

# White Paper

---

## **Automate and Optimize a Tiered Storage Environment – FAST!**

*By Bob Laliberte*

December, 2009

---

## Contents

Introduction .....	3
The Benefits of Tiering.....	4
Why Aren't More Enterprises Tiering More Effectively? .....	5
EMC's FAST and Ionix Solutions can Help.....	6
The Bigger Truth .....	8


All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482-0188. This ESG White Paper was developed with the assistance and funding of EMC.

## Introduction

Organizations of all sizes are actively transforming and virtualizing data centers to create a more dynamic and responsive infrastructure with the goal of delivering higher levels of service and performance to the business at lower costs. Many have already started the journey by virtualizing the server environment and building out networked storage environments. In fact, a large number, almost half, of these organizations have progressed to the point where they are running Tier 1 production applications in these environments (see Figure 1).<sup>1</sup>

**Figure 1. Organizations Deploying Tier 1 Applications in Virtualized Environments**

**Would you say that your organization currently runs "Tier 1" applications on virtual machines? (Percent of respondents, N=365)**



Source: Enterprise Strategy Group, 2007.

As organizations transfer more mission critical Tier 1 applications into virtual environments, it is imperative that the infrastructure be able to dynamically deliver the requisite performance and availability. Equally important, operations teams will need the appropriate management tools to deliver the visibility and control required for these environments. Additionally, as organizations consolidate multiple application tiers, the infrastructure will need to dynamically differentiate service levels to optimize performance, share resources, and reduce overall costs. In order to do that, innovative solutions are needed to drive higher levels of automation at every level.

In the server domain, companies like VMware have created technology like the Distributed Resource Scheduler (DRS) to detect when a VM's host resources are being overloaded and automatically move it to another host with better performance characteristics. In the storage domain, the concept of differentiating levels of performance for a variety of applications has traditionally manifested itself as storage tiering.

Storage tiering has evolved from early implementations based on multiple storage arrays providing different levels of service to tiering within the array. Moving to the array was a big step in right direction for many as it offered flexibility to deploy a mix of Flash drives for ultra high performance, Fibre Channel (FC) for high performance, and Serial ATA (SATA) for high capacity/low cost storage requirements.


The purpose of this paper is to describe the next major phase in the evolution of storage tiering and outline how EMC Symmetrix V-Max solutions employing Fully Automated Storage Tiering (FAST) and Ionix ControlCenter can enable organizations to simultaneously improve application service levels and reduce costs—automatically.

<sup>1</sup> Source: ESG Research Report, *The Impact of Server Virtualization on Storage*, December 2007.

## The Benefits of Tiering

Most organizations recognize the value of tiered storage. ESG’s research<sup>2</sup> (see Figure 2) indicates that over the next two years, tiering storage will among IT’s top five initiatives.

Figure 2. Most Important Storage Related Initiatives Over the Next 24 Months



Source: Enterprise Strategy Group, 2009.


Why? What benefits can be derived from a tiered storage environment? Respondents believe storage tiering will lead to:

- Higher levels of performance.** Businesses are always looking to gain an edge by ensuring the most demanding applications have the highest level of performance. Traditional magnetic media has not been able to keep pace with performance demands, but with the introduction of Enterprise Flash Drives, organizations can now achieve significantly higher levels. And this is not just for high-performance computing (HPC) environments—organizations would like to apply the benefits of Flash to many different applications. ESG asked 317 enterprises about the applications they thought would most benefit from the use of Flash technologies; the results in Figure 3 indicate that, while HPC environments are important, organizations see more value in deploying it for database/OLTP and e-mail applications. Business intelligence, data mining, and financial apps round out the top five.<sup>3</sup>

<sup>2</sup> Source: ESG Research Report, 2009 Data Center Spending Intentions Survey, March 2009.

<sup>3</sup> Source: ESG Research Report, 2008 Enterprise Storage Systems Survey, November 2008.

Figure 3. Business Applications that Would Benefit from Flash Drives



*Source: Enterprise Strategy Group, 2008.*

- **Reduced budgets for storage.** Even though Flash technology is still relatively expensive, a small number of Enterprise Flash Drives can effectively replace many FC drives that have intentionally been deployed under-utilizing capacity (short-stroked) to realize performance requirements for demanding applications. Additionally, by moving less active data from costly FC drives to high capacity/low cost SATA drives, organizations can drive significant costs out of the equation.
- **Reduced operational costs.** An effective tiering strategy may also positively impact the bottom line by only utilizing power hungry, high performing disks where they are required and energy efficient SATA for the majority of data, reducing power, cooling, and floor space requirements. Enterprise Flash Drives provide the lowest \$/IOPS while SATA drives provide the lowest \$/GB and are both very space and energy efficient.

## Why Aren't More Enterprises Tiering More Effectively?

It's clear that organizations could benefit from tiering their storage more effectively and many have it on their to-do lists over the next two years, so why haven't more already done it? The reality is that most organizations employ some form of tiering, but it tends to be on a macro level. For instance, most IT shops will leverage technology for their production environments that differ from what is used in a test and development setting or in a back up environment. While that is a good start, it leaves a lot of room for improvement.

Being tasked with providing a service to the business, IT needs to ensure that it can meet or exceed required performance levels 24 x7. Manually moving data between tiers can be a time consuming process, especially considering the massive volumes of corporate data that continue to grow exponentially. An organization could require additional headcount in order to effectively tier and manage storage in a dynamic and rapidly growing virtual environment. Specifically, management duties include constantly monitoring data usage and IO activity to identify changing application performance requirements and then scheduling time and spending the effort to move application data between tiers.

This effort has resulted in many IT organizations adopting an "it's all Tier 1" mentality for production. The problem is, as the business grows, so too do expensive Tier 1 storage environments despite the fact that 70-80% of the data housed there has probably not been accessed in the last 90 days.

Another factor holding back tiering is lack of visibility. At the array level, a storage administrator may be aware that only 20-30% of the data is actively being used, but not knowing exactly where the data from each application resides forces them to keep everything in Tier 1. If they move the wrong data, applications as well as company revenue could suffer. Also, again thanks to a universal Tier 1 mentality, most organizations do not invest in performance monitoring solutions. Baseline data may be taken in test and development, but once satisfied with the system, most IT organizations only worry about performance when the phone rings.


## EMC's FAST and Ionix Solutions can Help

EMC pioneered Flash technology in enterprise arrays and currently offers multiple tiers of storage consisting of different drive and protection types within a single array. In order for organizations to take full advantage of these multi-tier environments, it created an automated, policy-based software solution called FAST (Fully Automated Storage Tiering) and linked it to existing Ionix management solutions. EMC has recognized that, while the scale and scope of storage is increasing, administrators are forced to reduce their labor requirements yet still somehow maintaining complete control over complex, rapidly growing infrastructures. Dynamic workloads make it difficult to provide predictable, consistent performance levels. What these administrators need is an underlying infrastructure that provides full visibility, is self-managing, and is able to quickly adapt to change in tiered environments

FAST meets these challenges by helping simplify the management of the storage infrastructure in order to reduce storage costs and realize the full benefits of tiering. It does this by:

- **Leveraging technology designed to automate the dynamic allocation and relocation of data across different storage types based on the changing performance requirements of applications.** FAST helps customers maximize the value of in-the-box tiered storage by placing the right data on the right storage type at the right time. It will proactively monitor workloads and automatically move more frequently used data to higher-performing Flash drives and less frequently accessed data to higher-capacity SATA drives. This is all accomplished without any disruption to host operations, business continuity, or availability processes. Figure 4 illustrates the dramatic improvement that FAST tiering can have on an environment. On the left side of the figure, the heat map represents a Symmetrix V-Max with all its FC drives supporting an active ESX cluster. The red and orange areas indicate heavily utilized disks. On the right is the same environment after adding only 4 % Flash drives and applying FAST. EMC claims that in this environment, it created 68% less disk IO contention and enabled 2.5 times faster disk responses, as shown by the blue and green squares in the array.

Figure 4. Before and After Array Performance Mapping



Source: EMC, 2009.

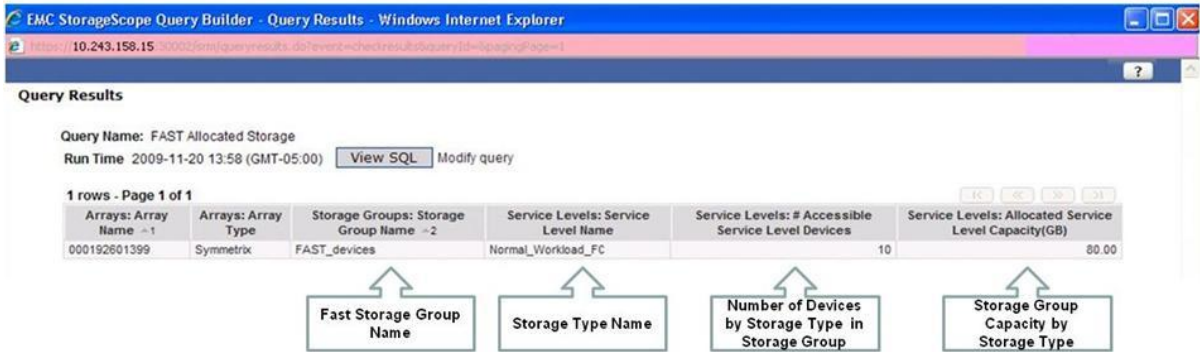
- **Enabling IT organizations to create different storage types or tiers that can be shared across multiple applications.** Once the tiers are set up, a policy can be created and associated to applications. FAST will manage against the policies to dynamically meet the service level requirements of applications.
- **Reducing total cost of ownership on several levels.** Customers can cut acquisition costs by using a small number of Flash drives to support the same IOPS that normally would require many short-stroked hard disk drives, at lower response times. Acquisition costs are also lower because more data can be placed on SATA drives at any given time based on usage frequency. And by reducing their drive count and maximizing use of Flash and SATA drives, customers can meet capacity and performance requirements while reducing power, cooling, and floor space requirements.

EMC has integrated the FAST technology with existing products from the Ionix and Symmetrix software families to deliver a solution that provides greater visibility, is easy to use, and requires minimal deployment and integration time. The result is:

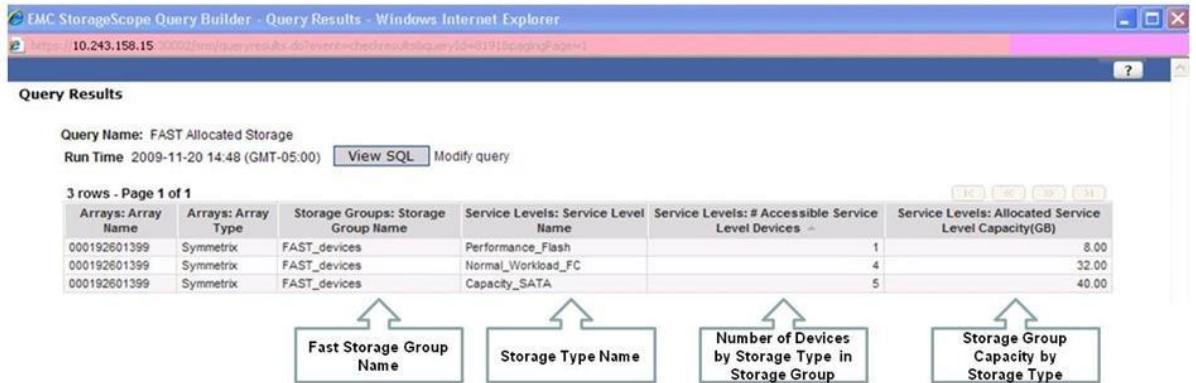
- **Greater visibility.** The combination of Ionix ControlCenter Storage Scope and FAST enables organizations to discover and report on service levels and capacity consumption of each application. Leveraging Ionix Control Center 6.1, organizations can build chargeback models for differentiated storage service levels. Additionally, the ability to track storage consumption by application helps storage administrators more effectively plan capacity upgrades. Symmetrix Performance Analyzer (SPA) has also been enhanced to support FAST and provides the storage administrator a view of performance trends (response time and throughput) by application as FAST dynamically moves applications between tiers. Ionix ControlCenter also provides visibility into the virtual machines so administrators can better understand the connection from the applications to the storage even when the application's data is spread out across multiple classes or tiers of storage.
- **Minimal deployment and integration times.** The key to implementing new technology is to make it easy to use while offering the user as much control as needed. EMC's Symmetrix Management Console (SMC) has been enhanced with a "FAST" wizard to help an IT organization to easily create tiers, policies, and associate policies to applications. The wizard minimizes the amount of time required to deploy and integrate FAST into existing environments. Additionally, FAST can run in fully automatic mode, but it also has a user approval mode. The user approval mode has all the same functionality, but before actually making any application moves between tiers, it will alert administrators so they can review and verify any potential changes before committing them to the system. ESG believes this "walk before you run" approach will actually accelerate the adoption of the fully automatic mode by establishing a pattern of trust. The software installs quickly and integrates seamlessly with existing environments and, depending on the size and complexity of an environment, FAST can discover and map an entire array and begin dynamically tiering in a matter of hours. Initial setup is as simple as creating tiers, defining the policies, and linking to the applications.
- **Enhanced reporting capabilities.** By integrating FAST with EMC management solutions, organizations can view detailed information regarding tiering and their configurations. The new reporting capabilities in Ionix StorageScope allow organizations to quickly understand how the storage environment is configured and to what level it is optimized. These reports include: raw and available, allocated, and accessible capacity by storage tier. This enables storage administrators to create chargeback reports, providing business owners the ability to realize the savings attained by moving to a tiered storage infrastructure. It also helps with capacity planning by providing a view into available resources and capacity consumption by storage type. And, it offers visibility into configuration changes over time to help storage teams make more effective tiering policy decisions. Figure 5 illustrates Ionix StorageScope reports that show capacity of a storage group before and after a FAST data move. In this example, FAST has moved performance sensitive data to Flash drives to enhance performance levels while moving less performance sensitive data to lower cost and lower power SATA drives.

Figure 4. Ionix ControlCenter Tracking Capacity Consumption by Storage Type

Before the move, StorageScope shows all data in the storage group, FAST\_devices, are on a storage type named Normal\_Workload\_FC which contains 15K rpm fibre drives



After the move, StorageScope shows the data in the storage group FAST\_devices are now on three storage types which include one 8GB volume of Flash, four 8GB volumes of fibre, and eight 8GB volumes of SATA.



## The Bigger Truth

Storage tiering has been around for a while and while it has had some success, it has largely been a labor intensive and manual process. Very large organizations have been able to generate some savings, but the potential exists to drive even more savings—both CAPEX and OPEX—and deliver higher levels of service. The keys to accelerating tiered storage deployments and efficiencies are based on having:

- Automated solutions.** Given today’s rate of storage growth and the need for the storage environment to maintain pace with a dynamic server environment, the only way to optimize is through automation. Trying to manually create and continuously optimize a tiered storage environment is not scalable. Organizations need to eliminate as many manual, time consuming activities as possible if they want to manage a rapidly expanding environment with the same number of resources
- End-to-end visibility.** Visibility includes understanding the interdependencies of application and storage tiers. Doing so allows an organization to better prioritize performance for mission critical business applications. Given the widespread and growing adoption of server virtualization solutions, it is imperative that storage teams be capable of identifying and tracking virtual machine relationships to storage. In addition to a detailed view, in order to provide effective chargeback, an enterprise level view is required. This would include not only Tier 1 storage, but should also scale to cover Tiers 2, 3, and so on.
- Dynamically tiered environments.** An environment has to be dynamic in order to extract the full benefits of storage tiering. Because we know that workloads can vary seasonally, monthly, or even daily and the data an application creates may be idle after 60 to 90 days, the storage infrastructure must be able to adjust to meet changing performance requirements.

By delivering Fully Automated Storage Tiering and end-to-end visibility through its Ionix and Symmetrix management solutions, EMC is enabling organizations to realize the full benefits of dynamic storage tiering. Its customers will have greater control, reduced operational and capital costs, and improved performance for a wider range of applications. Given EMC's focus on the journey to the cloud, one could also reasonably expect that eventually FAST will also enable dynamic tiering to cloud storage—driving costs down even further.



Enterprise Strategy Group | **Getting to the bigger truth.**