service requests.
- **Customizable dashboard views and reporting** capabilities that enable at-a-glance management by automatically presenting users with valuable information in terms of how they manage their storage. For example, customers can develop custom reports up to 18 times faster with EMC Unisphere.
- **Common management** provides a single sign-on and integrated experience for managing CLARiiON and Celerra® platforms.

The following image provides an example of the Unisphere Summary page that gives administrators a wealth of detailed information on connected storage systems, from LUN pool and tiering summaries to physical capacity and RAID group information.



Configuration

Overview of the physical components

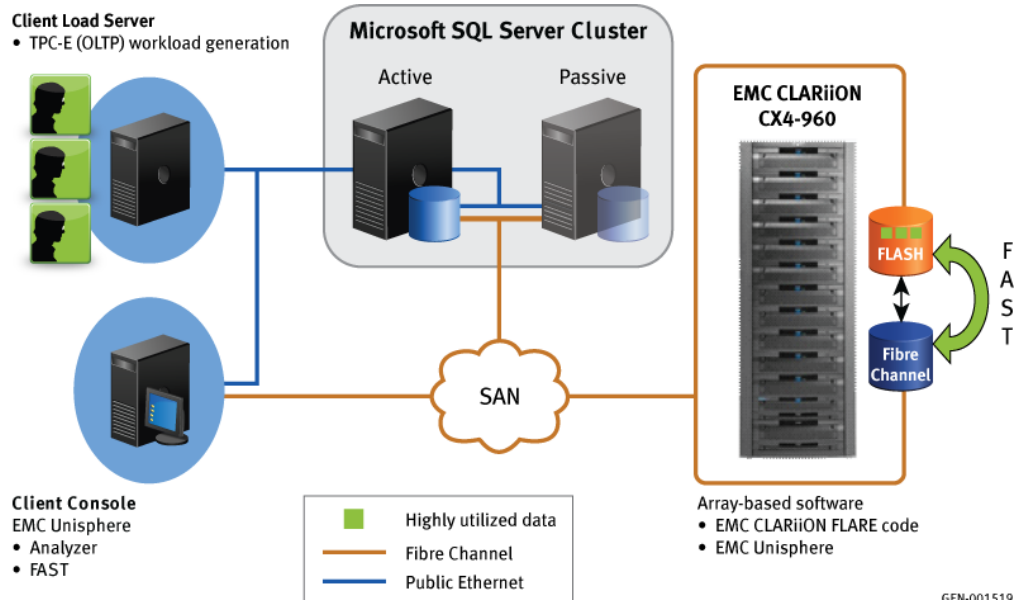
This tiered storage design consists of the following physical components:

- A two-node active/passive Microsoft Windows failover cluster
 - Storage connectivity provided by an 8 Gb/s FC switch
 - Network connectivity provided by a 1 GbE network switch
 - CLARiiON CX4-960 attached to servers through four front-end FC ports per SP
 - CLARiiON CX4-960 providing tiered storage (Fibre Channel and Flash disks)
-

Physical environment

The following diagram illustrates the overall physical architecture of the environment.

Test configuration



This section describes the test configuration used.

SQL Server 2008 test configuration

The SQL Server 2008 test configuration is based on the following profile:

- Number of SQL users supported: 75,000
- User data: 789 GB
- Expected throughput: 14,000 IOPS

SQL Server test application

The SQL Server load test tool used in this environment simulates an OLTP workload. It comprises a set of transactional operations designed to exercise system functionalities in a manner representative of a complex OLTP-type application environment.

OLTP workloads

The OLTP application used to generate the user load in this test environment is based on the TPC Benchmark-E (TPC-E) standard. TPC-E testing includes a set of transactions that represent the processing activities. The database schema, data, population, transactions and implementation rules have been designed to be broadly representative of modern OLTP systems. For this solution, the TPC-E application models the activity of a brokerage firm that:

- Manages customer accounts
- Executes customer trade orders
- Tracks customer activity with financial markets

Hardware resources

The hardware used to validate the solution is listed in the following table.

Equipment	Quantity	Configuration
EMC CLARiiON CX4-960	1	102 x 450 GB 15k FC disks 4 x 200 GB Flash disks
Fibre switch	2	FC 8 Gb switch 48-port
Network switch	1	Ethernet 1 Gb switch 48-port
Domain controller	1	2-core/4 GB RAM
TPC-E load servers	1	8-core physical server
Client console	1	8-core physical server
SQL Server - active	1	24-core/128 GB RAM
SQL Server - passive	1	24-core/128 GB RAM

Software resources

The software used to validate the solution is listed in the following table.

Software	Version
EMC FLARE	30
EMC PowerPath®	5.3
Microsoft Windows Server 2008	Enterprise Edition R2
Microsoft SQL Server 2008	Enterprise SP1 (clustered)

Storage design considerations

Introduction to storage design considerations

This section details the following storage design considerations:

- Manual tiering
 - With EMC FAST
 - Tier preferences
 - Sub-LUN tiering
-

Manual tiering

When considering implementing manual storage tiering for the CLARiiON CX4 series prior to FAST, administrators had to take into account the best practices and recommendations for:

- Microsoft SQL Server 2008
- Microsoft Windows Server 2008 R2
- EMC CLARiiON CX4 series
- RAID types

These best practices and recommendations are detailed in the white paper *EMC Tiered Storage for Microsoft SQL Server 2008 Enabled by EMC CLARiiON CX4 and Enterprise Flash Drives—A Detailed Review*.

While a good understanding of these best practices and recommendations is invaluable to any administrator, with the auto-tiering capabilities of FAST, it is the storage system that does the work.

With the introduction of Flash drives, and their obvious performance advantage over traditional spinning disks making them a first choice for use in the top tier, the latest challenge for administrators is now how to identify the most appropriate tier for their data to reside on, while at the same time making the optimal use of these new relatively expensive array resources.

Until now administrators have had to monitor workloads, then manually identify and migrate data to the correct tier, hoping that their allocation matched the business requirements. This is a challenging task and one that requires continuous intervention to maintain operational efficiency as data grows and performance requirements change.

Administrators would have to estimate I/O requirements for given workloads, and assign set numbers of disks to service these requirements, based not on capacity provided but on a configuration's ability to service given I/O requirements taking into account additional capacity for future growth. This leads to high numbers of disks being locked into servicing workloads that, in relation to the physical capacity of those disks, is disproportionate to the data footprint on the array.

This leaves administrators with the difficult task of explaining to senior management

why an array with multi-terabytes of seemingly available capacity is unable to be utilized without direct performance impact to service levels.

Ensuring that the best performing storage tier is used as effectively as possible on an ongoing basis is also both difficult and cumbersome at the file system (LUN) level. For example, on a 500 GB LUN, it would be inefficient to allocate 500 GB of the top tier when perhaps only a small percentage of the data within the LUN is hot.

To separate data into smaller, manageable pieces, administrators need to develop a strategy for partitioning data. This is made easier in SQL Server, which has the capability to partition database tables and their indexes over filegroups within a single database. Running a workload through the Database Tuning Advisor (DTA) results in recommendations being made for partitioning and also provides the code to do this.

Provided that it is done correctly, partitioning should always have a positive effect on the environment and most processes such as:

- Database backup and restore
- Index maintenance (rebuild), including index views
- Data management
- End-user workload
- Concurrency
 - Parallel partition query processing
 - Table partition lock-escalation
- Enhanced distribution or isolated database workloads using filegroups

Understanding all this technology is a challenge for an administrator and also a resource-intensive process. An administrator needs to have an excellent understanding of application behavior on the storage array before trying to optimize the I/O by selectively placing objects on separate spindles. Developing a strategy to cope with growing data sizes, partitioning, rebalancing data files, migrating LUNs, or transferring data to LUNs—all these tasks have an impact on array resources, service levels, and downtime.

Always consider that once introduced, complexity is hard to remove.

With EMC FAST With the introduction of FAST technology, administrators now have a method to alleviate the uncertainty of back-end storage design by giving them the ability to enable the EMC storage array to systematically tier data, based on I/O requirements. Once auto-tiering is enabled on the LUNs, EMC FAST technology continuously monitors and analyzes data workloads to generate the tiering recommendations to move colder (inactive) data to lower capacity optimized storage tiers and hotter (active) data to higher performing tiers.

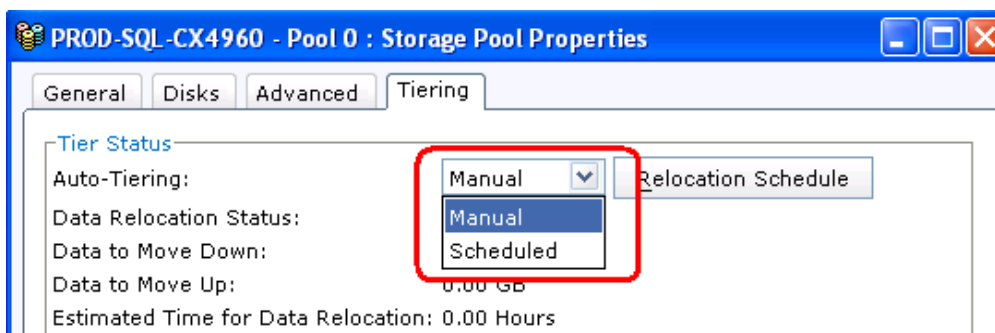
Not only is this tiering done at the LUN level but also at a very granular sub-LUN level. This ability greatly reduces provisioning uncertainty since data is moved according to the activity level and administrators are no longer locked into committing to a provisioning strategy that can quickly and unexpectedly change.

This ability to automatically locate data to the appropriate tier so that it is effectively

located **in the right place, at the right time**, is a major breakthrough in storage technology. Through investment in Flash and FAST technology, an administrator can now quickly show how workloads can be serviced through a smaller physical footprint on the array than traditional configurations required. This provides the additional benefits of lower investment, lower running costs and simplified administration, while maintaining or even increasing workload performance.

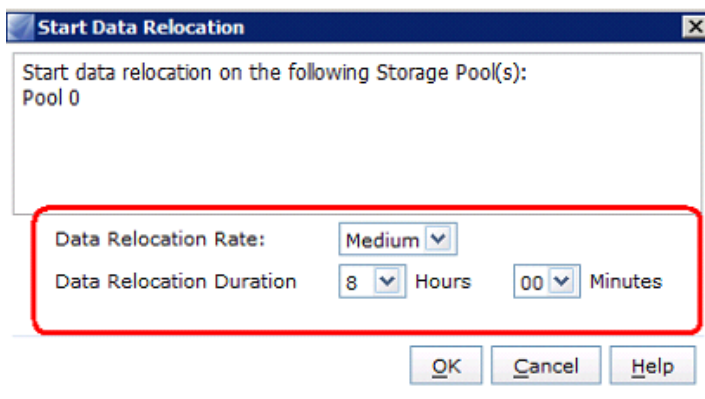
Tier preferences

Data relocation in FAST is governed by the global relocation setting within the **Tiering** tab of the Storage Pool Properties window, presenting the administrator with two options—manual or scheduled—as shown in the following image:



Manual

If the **Manual** option is selected, data relocation on the selected storage pool is initiated, and administrators can select the rate and the duration for the data relocation to complete as shown in the following image.

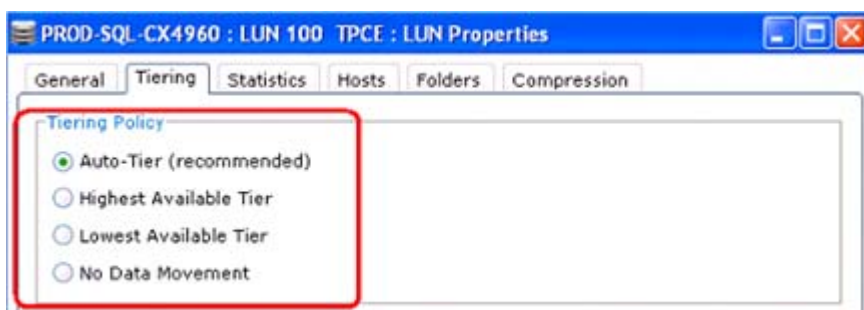


Scheduled

If the **Scheduled** option is selected as the auto-tiering policy on the storage pool, administrators can select the relocation rate, schedule the days on which relocation will occur, and set the start time and duration of the relocation. Data is moved in 1 GB segments to the appropriate tiers, depending on priorities FAST assigns each slice during the polling and analysis cycles, so as much critical data as possible is relocated over the given timeframe.

For example, if a large amount of data was available for relocation, then an administrator might schedule the relocation for a couple of hours each day so that by the end of the week data relocation has occurred for all the data.

FAST then continues to identify data that needs relocating and, by leaving this window open, administrators can ensure that FAST continues to tier the array automatically, based on the I/O activity of the LUNs within the pool that has been assigned the auto-tier policy. This is shown in the following image.



The following table details the available tiering policies.

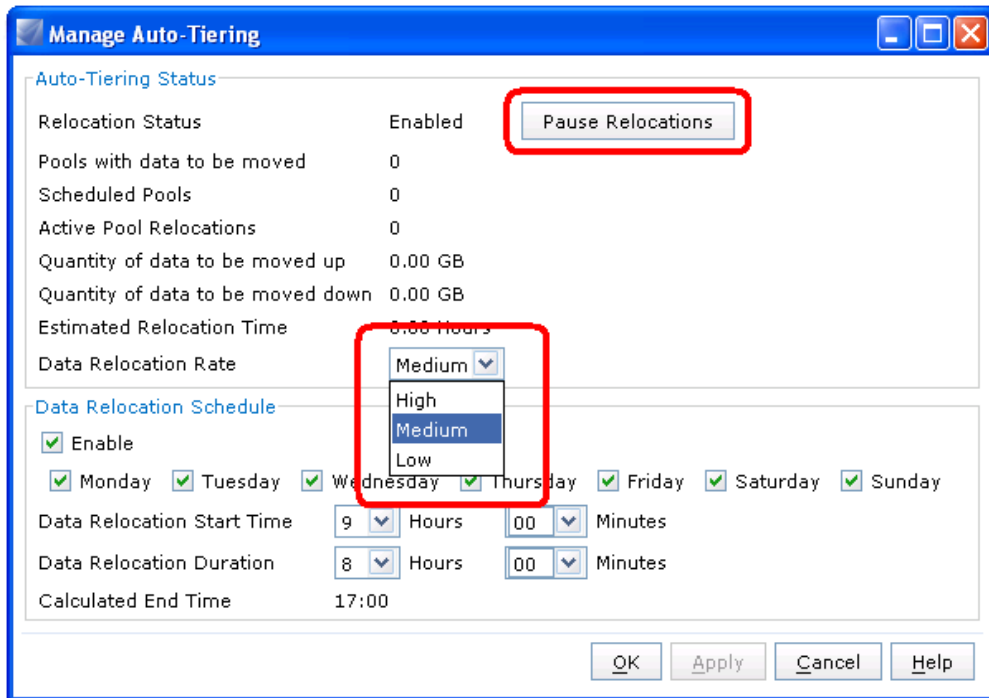
Policy	Description
Auto-Tier	Best for most users, adjusts the busiest data to the fastest tier.
Highest Available Tier	Sends the most critical data to the highest tier.
Lowest Available Tier	Sends the least performance-sensitive data to the lowest tier.
No Data Movement	Data is distributed evenly but is not moved after that.

Data relocation can occur at three different rates:

- High
- Medium
- Low

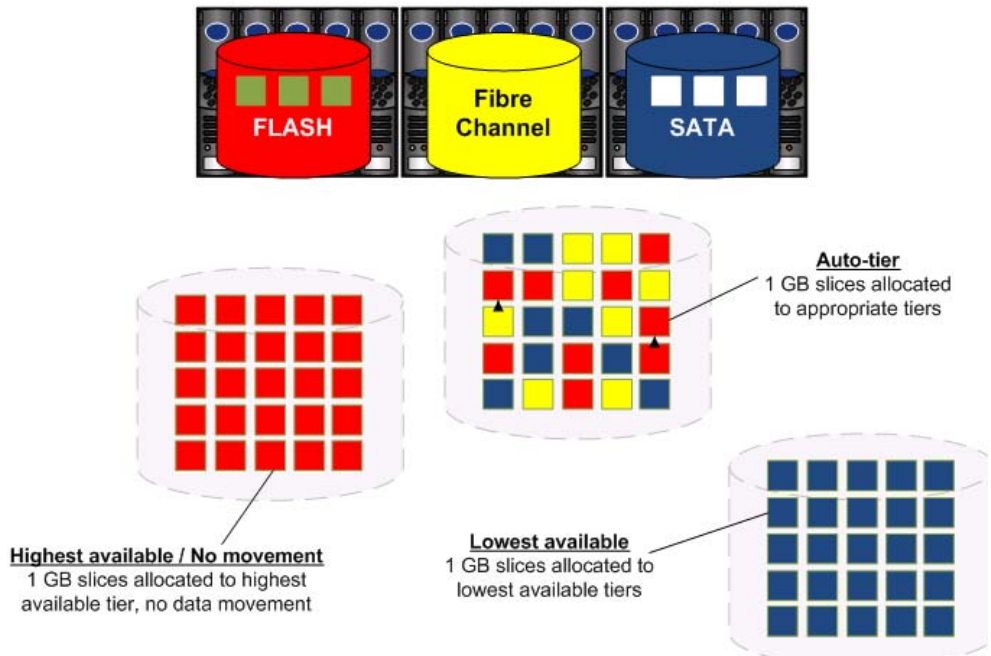
Relocation rates will differ between configurations and user workloads and administrators should set and monitor rates to find the most suitable setting.

Relocation can be started, paused, and stopped at any stage of the process, as shown in the following image.



Sub-LUN tiering

With EMC FAST data is moved and allocated in 1 GB slices, which is made possible through the introduction of sub-LUN tiering, one of the most exciting features of EMC's latest CLARiON FLARE code. This technology allows administrators to make the most of the space available to them on the high-performance Flash tier. FAST is able to profile the data on a LUN, and group this data into 1 GB segments, based on the I/O activity of particular chunks of data, as shown in the following image.



This functionality is made possible by wide striping data across multiple spindles. Most LUNs have varied I/O profiles, with areas of high and low activity that may not match the performance of the tier where they are located. With the option to create multiple tiers within the storage pool, FAST is able to allocate these segments to the appropriate tiers, the busiest segments being prioritized to the highest performing tiers, with the less busy segments going to the lower tiers, as space and priorities dictate.

This ability to spread a LUN across multiple tiers frees administrators from the previous restriction of having to move a whole LUN between tiers. Also, if administrators were confined to wide striping data across the FC and SATA drives only, they would need a large number of disks to maintain performance and, as a result, waste considerable storage space. With Flash allocated to the pool, the busiest segments can now be hosted on a small number of high-performance drives, while the hosting of data with lower I/O requirements can be met with fewer, lower-cost FC or SATA disks.

With FAST, administrators also have the option to lock data to its current allocation across the tiers, preventing further data relocation. This selection should only be made if the administrator has a clear understanding of the data's storage requirements, for example, locking critical data to the top tier that should always be available, or locking archived data to lower cost tiers for storage.

Implementation

Introduction to implementation

This section provides the following implementation details:

- Creating a storage pool
 - Homogeneous storage pools
 - Expanding a storage pool
 - Heterogeneous storage pools
 - Mixed-use storage pools
-

Creating a storage pool

With EMC FAST technology, storage provisioning is provided by storage pools, where disks of varying types are assigned to create a pool of disks, in numbers that match the configured RAID type assigned to the pool. FLARE 30 introduces a new type of pool LUN, fully allocated, whereby the LUN size is reserved with the pool. Pools can be thin (space-efficient) or thick (high-performance) and both LUN types can reside in the same pool.

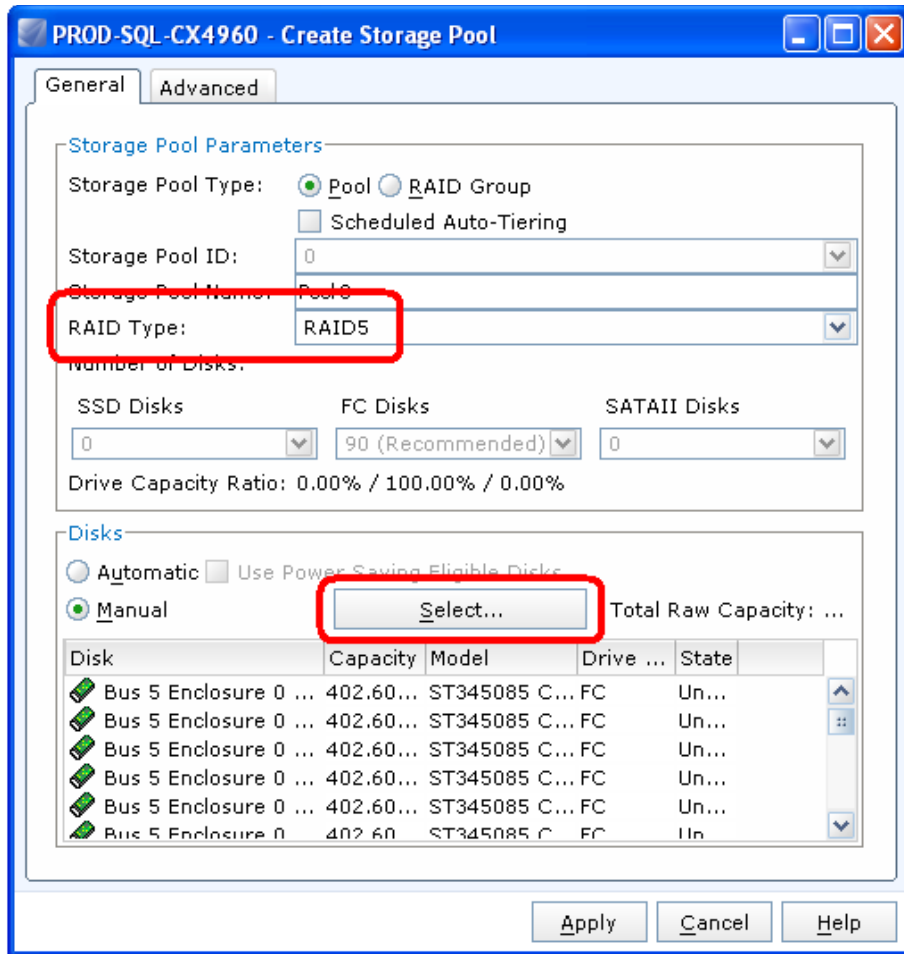
Traditional provisioning required the assessment of current storage requirements and then allocating space for future growth. This pre-allocation of resources can be wasteful and costly. EMC storage pools offer a more efficient way of providing for future growth, with new features in EMC FLARE that allow for pool expansions as requirements grow.

Storage pools are the new standard. Pools are easy to set up, they enable the distribution of the load over a collection of disks, and include the following features:

- Any disks, in any (back-end) location, except the **Vault Disks (0-4 in Bus 0)**
- Pools can be expanded
- Pool LUNs can be expanded
- Pool LUNs can be shrunk (supported by Windows 2008)
- Pool LUNs can be compressed
- Zero detection for more efficient storage (thin LUNs)


Homogeneous storage pools

To baseline the testing environment, a RAID 5 homogeneous storage pool was created in Unisphere, that is, where a single type of disk is assigned to the pool, as shown in the following image. In this instance, 90 FC disks were assigned. Homogeneous pools allow an administrator to provision storage and take advantage of auto-tiering at a later date by adding additional drive types.



The following image illustrates the layout of homogeneous storage pools.

Homogeneous Storage Pools

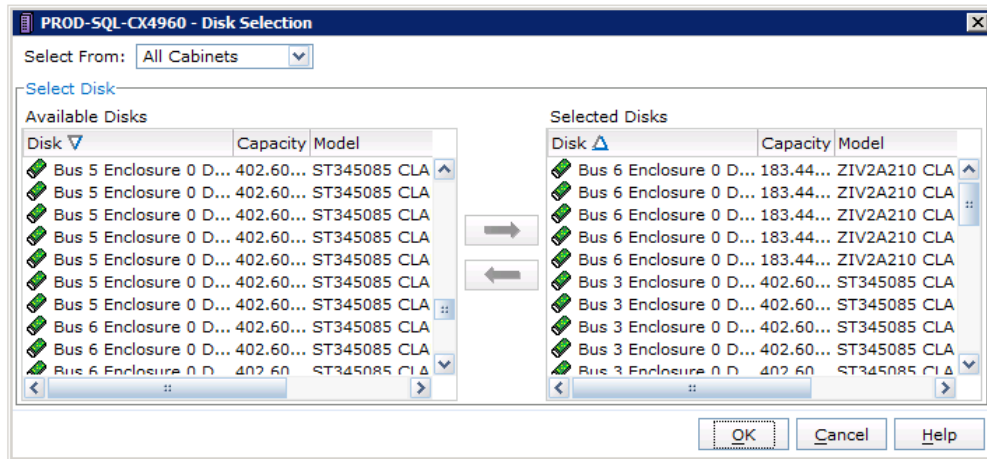


When creating a pool with a single type of disk, the auto-tiering options are visible, but auto-tiering cannot be enabled until additional disk types are allocated. For this 90 FC configuration, all the disks from Bus 1, 2, 3, 4, 5 and 6 were assigned to the storage group and the selection applied, after which the “Initiate Create Pool” operation is started.

Expanding a storage pool

The procedure for expanding storage pool properties in Unisphere is as follows:

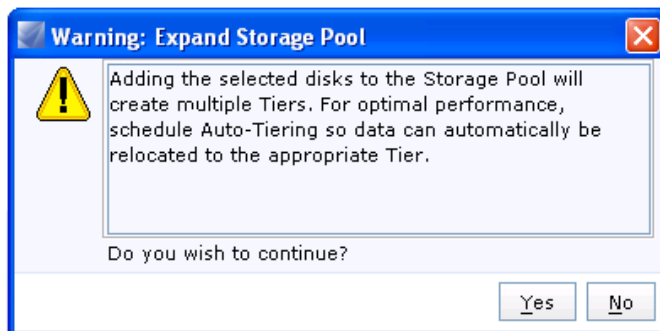
1. Choose **Disks** and click **Expand**.
2. Select the additional disks required, as shown in the following image.



Here additional disks are being selected.

As these changes are applied, a warning is displayed stating that adding the additional drive types will create multiple tiers, as shown in the following image. FAST automatically creates Tier 0 as the highest tier, with the Flash as the optimal performing drives. The already allocated FC drives become Tier 1 and the SATA drives form Tier 2, which is the lowest tier.

Once created, a storage pool can be easily expanded by right-clicking and selecting **Expand** or by clicking the **Expand** button when the pool is highlighted. Disks can then be selected as would be done when creating the pool, the only difference being that a different type of disk should be selected to those already in the pool. When you click **Apply**, a message is displayed warning that multiple tiers will be created, as shown in the following image.




Click **Yes** to continue. This has the effect of enabling auto-tiering as there are now multiple tiers to locate the data across.

Heterogeneous storage pools

With auto-tiering enabled on all LUNs, a heterogeneous storage pool with multiple disks types utilizing sub-LUN tiering is now available.

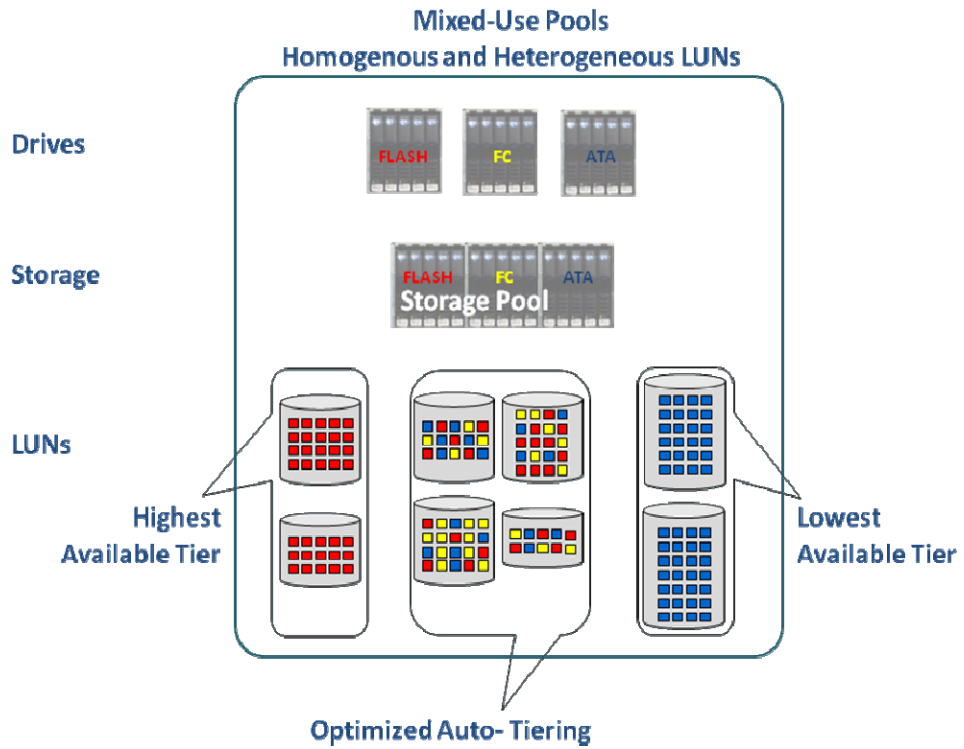
The following image illustrates the layout of heterogeneous storage pools.



Mixed-use storage pools

With the ability to assign LUNs to the highest and lowest tier, it is also possible to have mixed-use pools containing both homogeneous and heterogeneous LUNs. Administrators who have business-critical data they always want to pin to the highest tier can choose to place this data on the **Highest Available** tier and then, when the data is fully located on the tier, they can lock to the tier by choosing **No Data Movement**.

The following image illustrates the layout of mixed-use storage pools.



Testing

Introduction to testing This section provides the following details on the test run:

- Initial layout FC
- FC Baseline
- Layout with Flash and FC
- Flash and FC test
- Post relocation
- Test results
- Comparison of manual against FAST

Initial layout FC A disk storage pool consisting of 90 disks was configured prior to testing. The following table illustrates the layout of the production array as it was configured at this stage. This gave a Windows useable capacity of 28.98 TB for the pool.

BUS		CX4-960 PRODUCTION - DRIVE CONFIGURATION													
6.0	Pool 0	Pool 0 - FC													R5
	450GB 15k														
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
5.0	Pool 0	Pool 0 - FC													R5
	450GB 15k														
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
4.0	Pool 0	Pool 0 - FC													R5
	450GB 15k														
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
3.0	Pool 0	Pool 0 - FC													R5
	450GB 15k														
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
2.0	Pool 0	Pool 0 - FC													R5
	450GB 15k														
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1.0	Pool 0	Pool 0 - FC													R5
	450GB 15k														
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1.0	RG 300	RG 301	RG 302	spare	spare	spare	spare	spare	spare	spare	spare	spare	HS	HS	
	TD1/TD2/TL1	TD3/TD4	TPCE_Log												
	R10 2+2	R10 2+2	R10 2+2												
450GB 15k															
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	
0.0	RG 472	FLARE Drives			RG 300	RG 301	RG 302	spare	spare	spare	spare	spare	spare	HS	
	Master, Model, MSDB, Quorum				TD1/TD2/TL1	TD3/TD4	TPCE_Log								
	R5 4+1				R10 2+2	R10 2+2	R10 2+2								
146GB 15K				450 B 15k											
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	

Pool 0, a homogeneous storage pool, was configured using RAID 5 to hold six LUNs containing the test database, LUNs 100-106. LUN 107, which holds the test database log file, was configured on a traditional RAID 1/0 2+2 group.

Flash drive capacity is better targeted to house read-intensive I/O datasets of a random nature, with small I/O requests (up to 16 KB). Log files generally consist of sequential write operations and to this effect are better serviced by being housed on a dedicated spindle in either traditional RAID 1/0 groups or using dedicated RAID 1/0 pools.

The SQL Server tempdb database is generally read/write-intensive and is used for storing temporary tables, temporary stored procedures, and sub-queries, and for sorting aggregate operations. For the majority of workloads tempdb is also better suited to RAID 1/0.

The following table shows the configuration details.

LUN ID	Pool/RAID Group	Name	LUN size (GB)	Filename
30	300	LUN 30 TD1	35	tempdb.mdf
31	300	LUN 31 TD2	35	tempdb2.ndf
32	301	LUN 32 TD3	35	tempdb3.ndf
33	301	LUN 33 TD4	35	tempdb4.ndf
34	300	LUN 34 TL1	25	templog.ldf
100	Pool 0	LUN 100 TPCE	50	Market_1.ndf
				Misc_1.ndf
				MSSQL_tpce_root.mdf
101	Pool 0	LUN 101 TPCE_B1	180	Broker_1.ndf
102	Pool 0	LUN 102 TPCE_B2	180	Broker_2.ndf
103	Pool 0	LUN 103 TPCE_B3	180	Broker_3.ndf
104	Pool 0	LUN 104 TPCE_B4	180	Broker_4.ndf
105	Pool 0	LUN 105 TPCE_C1	100	Customer_1.ndf
				Customer_2.ndf
106	Pool 0	LUN 105 TPCE_C2	100	Customer_3.ndf
				Customer_4.ndf
107	302	LUN 107 TPCE_Log	100	MSSQL_tpce_log
TPCE DB Total			1070	

The table above shows that tempdb data files were allocated to the SQL instance. A best practice for SQL Server is to have an equal number of tempdb files per core, of equal size.

In SQL Servers which have many cores, a recent guideline established by Microsoft is to use a ratio of 4:1 CPU cores to tempdb data file as good starting point. This reduces the complexity of the solution because of having to manage so many tempdb data files, for example, 24 tempdb data files in this solution.

Only when it has been determined (due to blocking, GAM/SGAM locking, etc) that more tempdb data files are required in order to improve performance, should more tempdb data files be added in a controlled manner.

FC baseline

With the 90 FC pool in place the baseline test was run with the six data LUNs 100-106 spread across the pool and a workload was generated that drove disk utilization to 60 percent; this is a comfortable operating level for FC drives. It was also ensured that the physical disks: Avg. Disk sec/Read and Write SQL Performance counters did not exceed 20 milliseconds for data files and 10 milliseconds for log files, these are deemed to be safe operating levels for most OLTP environments.

Layout with Flash and FC

The pool was reconfigured to consist of four Flash and 30 FC drives (four Flash drives replacing 60 FC drives in the original configuration) and the dataset was restored to the FC tier.

CX4-960 PRODUCTION - DRIVE CONFIGURATION															
BUS															
6.0	RG 139 EFD-1 R5 3+1 200GB EFD				HS	empty	empty	empty	empty	empty	empty	empty	empty	empty	empty
	450GB 15k														
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
3.0	Pool 0 - FC R5 450GB 15k														
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
2.0	Pool 0 - FC R5 450GB 15k														
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1.0	RG 300 TD1/TD2/TL1 R10 2+2		RG 301 TD3/TD4 R10 2+2		RG 302 TPCE_Log R10 2+2		spare	spare	spare	spare	spare	spare	spare	spare	HS
	450GB 15k														
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
0.0	RG 472 FLARE Drives Master, Model, MSDB, Quorum R5 4+1 146GB 15K				RG 300 TD1/TD2/TL1 R10 2+2		RG 301 TD3/TD4 R10 2+2		RG 302 TPCE_Log R10 2+2		spare	spare	spare	spare	HS
	450 B 15k														
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14

This gave a formatted capacity of:

- Flash – 544 GB
- FC – 9.66 TB

FAST auto-tiering will always attempt to move the hottest data to the highest performing tier, based on the available capacity. The test database used in this configuration was 789 GB in size, compared to a top tier capacity of 544 GB. FAST will automatically reserve 10 percent of the tier to prevent over allocation.

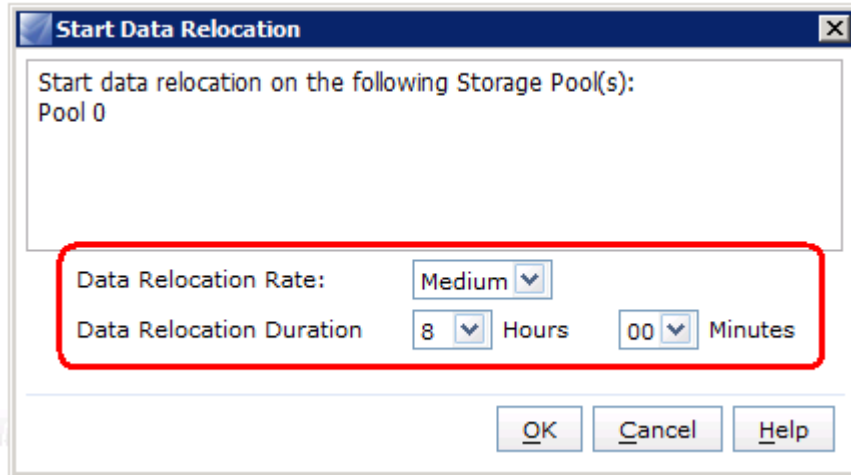
Flash and FC test

To allow FAST to make its relocation recommendations, a workload was applied to the environment using the exact same configuration setting as used for the FC baseline run. The load was allowed to run for four hours to allow the test tool to stabilize and relocation recommendation to be made. FAST recommend moving up 487 GB of data to the Flash tier.

Tier Name	Move Do...	Move Up (...)	User (GB)	Available
Flash	0.00	0.00	544.03	
FC	0.00	487.03	9660.59	42.03
				634.59

Data relocation was initiated by right-clicking the **Pool** and selecting **Auto-tiering > Start Data Relocation**, as shown in the following image.

During relocation, the default **Data Relocation Rate** (Medium) was used with the default **Data Relocation Duration** of an eight-hour window.



Relocation occurred over the eight-hour period with the workload continuing to run. Once relocation finished, the workload ran for a further two hours. This presents the data in its optimal configuration at the given time.

Post relocation

At the end of the relocation period, the LUN properties for each of the LUNs 100-106 were checked and it was seen that tiering occurred between the Flash and FC tiers in the percentages shown in the following table.

Pool 0	LUN 100	LUN 101	LUN 102	LUN 103	LUN 104	LUN 105	LUN 106
Flash	3%	79%	57%	77%	85%	12%	14%
FC	97%	21%	43%	23%	15%	88%	86%

The percentage in the Flash row indicates the percentage of the data residing in the LUN which FAST has determined to be hot data.

Once initial relocation occurred, FAST will continue to analyze the results of the hourly polling cycles and continue to make location recommendation. Moving data down to the FC tier and promoting data to the Flash tier as needed.

Test results

The testing of the FC baseline configuration and combined Flash and FC configurations was repeated several times with the same workload configurations for each test. To compare results the best run on the 90 FC pool was compared to the worst run on the four Flash and 30 FC pool.

The table below highlights the results identified between our configurations of 90 FC drives compared to four Flash and 30 FC drives.

	Before FAST	After FAST
Configuration	All Fibre Channel	Tiered Flash/FC
Disks	90 FC	4 Flash / 30 FC
Transactions /sec	2910.841	2983.441
Disk Transfers /sec	13785.326	14397.418
Avg Disk sec/Read & Write	all data files < 15 ms	
Avg Disk sec/Read & Write	log files < 5 ms	
CPU % Processor	41.92%	42.90%
SPA*/SPB Utilization	37.23% / 3.83%	38.45% / 3.82%
Disk Utilization	60%	91% / 62%

*SPA was the main allocation SP for the pool

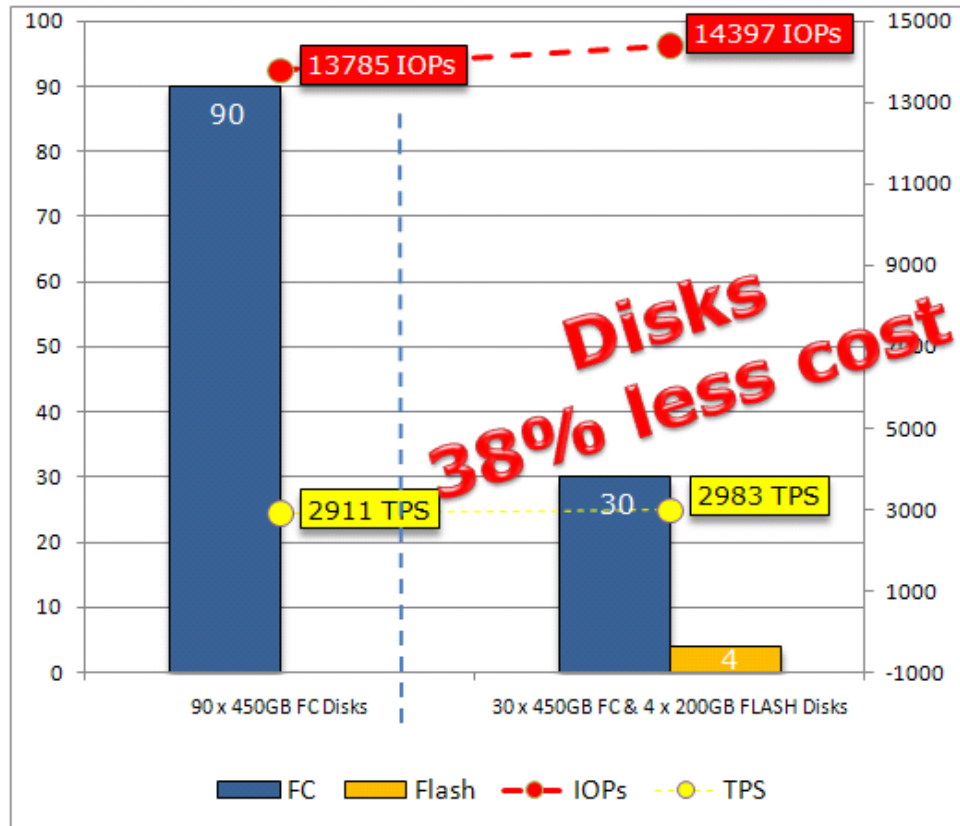
With four Flash drives now being used to replace 60 FC drives (**one Flash drive replacing 15 FC drives**) it was consistently seen that the new smaller footprint configuration provides improvement in performance for SQL Transaction/sec TPS and Disk Transfers/sec IOPs at a cheaper cost in drive acquisition, power, and cooling.

The table below highlights the power and cooling required on an annual basis for the two configurations. Implementing EMC FAST and Flash technology provides significant saving and a much greener footprint.

	Before FAST	After FAST	Change
Configuration	90 FC	4 Flash/30 FC	...
CX4-4PDAE	6	3	- 3 DAE
450 GB FC 15K	90	30	- 60 FC
200 GB Flash	0	4	+ 4 Flash
Power kVA	3.07	1.67	45.6% reduction
Heat Dissipation Btu/hr	10,000	5,500	45% reduction
Acquisition Cost	baseline	38% less	38% less

The diagrams below show the increase in both SQL Disks Transfers/sec (IOPS) and SQL Transaction/sec when employing EMC FAST and Flash technology together.

Disk IOPs and TPC before and after EMC FAST

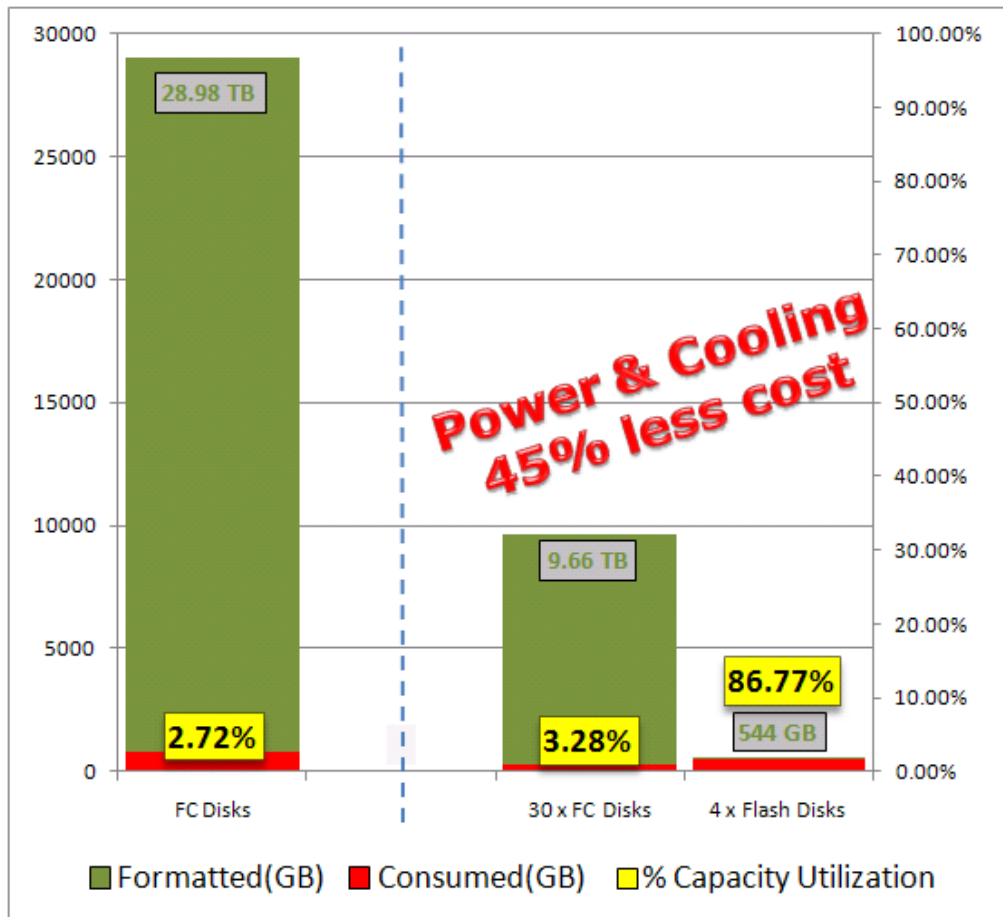


With such a saving it is easy to appreciate that spending more to increase the drive count on the smaller configuration would increase performance, while still saving money.

Disk Formatted Capacity Utilization before and after EMC FAST

The next graph shows the Disk Capacity Utilization between the FC and Flash disk configurations. In the 90 FC configurations you can see the extremely low utilization of the available 28.98 TB of formatted capacity, where 93.28 percent of the available capacity is unused. Even when moved, the new 30 FC and four Flash configuration shows only a marginal improvement in utilization of the FC disks at 3.28 percent, whereas the Flash drives are used to 86.77 percent. In fact, FAST puts a limit of 90 percent of the available space being used to keep some in reserve to prevent saturation of the disks.

This highlights the high capacity utilization of Flash drives, which when used in conjunction with EMC FAST waste little of this top tier of disks' capacity. At the same time, with FAST now relocating the hot data to the top tier we can now theoretically utilize the available capacity on FC to host data with low I/O requirements, such as archive data, files shares, and backups.

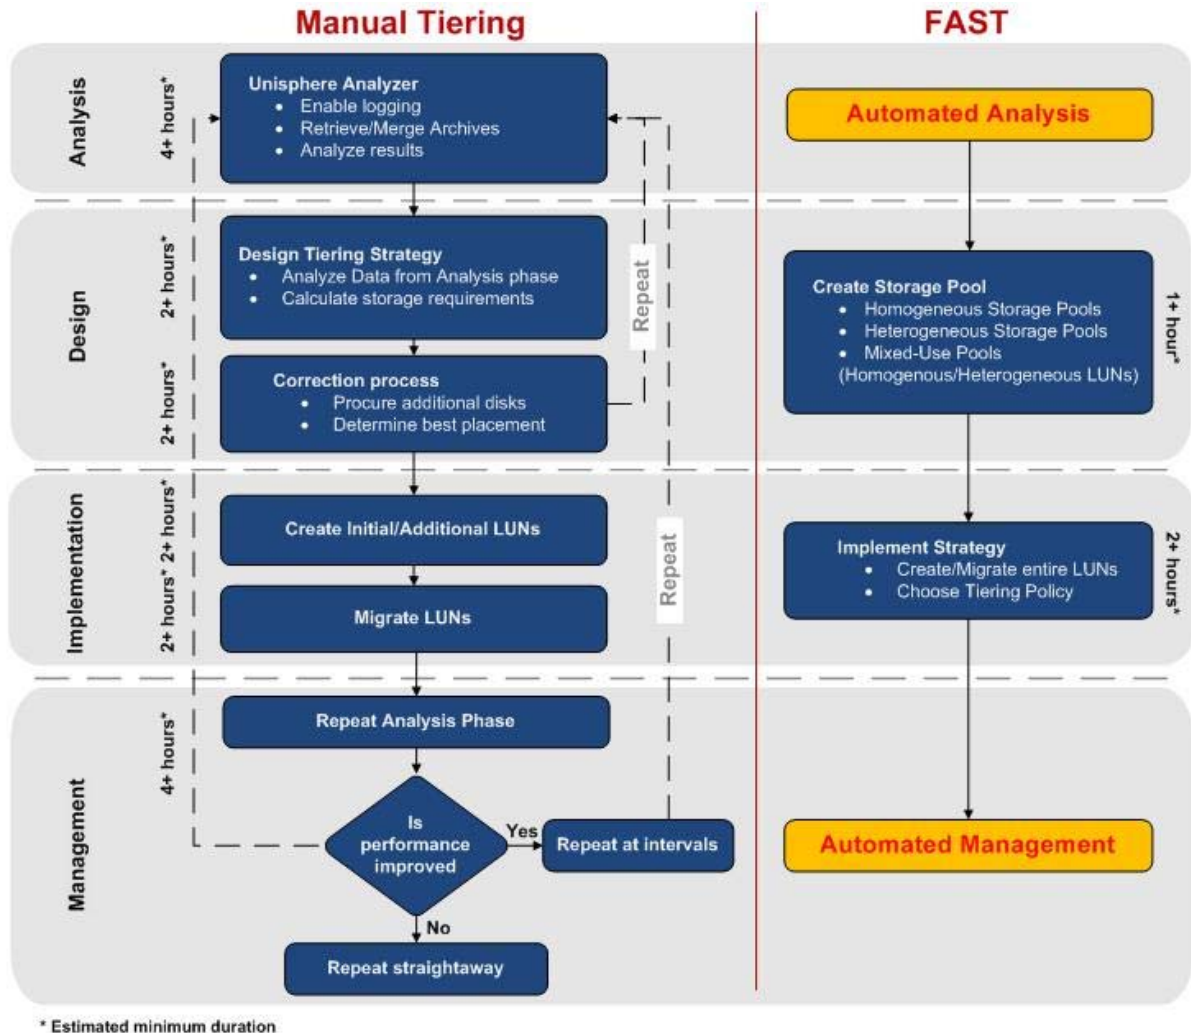


In a comparison between the configurations testing showed improvements in all the areas highlighted in the table below, most significantly being in the areas of reduced management, acquisition, and power/cooling costs.

<i>Configuration</i>	All Fibre Channel	Tiered Flash/FC
<i>Disks</i>	90 FC	30 FC / 4 Flash
<i>Tested TPS</i>	Baseline	2.4% Improvement
<i>Tested IOPs</i>	Baseline	4.2% Improvement
<i>Management</i>	Baseline	80% Less Time
<i>Acquisition Cost</i>	Baseline	38% Less Cost
<i>Power/Cooling</i>	Baseline	45% Less Cost

Comparison of manual against FAST

The following image outlines some of the activities that are required to implement a manual storage tiering strategy. It must also be remembered that administrators would need to complete this process for each application load being applied on the array.



In comparison to the complex manual process flow, it can be seen that EMC FAST technology provides a simple and speedy approach to implementing storage tiering. Administrator can now create a storage pool, create or migrate a LUN to the pool, then simply choose a storage tiering policy. This provides administrators with the ability to use a simple **three-step "set it and forget it"** approach to tuning application workloads such as SQL OLTP in **80 percent less time**.

Conclusion

Summary

Testing showed a typical customer environment with an existing storage configuration of 90 FC drives servicing an OLTP type workload on their CLARiiON CX4-960. This was then compared with a pool consisting of 30 FC drives and four Flash drives running exactly the same workload.

The functionality, testing, and observations documented in this white paper demonstrate and prove how:

- Employing a small number of Flash drives on the array can provide a better performing and cost-effective alternate to employing traditional practices of high spindle counts and short-stroking disks to service IO requirements.
- EMC FAST technology means investment in EMC Flash technology and provides immediate benefits by reducing hardware provisioning costs, decreasing the configuration footprints on the array, and reducing running costs by lowering power consumption. This ability to provide **leaner and greener disk configurations** is a major step forward in storage array technology.
- EMC FAST technology also simplifies storage management. FAST is quick and easy to implement, and provides a simple, automated process to tier storage pools to match I/O requirements, with the option to eliminate manual intervention completely.
- With Unisphere's management capabilities expanding pools to cover growing capacity and I/O requirements is now easy. For example, just one Flash drive can be quickly added to the current 30 FC and four Flash drive pool to increase the capacity of the top performance tier by 180 GB and support servicing another 2,500 disk IOPS.
- With sub-LUN tiering in FLARE reducing complexity on the array, LUNs can now be spread across the tiers at the block level. Through continuous polling and analysis, FAST continues to optimize the tiering of the environment, moving data either at the implementation of the administrator or through the scheduled windows. This is especially useful when workloads are unpredictable and variable from week to week.
- With the current four Flash drives in the configuration, an administrator can easily expand this tier by adding one or more Flash drives, increasing the capacity of the Flash tier to host more data and also service more I/O requests. Similarly, if an administrator decides to archive data SATA drives could be added to the pool.

Findings

The table below summarizes the key points that this solution addresses.

Key point	Solution objective
Automated tiering of storage pools <ul style="list-style-type: none">• Reduced administration• Quicker reaction to changes	Testing proved FAST's ability to tier storage pools at a sub-LUN level, based on I/O throughput, and demonstrated FAST's ability to analyze, recommend, and automatically implement data relocation plans through Unisphere's continuous analysis of the storage array. EMC FAST technology allows administrators to perform " set it and forget it " tuning for application workloads like SQL OLTP by more than 80 percent less time .
Increased ROI through <ul style="list-style-type: none">• reduced hardware costs• reduced running costs• reduced footprint	With FAST automatically moving only the high I/O, hot data to the Flash tier, fewer drives are now required to service the workload. This reduces TCO and provides quicker ROI. Testing proved how EMC FAST technology allows administrators to make optimal use of the investment in Flash technology. Testing showed more than a 38 percent saving between the two configurations, showing proving FAST and Flash technology can maintain performance while reducing costs .
Reduced complexity	Sub-LUN tiering reduced granularity by spreading single LUNs across tiers, negating the need to move whole LUNs. This optimizes the use of available capacity and greatly simplifies storage design layout.

Next steps

To learn more about this and other solutions contact an EMC representative or visit: www.emc.com.

References

White papers

For additional information, see the white paper listed below.

- *EMC Tiered Storage for Microsoft SQL Server 2008 Enabled by EMC CLARiiON CX4 and Enterprise Flash Drives—A Detailed Review – White Paper*