

# **EMC Tiered Storage for Microsoft SQL Server 2008 Enabled by EMC Symmetrix VMAX with FAST**

*A Detailed Review*

## **EMC Global Solutions**

---

### ***Abstract***

This white paper examines the configuration details, efficiency, and increased performance levels achieved through the integration of the EMC<sup>®</sup> Symmetrix<sup>®</sup> VMAX<sup>™</sup> Fully Automated Storage Tiering (FAST) capability into an enterprise online transaction processing (OLTP) environment running on Microsoft SQL Server 2008.

May, 2010

---

---

Copyright © 2009, 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 H6708.1

---

## Table of Contents

Executive summary .....	4
Introduction.....	5
Evolution of Storage Tiering .....	7
Key environment components.....	8
Configuration .....	10
Microsoft SQL Server 2008 test configuration .....	11
Storage design layout .....	15
FAST functionality .....	16
FAST design considerations .....	18
Configuring FAST .....	19
Validation of FAST technology with SQL Server OLTP workloads.....	25
Conclusion.....	29
References .....	30

---

## Executive summary

---

Businesses rely heavily on Microsoft SQL Server as an enterprise-class database solution to run demanding applications such as online transaction processing (OLTP). A key challenge for IT administrators is optimization and operational efficiency of available resources.

Storage tiering is one technique an IT administrator can use to optimize available resources. This method allows the movement of data to suitable Storage Types based on business requirements. For example, critical, frequently accessed data is automatically moved to a higher-performing Storage Type, while less frequently accessed data is automatically moved to a lower-cost, higher-capacity Storage Type. This is achieved through the use of key EMC® Symmetrix® VMAX™ storage array features including Enterprise Flash Drives (EFDs), SATA drives, and EMC Fully Automated Storage Tiering (FAST).

This white paper illustrates the benefits of Symmetrix VMAX FAST functionality in an enterprise-class SQL Server environment. FAST enables administrators to automatically optimize storage resources using three levels of storage media (or “Storage Types”) available on a Symmetrix VMAX:

- EFDs—Represent the tier used for critical, most heavily accessed table partitions.
- Fibre Channel (FC) drives—Represent the tier used for less critical, infrequently accessed table partitions.
- Serial advanced technology attachment (SATA) drives—Represent the tier used for inactive or static table partitions.

Testing will demonstrate how FAST technology automates the optimization of SQL Server database storage resources by providing:

- A mechanism whereby data is automatically moved to appropriate Storage Types, based on configured FAST Policies. For example, frequently accessed Storage Types (EFDs) versus inactive Storage Types (SATA drives)
- Consistently high storage performance levels for SQL Server OLTP workloads

---

### Business case

OLTP systems hosted on a SQL Server platform represent one of the most common data processing systems in today’s business environments. As the storage and business requirements of data in OLTP systems evolve over time, they will require constant manual intervention to maintain peak operational efficiency leading to high maintenance costs. Businesses require automated mechanisms that keep pace with these fast-changing environments.

The FAST functionality available in the Symmetrix VMAX storage system meets these requirements by implementing mechanisms that:

- Facilitate the nondisruptive movement of data to the most effective Storage Type.
  - Reduce administrative costs by automatically creating configuration change plans that provide built-in data analysis.
  - Provide automated tiering mechanisms that drive optimum resource utilization, increased efficiency, and optimal performance.
-

---

## Technical solution

FAST technology enables organizations to automate the movement of critical, frequently accessed data to maximum performance storage, and less frequently accessed data to lower-cost, higher-capacity storage.

This automated tiered storage design is comprised of the following elements:

- Tiering within the array by using different drive technologies:
  - High-performing EFDs
  - Traditional FC drives that deliver reasonable performance at lower cost
  - High-capacity, low-cost SATA drives
- Symmetrix FAST controller:
  - Provides an automated data analysis function based on preset FAST algorithms
  - Makes recommendations to migrate the data nondisruptively to the appropriate Storage Type
  - Moves workloads automatically, or with user approval
- Symmetrix Management Console (SMC) is a web-based interface that provides a simple way to configure FAST settings and policies.
- Symmetrix Performance Analyzer is a web-based interface that provides details of performance for each application, by monitoring performance before and after the FAST move.

---

## Introduction

This white paper illustrates the benefits of automated Storage Tiering in an active, enterprise-class SQL Server OLTP environment. To help demonstrate the impact of Storage Tiering, a performance baseline of the OLTP environment is determined prior to implementing FAST technology.

Next, SQL Server database volumes containing partitioned table data files are identified within the storage layer as candidates for relocation. FAST can automatically relocate these volumes to a more appropriate Storage Type.

Lastly, the SQL Server application performance is once again measured to compare against the baseline performance results.

This white paper details the following:

- Test environment configuration and test profile
- FAST functionality
- FAST design considerations
- Configuring FAST
- FAST validated test results
- FAST findings and conclusions

---

**Purpose**

The purpose of this solution is to:

- Demonstrate the efficiency, and storage optimization benefits that can be achieved by integrating FAST technology into an enterprise SQL Server OLTP environment.
  - Highlight the time-saving benefits by comparing FAST automation processes with manual Storage Tiering practices.
  - Demonstrate the benefits of Storage Tiering on performance for SQL Server OLTP workloads.
- 

**Scope**

Consider the following before proceeding:

- Testing is based on a high volume, enterprise-class SQL Server OLTP environment.
  - Processes outlined during each stage are high-level in nature and should be read in conjunction with the documentation referenced in this white paper.
  - The information in this document is not intended to replace existing, detailed product implementation guides or best practices.
- 

**Audience**

This white paper is intended for:

- Field personnel who are tasked with deploying the EMC Symmetrix VMAX as the storage platform
- Customers, including IT planners, storage architects, and database administrators
- Engineering and product development groups
- EMC staff and partners, for guidance and development of proposals

It is assumed that the reader is familiar with:

- Microsoft SQL Server, and SQL table partitioning
  - EMC Symmetrix storage
  - EMC Solutions Enabler
  - EMC Symmetrix Management Console
- 

**Terminology**

Review the following for a better understanding of the terms relative to FAST.

**Configuration Change Plan:** A list of data migration recommendations provided by FAST to achieve improved performance.

**Disk Group:** Physical disk drives within the Symmetrix VMAX that are grouped by technology (EFD, FC, SATA), rotational speed, and size. Drive groups form the base unit for Storage Types. A disk group may be part of multiple Storage Types.

**Dynamic Reallocation Volumes (DRV):** A non-user-addressable logical volume used by the Symmetrix VMAX to temporarily hold data while transferring between Storage Types.

**Dynamic Storage Tier:** Physical disk drives that are grouped by technology and protection type (for example, RAID 1, RAID 5, RAID 6). A Storage Tier can contain one or more disk groups of the same technology. A dynamic group will automatically include all drives of a particular technology.

---

**Static Storage Tier:** Physical disk drives that are grouped by technology (EFD, FC, SATA) and protection type (RAID 1, RAID 5, RAID 6). A static Storage Tier can contain one or more disk groups of the same technology.

---

## Evolution of Storage Tiering

---

From a business perspective, “Storage Tiering” generally means that policies coupled with storage resources having distinct performance, availability and other characteristics are used to meet the service level objective (SLO) for a given application. (By SLO we mean a targeted I/O service goal, that is, performance for an application.) This remains the case with FAST.

For administrators, the definition of “Storage Tiering” is evolving. Initially, different storage platforms met different SLOs. For example:

- Gold Tier – Symmetrix; Silver Tier – CLARiiON®; Bronze Tier – Tape

More recently, “Storage Tiering” meant that multiple SLOs are achievable in the same array:

- Gold Tier – 15k, FC RAID 1; Silver Tier – 10k, FC RAID 5; Bronze Tier – 7.2k, SATA, RAID 6

For example, users might establish a Gold Storage Class as follows:

Service level objective	Storage Class	FAST Policy	Storage Type
Read/Write response time objective	Gold	10%	15k rpm, RAID 1
		40%	10k rpm, RAID 5
		50%	7.2 k SATA RAID 6

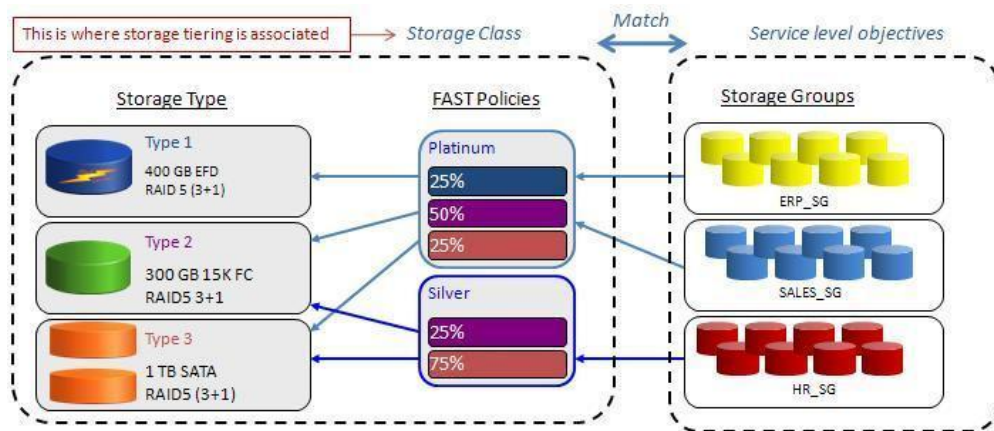
---

**Relationship between FAST Policies, Storage Class, Storage Group, and Storage Type**

Note, because multiple Storage Types can support the same application, "tier" is not used to describe a category of storage in the context of FAST. Rather, EMC is using new concepts defined next.

- **FAST Policies:** Policies that manage data placement and movement across Storage Types to achieve service levels for one or more Storage Groups.
- **Storage Class:** A combination of Storage Types and FAST Policies to meet service level objectives (SLO) for Storage Groups.
- **Storage Group:** A logical grouping of volumes (often by application) for common management.
- **Storage Type:** A shared storage resource with common technologies, namely drive type and RAID scheme.

The following figure illustrates how these mechanisms work together seamlessly to deliver FAST.



**Key environment components**

The following sections examine how key components work to integrate the FAST capability into the SQL Server application providing high-performing database management.

**EMC Symmetrix VMAX storage array**

The EMC Symmetrix VMAX with Enginuity™ version 5874 Q4 service release is the newest addition to the Symmetrix family. The tiered storage configuration used in the test environment is based on the following Symmetrix VMAX features:

- FAST functionality
- Three levels of storage media (EFDs, FC drives, SATA drives) that comprise the tiered storage

---

**Fully Automated Storage Tiering (FAST)**

EMC Symmetrix VMAX FAST technology is designed to automate the relocation of application data across different performance tiers based on the changes in application requirements.

The FAST controller performs four basic, automated functions:

- Collects performance data
  - Analyzes the data in conjunction with capacity data
  - Generates a configuration change plan based on the analysis
  - Relocates the data to the appropriate tier
- 

**FAST managed objects**

FAST technology is driven by the following managed objects:

- Storage Types
  - Storage Groups
  - FAST Policies
- 

**EFDs**

EFDs can dramatically increase performance for demanding Microsoft SQL OLTP database applications because they can deliver single-millisecond application response times, and significantly higher IOPS as compared to traditional FC disk drives. Energy consumption can be significantly reduced using EFDs.

The high-performance characteristics of EFDs eliminate the need for organizations to purchase large numbers of traditional hard disk drives, while only utilizing a small portion of their capacity to satisfy the IOPS requirements of a SQL Server database environment.

---

**EMC Symmetrix Management Console (SMC)**

The Symmetrix VMAX storage system provides a web browser interface, Symmetrix Management Console (SMC). SMC provides centralized management to the entire Symmetrix VMAX storage infrastructure.

In the context of FAST, SMC integrates easy-to-use wizards for:

- Creating Storage Tiers
- Associate FAST Policies to Storage Groups

**Note:** The Solutions Enabler Command Line Interface (SYMCLI) may also be used to manage FAST objects.

---

---

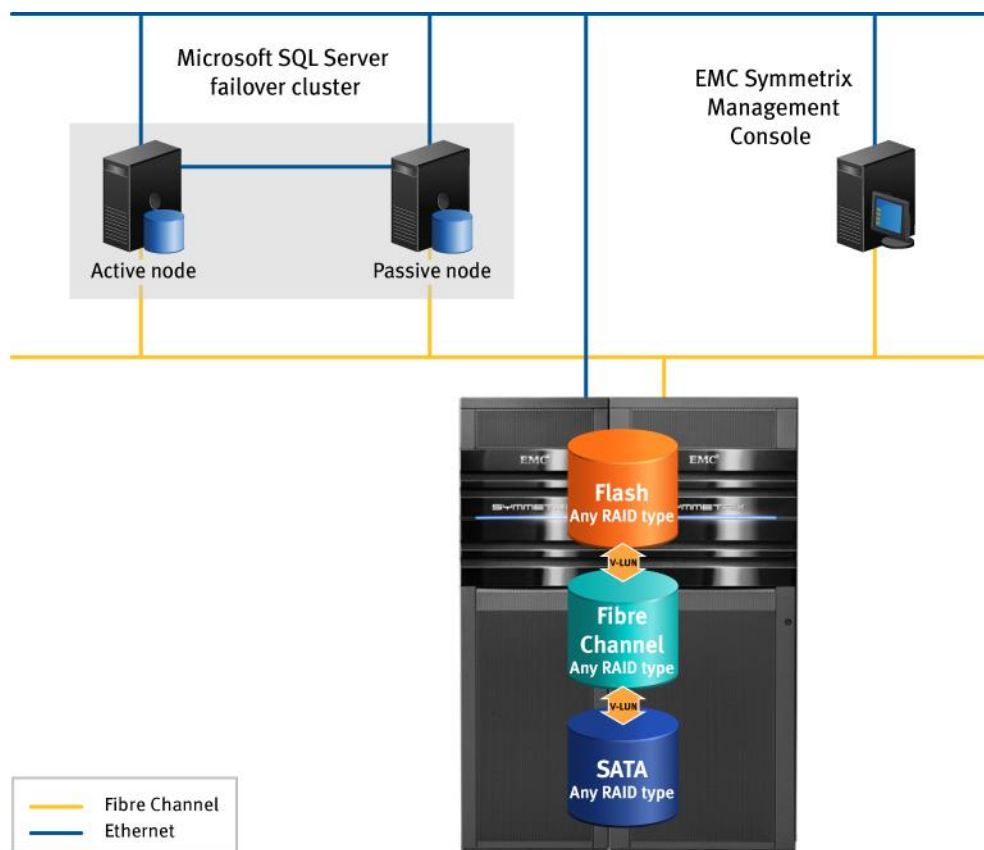
## Configuration

The FAST design implements the following physical connections:

- A two-node Microsoft failover cluster:
    - One active node with four quad-core CPUs and 64 GB of RAM
    - One passive node with four quad-core CPUs and 64 GB of RAM
  - A utility server hosting the SMC server components
  - FC connectivity provided by a 4 Gb/s director level enterprise-class switch
  - The Ethernet backbone provided by a Gigabit Ethernet network switch
  - The Symmetrix VMAX is attached to the servers through four of its front-end ports
  - The Symmetrix VMAX provides EFD, FC, and SATA drive technologies
- 

## Physical architecture

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



SYM-002231

---

---

## Microsoft SQL Server 2008 test configuration

---

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

- Number of SQL users supported: 75,000
  - Simulated user workload with 1% concurrency rate and zero think time consistent with Microsoft testing methodologies
  - User data: 1 TB
- 

### SQL Server test application

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

---

### OLTP workloads

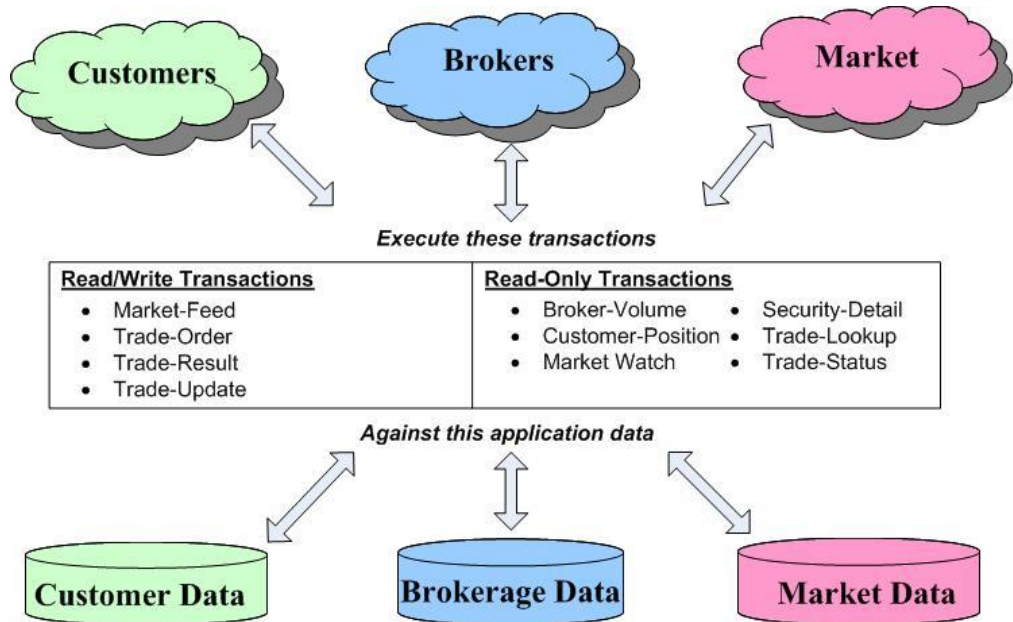
The OLTP application used to generate user load in this test environment is based on the TPC Benchmark-E (TPC-E) standard. TPC-E testing is composed of a set of transactions that represent the processing activities. The database schema, data population, transactions, and implementation rules have been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

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

---

**Key components of SQL Server testing**

This benchmark is composed of a set of transactions that are executed against three sets of database tables that represent market data, customer data, and broker data. A fourth set of tables contains generic dimension data such as zip codes. The following diagram illustrates the key components of the test environment.



---

**Partitioning the SQL database**

SQL table partitioning is used to segment data into smaller, more manageable sections. Table partitioning can lead to better performance through parallel operations. The performance of large-scale operations across extremely large data sets (for instance many millions of rows) can benefit by performing multiple operations against individual subsets in parallel.

The individual subsets can be moved to several disk drives to effectively reduce I/O contention.

The number of table partitions to allocate depends on:

- Table size
- LUN utilization

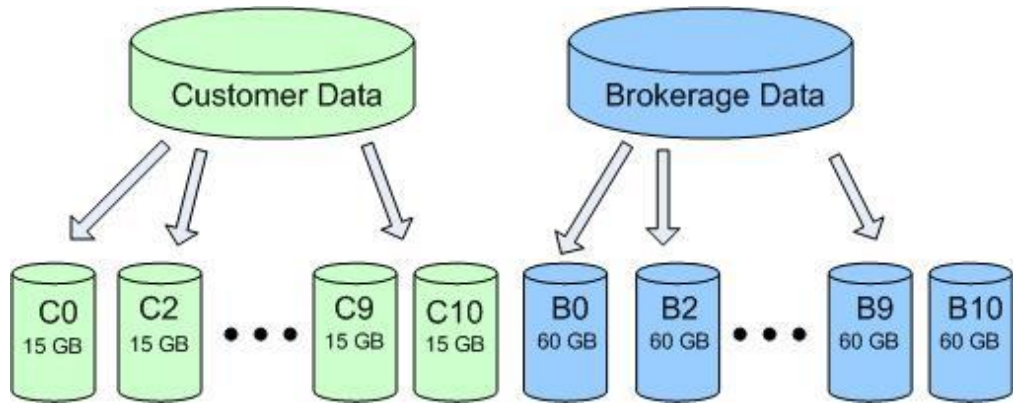
The broker and customer file groups for this application are the largest and best candidates for partitioning.

---

---

**Broker and customer file groups partitioning**

The broker and customer file groups are each divided into 11 partitions. Each partition is stored on a separate LUN. The first 10 partitions hold data generated during the initial data population. The 11th partition holds new data generated during simulated user activity; see the following figure.



**Broker and customer file groups**

The following table details the file groups used in testing.

**Note:** The OLTP application's storage is configured across the FC drives using RAID 1 protection. This serves as a base configuration utilizing FAST to relocate the storage.

File group name	Table name	Drive (directory with mount point)
broker_fg1-10	CASH_TRANSACTION	M:\Broker\B1~B10
	SETTLEMENT	
	TRADE	
	TRADE_HISTORY	
customer_fg1-10	HOLDING	M:\Customer\C1~C10
	HOLDING_HISTORY	
broker_fg	CHARGE	M:\Broker\B0
	COMMISSION_RATE	
	TRADE_TYPE	
	TRADE_REQUEST	
	BROKER	
customer_fg	ACCOUNT_PERMISSION	M:\Customer\C0
	CUSTOMER	
	CUSTOMER_ACCOUNT	
	CUSTOMER_TAXRATE	
	HOLDING_SUMMARY	
market_fg	EXCHANGE	M:\MKT DB
	INDUSTRY	
	SECTOR	
	STATUS_TYPE	
	COMPANY	
	COMPANY_COMPETITOR	
	DAILY_MARKET	
	FINANCIAL	
	LAST_TRADE	
	NEWS_ITEM	
	NEWS_XREF	
	SECURITY	
	WATCH_ITEM	
	WATCH_LIST	
misc_fg	TAXRATE	M:\MKT DB
	ZIP_CODE	
	ADDRESS	
Tempdb		Y:\TEMPDB1~4
Transaction Log		N:\

---

## Storage design layout

---

This test environment uses the following drive configurations and storage design guidelines.

---

### Physical drive configuration

The following table details the physical drives used in the test environment.

Drive type	Number of drives	Specifications	RAID type
EFDs	8	400 GB	RAID 5
FC drives	96	450 GB, 15k rpm	RAID 1
SATA drives	32	1,000 GB, 7.2k rpm	RAID 6

---

### Logical drive configuration

The following table details the logical drives used in the test environment.

Broker file group	Customer file group	TempDB files	Market and misc. file group
11 LUNs, 90 GB	11 LUNs, 25 GB	4 LUNs, 25 GB	1 LUN, 45 GB

---

### Storage design considerations

Consider the following when configuring storage:

- Multiple drive technologies are not a requirement for Storage Tiering.
  - Allow for unconfigured storage in each tier for use by data movement.
  - The Virtual LUN migration feature can be used to manually relocate LUNs as required.
  - Plan for an adequate number of DRV devices in cases where device swaps take place. The number of DRV devices should be equal to the Maximum Simultaneous Volumes Moved parameter.
  - DRV devices should not reside on SATA drives. Placing the DRV devices on SATA drives could lead to poor application performance during a swap operation.
-

---

## FAST functionality

---

### Manual tiering operation versus FAST operation

In our simulated, typical customer OLTP environment testing shows that the number of support hours required to maintain manual tiering far exceed the amount of time required when FAST is implemented, as detailed in the following table.

Storage administrator tasks	Support level required for manual tiering	Support level required for FAST
Carefully monitor storage array performance	High level of support	None required
Manually collect performance data	High level of support	None required
Spend hours analyzing key metrics	High level of support	None required
Reconfigure storage based on a snapshot of the storage array's performance	High level of support	None required
Device reallocation	High level of support	None required

---

### How FAST identifies the appropriate Storage Tier

FAST uses algorithms to analyze data and identify the most appropriate tier, as described next:

- During the user specified working window for performance, the FAST controller samples performance statistics for Storage Groups associated with FAST Policies at 10-minute intervals.
  - When enough performance samples (as defined by the "Initial Period" parameter in the FAST settings) have been collected, the FAST controller will begin to analyze the data to determine the appropriate tier for a device. Three algorithms are considered when making this decision. The algorithms are listed in order of descending probability, as follows:
    - **EFD promotion/demotion.** This algorithm is designed to maximize EFD utilization in the Symmetrix VMAX. The algorithm works by modeling EFD performance for each device being monitored by FAST. Each device is given an "EFD performance score" based on average reads per second and read/write ratios. The devices with the highest scores are considered good candidates for promotion to EFDs.
    - **Capacity-based.** This algorithm enforces the percentages configured in the FAST Policy. If a Storage Group has a higher percentage of devices on a tier than the policy allows, a recommendation is made to move devices to a different tier in compliance with the FAST Policy
    - **FC/SATA Cross Tier.** Used to balance utilization across FC and SATA drives. Only devices on these two technologies are considered by this algorithm. The devices are ranked based on utilization, the most utilized devices recommended to move to the least utilized drives.
-

---

### Configuration change plan summary

The end result of FAST's analysis is the creation of a configuration change plan that lists recommendations for device movement to a more appropriate Storage Type. SMC triggers an alert when a configuration change plan is created. If the data movement mode is set to **Automatic** the device migrations will occur during the next available data move window. If the data movement mode is set to **User Approved**, data migrations will occur during the next data movement window after the user has approved the configuration change plan.

---

### FAST device migration

One of the most critical functions of FAST is device migration, described next.

- The FAST controller uses one of two methods (a move or a swap) to relocate devices:
    - A move occurs when there is enough unallocated space available in the target Storage Tier to accommodate the move.
      - A move is considered as the preferred method of relocating a device. This process uses the virtual LUN migration feature to relocate the device to a different tier. Virtual LUN migration creates a new device from unconfigured space in the target tier.
      - The new device is set as a secondary mirror and synchronized to the first device, or the primary mirror.
      - When synchronization is complete, the mirrors swap priority with the secondary mirror becoming the primary.
      - The original device, now the secondary mirror, is detached and deleted becoming free space.
    - A swap takes place when there is no unconfigured space in the target tier and a corresponding device of similar size is available.
      - A swap involves three data copies and uses a DRV device to preserve the data on the devices involved in the swap.
      - A DRV device is selected and is associated with the first device in the swap. The DRV is completely synchronized with the device. Once synchronization is complete, the original device is associated with the second device.
      - The original device is then synchronized from the second device. When synchronization is complete, the second device is associated with the DRV device and data is synchronized from the DRV device.
      - Once synchronization is complete, the DRV device returns to an available state.
-

---

## FAST design considerations

---

Consider the following prior to configuring FAST:

- In FAST, **Automatic mode** executes the data movement plans during the window provided. **User approved mode** requires the additional step of approving the configuration change plan.

**Note:** EMC recommends selecting **User approved mode** option at least until the administrator understands the FAST controller's recommendations and feels comfortable with them.

- The maximum moves per day and the maximum simultaneous move parameters count a device move or a device swap as a single occurrence even though a device swap involves multiple copy tasks.
- Once a device move or swap is initiated, it cannot be cancelled or stopped.
- When setting the device move window, take into account the time window allowed. Device relocation will not start within the last half hour.
- Any moves in progress will run to completion even though the window time limit is exceeded.
- Quality of service (QoS) can be used to control the pace of the synchronization phase of a move/swap; see the *EMC Solutions Enabler Symmetrix Array Controls Command Line Interface Guide*.
- Use FAST in conjunction with the optional EMC Symmetrix Optimizer data management software solution to further enhance the load balancing capability across tiers. For example, While FAST algorithms balance workloads across Storage Types, Symmetrix Optimizer algorithms achieve a balanced workload within a Symmetrix disk group. See the *EMC Symmetrix Optimizer—A Detailed Review* white paper for additional information.

**Note:** When using FAST in conjunction with the Symmetrix Optimizer the parameters listed in [Set FAST settings](#) are shared between both of these products. Administrators cannot set these parameters separately.

---

---

## Configuring FAST

---

This section describes the configuration tasks for setting up FAST in the test environment using the following sequence:

- Set FAST settings
- Set FAST working windows
- Create Symmetrix tiers
- Create a FAST Policy
- Associate the Storage Group

**Note:** The steps outlined during each stage are high-level in nature. For detailed procedures, please see the *Implementing Fully Automated Storage Tiering (FAST) for EMC Symmetrix VMAX Series Arrays Technical Note*.

---

### Management tools

The following management tools are available to configure FAST:

- Solutions Enabler Command Line Interface (SYMCLI)
  - Symmetrix Management Console (SMC)
-

---

## Set FAST settings

**Data Movement Mode** is an important setting for FAST. Consider the following before setting the mode to **Automatic** or **User Approved**:

- The **Automatic** setting will execute data movement plans during the window (timeframe) selected. See [Set FAST working windows](#).
- The **User Approved** setting requires the additional step of approving the move plan. Additional **Data Movement Mode** parameters include:
  - **Maximum Number of Volume(s) Moved per day**: The maximum number of moves or swaps FAST may perform in a 24-hour period. The maximum is 200.
  - **Maximum Simultaneous Volume(s) Moved**: The maximum number of concurrent moves or swaps. The maximum allowed is 32.
  - **Workload Analysis Period**: Defines the amount of historical statistical information used for analysis. Ideally, this parameter should be set for a normal business day. This parameter can be set for hours, days, and weeks. The maximum is four weeks.
  - **Initial Period**: Defines the minimum amount of samples to be collected before analyzing data. Set based on the application's normal workload.

**FAST - Settings ( 000192601343 )**

General | Time Window

**Shared Parameters**

Data Movement Mode:  Automatic  User Approved

Maximum Number of Volume(s) Moved per Day:

Maximum Simultaneous Volume(s) Moved:

Workload Analysis Period:  hour(s) ▼

Initial Period:  hour(s) ▼

---

## Set FAST working windows

FAST uses time frames (referred to as working windows) to gather performance data and to schedule data movement.

The first type of window is the performance window. During this time period the FAST controller will sample performance data to be included in the analysis of what data to move.

The second type of window is the data movement window. This window specifies the time when data movement is allowed, or not allowed. FAST data movement runs at a low priority, but still produces overhead on the EMC Symmetrix.

### Recommendations

- Schedule the performance window during regular business hours.
- Schedule the data movement window during off peak hours.

**Note:** An important setting for the FAST controller is the **Data Movement Mode**, which may be configured as **Automatic** or **User approved**. The Automatic setting will execute data movement plans during the window provided and **User Approved** setting will require the additional step of approving the move plan.

The screenshot shows a dialog box titled "FAST/Optimizer - Select a time window". It contains the following fields and options:

- Time Window Name:** A text box containing "Test Performance Window".
- Time Window Type:** Radio buttons for "Inclusive" (selected) and "Exclusive".
- Recurrence:** Radio buttons for "One Occurrence" (selected) and "Weekly Recurrence".
- Recurrence Range:** Two date pickers. The "Start Date" is "Nov-02-2009 11:34" and the "End Date" is "Dec-31-2020 00:00".
- Buttons:** "OK", "Cancel", and "Help" buttons at the bottom right.

## Create a Symmetrix tier

Creating a Symmetrix tier requires setting the:

- Tier name
- Disk technology
- RAID protection type
- The default action is to create a dynamic tier that includes all of the disks of the selected drive technology. Selecting the **Manually select disk groups** checkbox will create a static tier.
- Disk groups can be included in the static tier by choosing the **Selection** checkbox.

**Recommendation:** Use static tiers to control the RAID protection of multiple disk groups of the same technology.

**Tier - Create Dialog (000192601343)**

Tier Name:

Technology:

Protection:

Manually select disk groups.

**Capacity Numbers for new Tier:**

Symmetrix Capacity:	443146 GB
Total Tier Capacity:	2980 GB
Total Used:	0 GB
Total Free:	2980 GB

Note: As this tier is static, only disk groups which have been manually added will be included.

Selection	# Disk Group	Disk Group Name	Disk Speed	Disk Count	Capacity (GB)	Us
<input checked="" type="checkbox"/>	2		0	8	2980	0

Select All Deselect All

OK Cancel Help

---

## Create a FAST Policy

A FAST Policy is a set of rules that govern how automated Storage Tiering is applied to a particular Storage Group using FAST technology. Each policy can include up to three tiers. The administrator determines what percentage of the Storage Group's data resides on any particular tier.

Consider the following before creating a FAST Policy:

- Three tiers are specified in this example. The percentage of the data that can be on the tier can be anywhere from 0 percent to 100 percent. The total of all three tiers needs to total at least 100 percent.
- Note that it is possible to set every tier at 100 percent. In this case, FAST is free to place data on any tier.
- When determining the percentages, it is important to take into account:
  - The possible impact of moving data to a lower-performing tier.
  - Other applications may be sharing resources on a particular tier.
  - Multiple FAST Policies may reference the same tier.

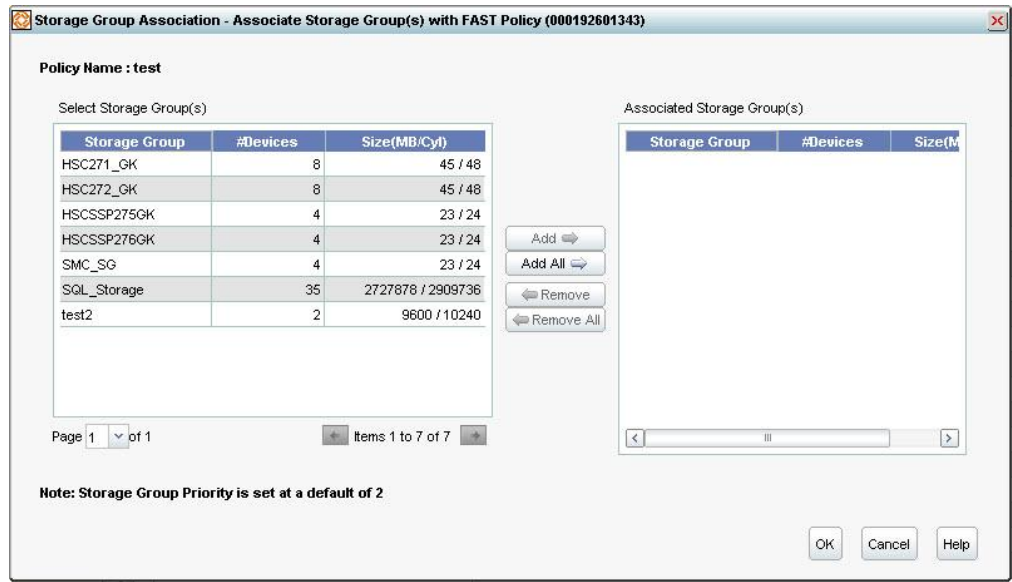
The screenshot shows a dialog box titled "FAST Policy Management - Create FAST Policy (000192601343)". It contains the following fields and controls:

- Policy Name :** A text input field containing "Test Policy".
- Tier :** Three rows, each with a dropdown menu and a percentage input field:
  - Row 1: Dropdown "EFD\_Tier0\_R5", percentage input "50".
  - Row 2: Dropdown "FC\_450\_15k\_Tier1\_R1", percentage input "100".
  - Row 3: Dropdown "SATA\_Tier2\_R6", percentage input "50".
- Associate Storage Group(s)**: A button.
- OK**, **Cancel**, and **Help**: Three buttons at the bottom right.

## Associate the Storage Group

Consider the following before associating Storage Groups to a FAST Policy:

- Once the percentages are determined the policy is associated to a Storage Group. In this case, the Storage Groups are existing auto-provisioning Storage Groups.
- Note that a **Group Priority** is assigned. The **Group Priority** determines which application (Storage Group) will receive resources on a Storage Tier in case of a conflict or overprovisioning.

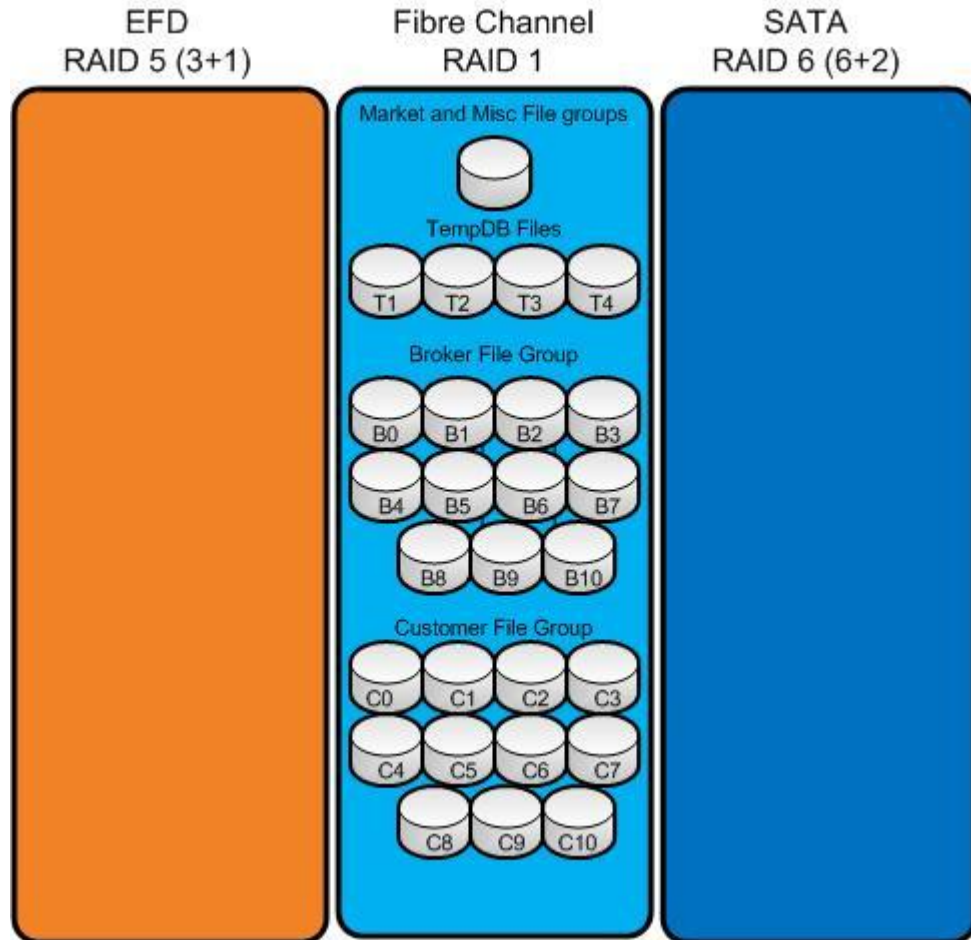


## Validation of FAST technology with SQL Server OLTP workloads

### Data allocation and performance metrics before FAST

An OLTP performance baseline was determined prior to enabling FAST based on the initial storage configuration, as detailed in the diagram and table below.

As shown, the OLTP application was originally configured with all of the data on the FC Symmetrix tier. Each disk in the figure represents a metadevice with eight members.



Descripton	Value
SQL Server CPU utilization	60%
OLTP Transactions Per Second	2,184.86
Workload I/O Read/Write Ratio	80:20
Workload IOPS	23,000

---

**IOPS and disk latency recorded before FAST**

IOPS and disk latency statistics for the broker file group before implementing FAST are detailed in the following graph.



- Most of the disk activity during the simulated user load was skewed to the broker file group. The total workload was approximately 23,000 IOPS with 13,585 IOPS targeted on the 11 LUNs in the broker file group.
- The graph above represents the I/O load on the 11 LUNs in the broker file group. The IOPS and disk latency values for each of the individual LUNs are listed in the [Brokerage table comparison chart](#). Individual LUNs are identified as B0 through B10.
- The average workload for the broker file group is near 1,300 IOPS with disk latencies in the range of 4 ms to 10 ms.

Additionally, see the [Brokerage table comparison chart](#) for a detailed comparison of IOPS and disk latency achieved before and after FAST is introduced into the test environment.

---

**FAST configuration change plan for the test environment**

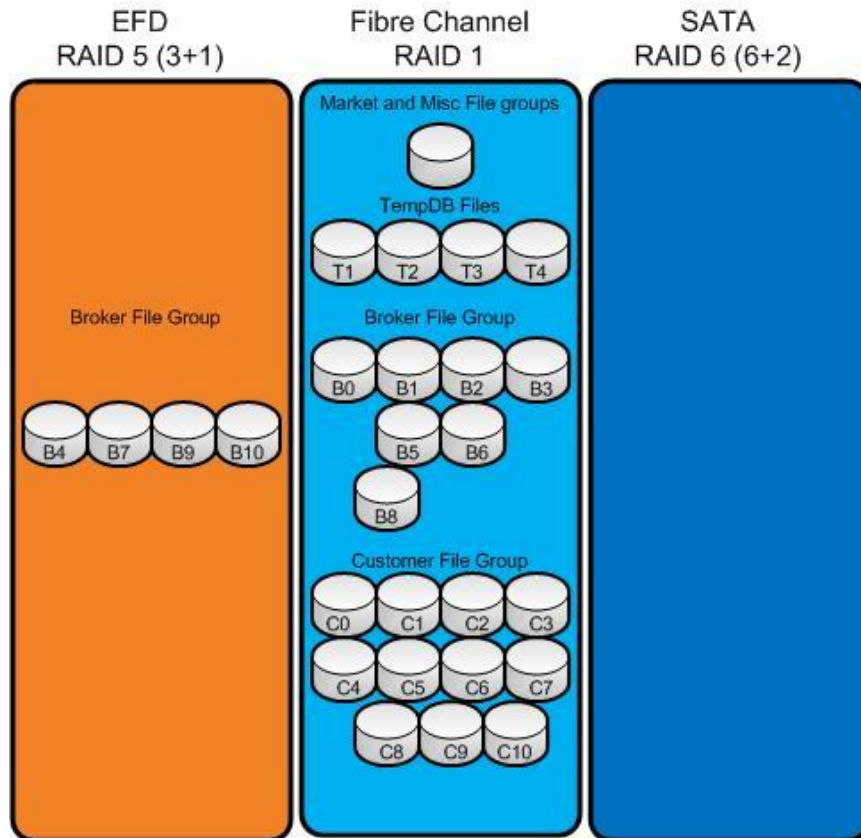
Once the FAST controller is enabled, initial workload analysis indicates that:

- The configuration change plan recommends moving LUNs B4, B7, B9, and B10 to the EFD\_R5 Symmetrix Tier. These are the volumes with the highest IOPS and disk latency numbers.
- Since LUNs B4, B7, B9, and B11 are eight member metadevices, FAST has reached the limit for concurrent disk moves (32).

Additionally, after the configuration change plan completes relocating data, the effect of FAST automatic balancing is seen; as detailed in the [Data allocation and performance metrics after FAST](#) section of this white paper. For example, the disk configuration after the first configuration change plan is approved.

**Data allocation and performance metrics after FAST**

As shown, the OLTP application that was originally configured with all of the data on the FC Symmetrix tier, has now relocated LUNs B4, B7, B9 and B10 to the high performing EFD tier.



Descriptor	Before FAST	After FAST
SQL Server CPU utilization	60%	70%
OLTP Transactions Per Second	2,184.86	2,843.21
Workload I/O Read/Write Ratio	80:20	80:20
Workload IOPS	23,000	25,000

---

**IOPS and disk latency recorded after FAST**

IOPS and disk latency statistics for the broker file group after FAST is implemented are detailed in the following graph.



- Most of the disk activity during simulated user load remains skewed to the broker file group. The total workload totaled approximately 25,000 IOPS with 14,842 IOPS targeted on the 11 LUNs in the broker file group. The broker file group IOPS increased by 1,257 IOPS and the TPS increased by 685.
  - The graph represents the I/O load on the 11 LUNs in the broker file group. The IOPS and disk latency values for each of the individual LUNs are listed in the [Brokerage table comparison chart](#).
  - The individual LUNs are identified as B0 through B10.
  - The average workload for the broker file group is near 1,400 IOPS with disk latencies in the range of 1 ms to 4 ms.
-

---

## Conclusion

---

### Brokerage table comparison chart

In this environment, the impact of implementing FAST yielded immediate, impactful performance results, as illustrated in the table and following results summary.

LUN	Before FAST			After FAST		
	IOPS	Latency	Technology/RAID type	IOPS	Latency	Technology/RAID type
B0	655	4 ms	FC 15k/RAID 1	704	4 ms	FC 15k/RAID 1
B1	1,260	4 ms	FC 15k/RAID 1	1,389	4 ms	FC 15k/RAID 1
B2	1,269	4 ms	FC 15k/RAID 1	1,411	4 ms	FC 15k/RAID 1
B3	1,267	4 ms	FC 15k/RAID 1	1,397	4 ms	FC 15k/RAID 1
B4	1,275	6 ms	FC 15k/RAID 1	1,403	1 ms	EFD/RAID 5
B5	1,274	4 ms	FC 15k/RAID 1	1,396	4 ms	FC 15k/RAID 1
B6	1,270	4 ms	FC 15k/RAID 1	1,391	4 ms	FC 15k/RAID 1
B7	1,274	6 ms	FC 15k/RAID 1	1,404	2 ms	EFD/RAID 5
B8	1,267	4 ms	FC 15k/RAID 1	1,400	4 ms	FC 15k/RAID 1
B9	1,277	9 ms	FC 15k/RAID 1	1,408	2 ms	EFD/RAID 5
B10	1,497	10 ms	FC 15k/RAID 1	1,539	2 ms	EFD RAID 5
<b>Totals</b>	13,585			14,842		

- In the initial configuration, the broker file group partitioned on LUNs B4, B7, B9 and B10 were experiencing high disk latency times with the highest level of application activity.
- The FAST controller correctly identified the high activity LUNs and relocated them to a higher performing Storage Type (EFD).
- The [Brokerage table comparison chart](#) shows a significant drop in disk latency for LUNs B4, B7, B9 and B10. The most significant drop being LUN B10, which became 8 ms faster.
- Relocating the four LUNs with high disk latency times increased the number of IOPS being serviced by the same number of simulated users by 1,257 IOPS in the broker file group.
- The positive effect of relocating these four LUNs also provided benefits for the LUNs on the FC Storage Type. Moving the demand to another Storage Type yielded higher IOPS on the remaining LUNs.

---

## Findings

Interoperability and performance testing confirms that:

- Database environments are good candidates for the application of FAST technology. FAST effectively supports dynamic applications where some of the storage is heavily utilized, some is moderately utilized and some is barely utilized.
- Comparing the before and after data in the [Brokerage table comparison chart](#) shows that FAST improved IOPS by 10 percent, and average disk latency by nearly 50 percent for the broker file group.
- With easy configuration and very little maintenance to the automation parameters, FAST can effectively make intelligent recommendations and nondisruptively relocate data to appropriate storage resources.
- FAST is fully interoperable with all Symmetrix replication technologies—SRDF<sup>®</sup>, TimeFinder<sup>®</sup>/Clone, TimeFinder/Snap, and Open Replicator. Any active replication on a Symmetrix device remains intact while the device is being moved or swapped. Similarly, all incremental relationships are maintained for the moved or swapped devices.

---

## Summary

This white paper highlights the ease of use, efficiency and resource savings realized by utilizing the FAST capabilities available on the Symmetrix VMAX storage system.

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

- FAST helps customers maximize the benefits of Storage Tiering by automating the process of monitoring, analyzing, and relocating SQL Server application data.
- FAST technology can save hours of storage administration as compared to manual tiering.
- The FAST controller proactively monitors device workloads and performs analysis to better utilize resources on the Symmetrix VMAX.
- FAST relocates storage resources dynamically to ensure reliable performance levels for mission-critical applications.

---

## References

The following EMC documents provide supporting information.

---

### Technical notes and white papers

For additional information, see the white papers listed below.

- *Implementing Fully Automated Storage Tiering (FAST) for EMC Symmetrix VMAX Series Arrays Technical Note*
- *EMC Symmetrix DMX-4 Enterprise Flash Drives with Microsoft SQL Server Databases—Applied Technology white paper*
- *EMC Symmetrix Optimizer—A Detailed Review white paper*

---

**Product  
documentation**

For additional information, see the product documents listed below.

- *EMC Symmetrix Management Console (SMC) online help* (integrated with the Symmetrix VMAX)
  - *EMC Solutions Enabler Symmetrix Array Controls Command Line Interface guide*
-