

REDUCE COSTS AND OPTIMIZE MICROSOFT SQL SERVER PERFORMANCE IN VIRTUALIZED ENVIRONMENTS WITH EMC SYMMETRIX VMAX

An Architectural Overview

EMC GLOBAL SOLUTIONS

Abstract

This white paper demonstrates a solution to validate the use of Fully Automated Storage Tiering for Virtual Pools (FAST VP) in a virtualized Microsoft SQL Server 2008 R2 environment, supporting an online transaction processing (OLTP) workload. The solution also validates VMware® vStorage API for Array Integration (VAAI) with sub-LUN automated and EMC® Virtual Storage Integrator (VSI) capabilities for self-service provisioning of storage for VMware environments.

February 2011



Copyright © 2011 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate 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.

VMware, ESX, ESXi, vMotion, VMware vCenter, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

Part Number h8037.1

Table of Contents

Executive summary	6
Introduction to the EMC Symmetrix VMAX and Enginuity 5875.....	6
Business case	6
Solution overview.....	7
Key results	7
Introduction	8
Purpose.....	8
Scope.....	8
Audience.....	9
Terminology	9
Technology overview	11
Introduction to the key components	11
EMC Symmetrix VMAX storage array	11
EMC Symmetrix Management Console.....	11
FAST VP	12
EMC Virtual Storage Integrator.....	12
VMware VAAI.....	12
Configuration	14
Configuration overview.....	14
Physical architecture	14
Environment profile	15
Hardware resources.....	15
Software resources.....	16
SQL Server 2008 R2 test application	16
PowerPath/VE configuration.....	16
Storage design consideration	16
Storage design configuration.....	17
VMware vSphere configuration.....	17
SQL Server 2008 R2 configuration	18
Physical device configuration	19
SQL server logical layout	19
SQL storage design without FAST VP.....	20
SQL storage design with FAST VP.....	20
FAST VP	23
Introduction to FAST VP	23
How FAST VP works	23
FAST VP configuration.....	23

Management tools	24
Set Policy time windows	24
Create a virtual pool tier	25
Create a policy.....	25
Associate a storage group	26
FAST VP test and validation.....	28
FAST VP test introduction.....	28
Test application I/O characteristics	28
Baseline performance statistics before FAST VP	29
Migrating into multi-tier thin pool	29
Implementing FAST VP on a SATA and Flash multi-tier thin pool.....	30
Storage tiering.....	30
Performance statistics after FAST VP data relocation on the Flash and SATA tiered storage30	
ESX server performance statistics on tiered storage with FAST VP enabled	31
SQL server performance statistics on tiered storage with FAST VP enabled	32
Performance comparison.....	33
Cost comparison	34
VAAI validation	36
VAAI test introduction.....	36
VAAI enabled SUB-LUN cloning.....	36
VAAI test conclusion.....	36
VSI integration validation	37
VSI introduction	37
VSI validation	37
VSI test conclusion	37
Pool management via SMC	38
Pool LUN expansion and shrinkage management via SMC.....	38
Shrink thin pool	38
Expand thin pool	39
Thin pool monitor - Set pool utilization threshold via SMC.....	40
Pool LUN management via SMC.....	40
Conclusion	41
Summary.....	41
Findings	41
References	42
White papers.....	42
Other documentation	42
Additional information.....	43
Overview	43

Baseline ESX server performance statistics	43
Baseline SQL server performance statistics	45

Executive summary

Introduction to the EMC Symmetrix VMAX and Enginuity 5875

EMC® Enginuity™ 5875 is the latest release for the Symmetrix VMAX™. This Enginuity version on Symmetrix VMAX carries the extended and systematic feature development forward from previous Symmetrix generations. It further advances the ability of EMC self-optimizing intelligence to deliver performance, array tiering, availability, and data integrity that now define advanced storage functionality. A prerequisite for complex, demanding, risk-intolerant IT infrastructures, Enginuity—coupled with Symmetrix VMAX—is the essential foundation technology for delivering cost-effective high-end storage services.

This solution uses new and enhanced features of Symmetrix VMAX with Enginuity 5875 to validate storage efficiency and scale for organization growth:

- EMC Symmetrix® Fully Automated Storage Tiering for Virtual Pools (FAST VP)—providing better optimization of high-performance Flash drives, and better responsiveness to changes in data activity through sub-LUN automated storage tiering.
- VMware® vStorage API for Array Integration (VAAI)—enabling storage administrators to maintain control of the infrastructure and managing virtual server storage.
- EMC Virtual Storage Integrator (VSI)—providing rich functionality in easing the management of storage and providing a simple and effective view of storage layout at all layers.

Business case

For database administrators (DBAs) and SAN architects, it is a constant, critical business challenge to maintain or improve the performance of a company's Microsoft SQL databases in both physical and virtual server environments.

In many data centers today, the most common solution for constrained storage performance is to order additional Fibre Channel (FC) drives that have to be formatted and configured before SQL databases can be restriped across the new LUNs. This increases costs, takes a longer time to deploy, and drives up power, cooling, and rack space requirements for the overall solution.

In addition to the high acquisition and operational costs of the drives, customers who perform these procedures need to spend hours, and potentially days, in analyzing the current environment, designing the new layout, implementing the solution, and managing it over time.

EMC Symmetrix manual storage tiering and Virtual Provisioning™ are good ways to reduce costs and eliminate bottlenecks.

Now with FAST VP, storage tiering and Virtual Provisioning can be built into a company's application deployment model and can be fully automated. FAST VP can greatly improve performance while significantly reducing storage-related costs.

Solution overview

This solution focuses on proving the feasibility of FAST VP in an online transaction processing (OLTP) environment with virtualized Microsoft SQL server environments. A multi-tiered storage pool with high-performance Flash drives and cost-effective Serial Advanced Technology Attachment (SATA) drives is demonstrated to provide higher I/O requests per second (IOPS) and lower cost-per-gigabyte (GB) than a commonly used single-tier FC environment.

FAST VP automates the optimization of the back-end storage for SQL server databases, relocating hot and cold data to the appropriate tier at a sub-LUN level. It eliminates the need for manual intervention by the database and storage administrator while providing much more flexibility and finer granularity for database optimization.

The new VMware VAAI for the Symmetrix function in this solution is also demonstrated by storage vMotion® and cloning before and after VAAI is enabled. With VAAI:

- Hardware acceleration on the storage array is enabled, allowing storage migrations within the SAN and off the network.
- CPU processing is significantly reduced for both the array and ESX® central processing unit (CPU).

EMC Virtual Storage Integrator (VSI) enables self-service for the VMware administrator. It makes storage management in a VMware environment (such as finding the correlation between storage, virtual machines, and LUNs) much easier.

Key results

A key validation proof point for this solution was to perform a comparison between a single-tiered storage configuration and a multi-tiered storage configuration.

Moving from a baseline of a single-tiered FC drive configuration to a multi-tiered Flash drive and SATA drive configuration can significantly reduce costs. Compared with the costs in a single-tiered environment, the costs in a multi-tiered configuration are greatly reduced, as shown in Table 1.

Table 1. Cost and performance comparison

Configuration	All FC	Tiered FLASH+SATA
Drives	64	32 (4 Flash+28 SATA)
Acquisition costs	baseline	32% less
Power/cooling costs	baseline	60% less
Management time	baseline	Greatly reduced
SQL TPS	baseline	5% improvement
SQL IOPS	baseline	8% improvement

Note Acquisition costs and power/cooling costs are calculated for drive only.

Introduction

Purpose

The purpose of this solution is to validate the use of FAST VP in a virtualized Microsoft SQL Server 2008 R2 enterprise-class environment supporting multiple concurrent OLTP-type workloads. The solution also validates the VAAI feature with sub-LUN automatic migration and replication availability, and the EMC VSI capability of self-service for storage administration in VMware.

Scope

This solution demonstrates how to use storage pools for auto-tiering at a sub-LUN level (sub-LUN is an extent of data within the LUN), and can be as granular as one chunk (or 768 KB) with a simple and fully automated storage tiering mechanism.

A storage pool with two tiers (Flash and SATA drives) is configured to run the same load of a commonly used FC-only storage configuration. This solution highlights the simple approach to storage tiering that FAST VP enables.

- Comparing database performance with:
 - Locating data on a single FC drive thin pool (also known as a virtual pool)
 - Locating data on storage with both Flash and SATA drive thin pools
- Showing auto-tiering through scheduler—The time windows are used by FAST VP to specify when data can be collected for performance analysis and when moves can execute.

The solution also demonstrates the improved granularity of Enginuity 5875, showing how FAST VP's sub-LUN tiering is capable of placing hot (very frequently accessed) and cold (less actively accessed) data across different storage tiers at a sub-LUN level. This only requires a single pool LUN rather than additional LUNs to be created within each tier.

EMC Symmetrix Enginuity 5875 microcode also provides the following key features:

- Thin provisioning
- VAAI
- VSI
- Sub-LUN cloning using VAAI with TimeFinder®/Clone/Snap
- Pool LUN expansion
- Pool LUN shrink

Audience

The intended audience for the white paper is:

- Customers
- EMC partners
- Internal EMC personnel

Terminology

Table 2 defines terms used in this document

Table 2. Terminology

Term	Definition
Data device	Virtual Provisioning term for devices (not mapped to the host) that provide physical storage for thin devices. Data devices must be contained in a virtual pool before they can be used.
Disk Transfers/sec	The rate of read and write operations on the disk.
Enginuity	The operating environment that provides the intelligence that controls all components in a VMAX array.
Fully Automated Storage Tiering (FAST)	A feature of Symmetric VMAX that automates the identification of data volumes for the purpose of allocating or reallocating business application data across different performance and capacity tiers within the storage array.
FAST VP	Fully Automated Storage Tiering for Virtual Pools supports automatic storage tiering at the sub-LUN level. VP denotes virtual pools that are Virtual Provisioning thin pools.
FAST VP policy	Contain a set of tier usage rules that defines how an application's disks are most effectively distributed across storage tiers.
Logical Unit Number (LUN)	A unique identifier used to distinguish storage devices.
PowerPath®	EMC server-resident software that enhances performance and application availability by supporting multiple I/O paths to logical devices, and load balancing. It also provides automatic failover.
Response time	The interval of time between submitting an I/O request and receiving a response.
Serial Advanced Technology Attachment (SATA)	A hard disk technology that focuses on greater data density, less power consumption, and less cost with lower data transfer rates.
SMAS	Symmetrix Management Application Server is a combined product installer for Symmetrix Management Console (SMC) and SPA.
SMC	Symmetrix Management Console is a powerful and intuitive GUI that configures and manages multiple Symmetrix arrays. It presents the functionality of the SYMCLI in a browser interface and simplifies storage administration tasks through the use of built-in wizards. It includes the functionality for configuring FAST VP.

Thin device (TDev)	A logical device that is configured and presented to the host that is allocated physical storage as needed. The physical storage used comes from the virtual pool to which the thin device is bound. Thin devices look like and interact with the host just like standard devices with configured capacity.
Virtual LUN (VLUN)	Virtual LUN technology enables the migration of data between storage tiers within the same array. VLUN provides the ability to move data without disrupting applications.
Virtual LUN VP Mobility	Virtual LUN VP Mobility enables administrators to manually move thin LUNs between pools. This new feature of Enginuity 5875 includes the ability to move a thin volume's chunks from multiple virtual pools to a single pool, regardless of the underlying disk technology or RAID protection. Movements are completely transparent to the environment, having no host or application impact.
Virtual Provisioning	EMC's implementation of thin provisioning.

Technology overview

Introduction to the key components

This section provides an overview of the technologies that are used in this solution:

- EMC Symmetrix VMAX storage array
- EMC Symmetrix Management Console
- FAST VP
- EMC Virtual Storage Integrator
- VMware VAAI

EMC Symmetrix VMAX storage array

EMC Symmetrix VMAX with Enginuity version 5875, with the strategy of simple, intelligent, modular storage, incorporates a new, highly scalable Virtual Matrix Architecture™ that enables VMAX arrays to grow seamlessly and cost-effectively from an entry-level configuration into the world's largest storage system. It offers:

- More efficiency: New zero-downtime technology for migrations (technology refreshes) and lower costs with automated tiering.
- More scalability: Up to twice the performance, with the ability to manage up to 10 times more capacity per storage administrator.
- More security: Built-in encryption, RSA-integrated key management, increased value for virtual server and mainframe environments, replication enhancements, and a new eLicensing model.

The tiered storage configuration used in the test environment is based on the following VMAX features:

- Sub-LUN FAST VP
- Virtual LUN VP Mobility
- VSI

EMC Symmetrix Management Console

The Symmetrix VMAX storage system provides a built-in web browser interface, Symmetrix Management Console (SMC). SMC provides a centralized management to the entire VMAX storage infrastructure. In the context of FAST, SMC integrates easy-to-use wizards to:

- Create thin pools
- Create storage tiers
- Define and associate storage groups
- Setup FAST policies

FAST VP

FAST VP is a new technology introduced in the latest microcode version (5875). FAST VP combines a "sub-volume" auto-tiering technology from EMC and leverages Virtual Provisioning technology. It enables storage administrators to implement automated policy-driven plans that perform dynamic nondisruptive changes to storage layouts of different applications by ensuring that the hot spots of a volume or LUN are served by high-performance drives and the inactive data is served by cost-effective drives. With FAST VP, customers can achieve:

- Maximum utilization of Flash drives for high-performance workloads
- Lower cost of storage by placing the less accessed data on SATA drives
- Better performance at a lower cost, requiring fewer drives, less power and cooling, and a smaller footprint
- Radically simplified automated management in a tiered environment

FAST VP can move data among (up to) three tiers (Flash drive, FC, and SATA drive) to meet the performance and capacity demands of a broad range of applications. Frequently, accessed data is moved to or kept at proper storage tiers, based on the access patterns of sub-volumes and the defined FAST policy. Based on the changing performance requirements of applications, FAST VP only promotes those active hot spots of a volume or LUN to high-performance drives such as Flash drives, but not the entire volume or LUN. At the same time, FAST VP also moves less accessed portions of a volume or LUN to low-cost drives such as SATA drives. Customers thus get the best of both worlds: high performance and low cost.

EMC Virtual Storage Integrator

EMC Virtual Storage Integrator (VSI) is a VMware vCenter™ plug-in available to all VMware users with EMC storage in their environment.

VSI simplifies mapping VMware vSphere™ data stores to LUNs and NFS shares on EMC storage, and helps pinpoint the location of virtual machines and raw device mapping files on the array. The VMware administrator's visibility into the storage layer can help troubleshoot storage performance issues and simplify communication among server, storage, and virtualization teams.

VMware VAAI

VMware vStorage APIs for Array Integration (VAAI) enable VMware vSphere to offload certain operations to Symmetrix VMAX or other arrays, thereby accelerating tasks, supporting high-scale environments, and freeing up server resources for other activities. VAAI includes three components with this release:

- Full Copy— Enables the storage array to directly copy data within the array without requiring the ESX server to read and then write the data. With the feature enabled, creation of VMs and Storage vMotion with VAAI enabled could be up to 10 times faster.
- Block Zero—Initializes the file system block and virtual disk space more rapidly. With the feature enabled, the common tasks, such as Virtual Machine File System (VMFS) formatting and relocation, could have up to 10 times fewer I/O.

- Hardware-Assisted Locking—Enables more efficient metadata updates and assists virtual desktop deployments. The feature could support up to 10 times more virtual machines per data store.

Configuration

Configuration overview

The solution design implements the following physical components:

- Two VMware ESX servers forming a VMware cluster and containing:
 - Four SQL Server 2008 R2 instances on four Windows 2008 R2 virtual machines
 - One virtual machine hosting the SMC server components
 - One virtual machine hosting the vCenter server components
- One VMware ESX server contains four test client virtual machines
- The Symmetrix VMAX is attached to the servers through four of its ports
- Fibre Channel connectivity is provided by a 4 Gb/s director-level enterprise-class switch
- The Ethernet backbone is a Gigabit Ethernet network switch
- The Symmetrix VMAX provides Flash, FC, and SATA drive technologies to the OLTP application
- PowerPath/VE is implemented in the ESX servers to provision four paths for the thin device

Physical architecture

Figure 1 illustrates the overall physical architecture of the environment.




Figure 1. Physical architecture

Environment profile

The environment profile is listed in Table 3.

Table 3. Environment profile

Profile	Quantity/Size/Type
Symmetrix VMAX <ul style="list-style-type: none"> FC configuration for baseline FAST VP deployed in a two-tier solution leveraging Flash and SATA drives VAAI for sub-LUN cloning 	1 x Enginuity 5875
User database	4 x TPCE databases totaling approximately 1.4 terabytes (TB)
User database log and tempdb (data and log)	Eight 600 GB 10k RAID 1
SQL Tier 0 RAID type, physical drive size	Four 200 GB Flash RAID 5 (3+1)
SQL Tier 1 RAID type, physical drive size	Sixty-four 450 GB 15k rpm FC RAID 1
SQL Tier 2 RAID type, physical drive size and speed	Twenty-eight 1 TB 7.2k rpm SATA II RAID 1

Hardware resources

Table 4 shows the hardware resources used in this solution.

Table 4. Hardware resources

Equipment	Quantity	Configuration
Symmetrix VMAX	1	Enginuity 5875
		200 GB Flash drives x 4
		1 TB SATA 7.2k rpm drives x 28 (FAST VP)
		450 GB FC 15k rpm drives x 64 (baseline)
		600 GB FC 10k rpm drives x 8 (-tempdb and log files)
FC switches	1	Multilayer SAN switch: 7 switch slots, 48 ports/slot, two slots for engine II, 2/4/8/10 Gb/s FC switching modules
Network switch	1	Layer-three network switch: 241 10/100/1000 Ethernet ports
ESX server	3	4 sockets, quad core, X7350 @ 2.93 GHz, 128 GB RAM

Software resources

Table 5 shows the software resources used in this solution.

Table 5. Software resources

Equipment	Version
EMC PowerPath/VE	5.4 SP1
EMC Solutions Enabler	7.2.0.0
EMC SMAS	7.2.36.4
EMC Virtual Storage Integrator	4.0.0
Microsoft SQL Server 2008 R2 Enterprise	RTM, build 10.50.1600
Microsoft Windows 2008 R2 Enterprise	RTM, build 6.1.7600
VMware Virtual Center	4.1.0, build 233726
VMware Virtual Client	4.1.0, build 233726
VMware vSphere	4.1.0, build 260247

SQL Server 2008 R2 test application

The test environment is based on a TPC-E-like workload. It is composed of a set of transactional operations simulating the activity of a brokerage firm, such as managing customer accounts, executing customer trade orders, and other interactions with financial markets.

PowerPath/VE configuration

PowerPath/VE uses multiple physical path components: host bus adapters (HBAs), switches, storage processors (SPs), and cables between an ESX/ESXi™ host and the VMAX to provide multipath and load balance. If one or more path components fail, the ESX and ESXi host can use a viable alternate path to access the VMAX storage device.

In the solution, four VMAX FA ports are zoned to the four ESX server HBA ports directly for PowerPath/VE to provide multiple paths and load balance.

Storage design consideration

Storage design is an important element to ensure successful deployment for Microsoft SQL Server 2008 R2 on Symmetrix VMAX. The process is essentially the same as a physical environment from a disk perspective. Additionally, in virtual environment, the volume for the operating system of the virtual machine also needs to be accounted in. The virtualized SQL Server 2008 R2 storage design includes a number of different parts including:

- Storage design—Disk requirements and layout.
- Virtual Provisioning—virtual pool design. This provides flexibility to satisfy disk space requirement dynamically for applications.

- Automated Tiering—FAST VP policy design for facilitating sub-LUN movement with performance and cost optimization. This is a policy-enabled technology that automates data movement between the two tiers based on the performance statistics it gathers.

Storage design configuration

The storage layout for this solution is designed to address the IOPS requirement, capacity utilization, and disk performance. Additionally this solution was designed to account for ease of use, which allows for easy configuration and maintenance for storage administrators and DBAs:

- Easy configuration for storage administrators and DBAs:
 - One large thin meta volume for the data store volume.
 - One data store for all database files of the four SQL Server virtual machines.
 - Four virtual machine disk formats (VMDKs) on the same data store for the four SQL Server virtual machines.
- FAST VP policy:
 - The policy enables automated data movement between the two tiers based on the performance statistics it gathers. The FAST VP policy can promote the data to the Flash drive tier or demote it to the SATA tier.
- Thin provisioning: with flexibility to satisfy disk space requirements dynamically for applications:
 - One thin FC pool for all SQL data files for the baseline test.
 - One thin Flash drive pool and one thin SATA pool for all the SQL data files for the FAST VP test.
 - One thin FC meta volume from a separate FC pool.
- FC thin pool: with mirrored protection:
 - Configured for the baseline performance test.
- Flash drive thin pool: with RAID 5 (3+1) protection and added to the Flash drive tier:
 - Configured for the FAST VP performance test.
- SATA thin pool: with mirrored protection and added to the SATA tier:
 - Configured for the FAST VP performance test.

VMware vSphere configuration

Two ESX servers are configured to support four SQL server virtual machines with two virtual SQL servers on each ESX server. Each SQL server has its own data store and virtual disk for its operating system. All virtual disks are of VMDK.

The tempdb and log file of the production database are allocated from a single data store from the same thin device.

The four production databases are put on four virtual disks allocated from a single data store from the same thin device.

SQL Server 2008 R2 configuration

Four Microsoft SQL Server 2008 R2 servers are consolidated through the use of VMware vSphere 4 and Symmetrix VMAX. This virtualized SQL Server 2008 R2 environment provides similar performance and throughput as multiple physical hosts in a reduced footprint. This results in a greater efficiency in the use of server resources and it reduces cost in terms of power and cooling.

The SQL server configuration is based on the following profile in Table 6.

Table 6. SQL Server profile

Profile	Value
Total SQL database capacity	1.4 TB
Number of SQL instances	4 (1 per virtual machine)
Number of user databases per instance	1
Number of virtual machines	4
Type of data store	VMFS
SQL virtual machine configuration	8 virtual processors (vCPUs) with 16 GB memory (no overcommitment)
Concurrent users	Mixed workloads to simulate hot, warm, and cold applications

Windows and SQL Server 2008 R2 configuration of each virtual machine are as follows. The other settings keep SQL default values.

- Grant "Lock pages in memory" to a SQL startup account
- Use a 64 KB NTFS allocation unit size for the user data device
- Four tempdb data files with an equal initialization size for every SQL instance

The four different databases are designed to handle different loads to simulate hot, warm, and cold workloads. Each database has a client to generate a TPC-E-like workload, which attempts to simulate a stock trading brokerage firm. By adjusting the workload, the throughput (measured in transactions per second) of each database can be varying to mimic the hot, warm, and cold workloads of the real world production applications. Database activity is measured by the percentage of TPS for each database out of the total TPS supported by the designed thin device. The database and workload profile is displayed in Table 7.

Table 7. Database and workload profile

Database name	Database profile	Database size	Designed database activity (percentage of total transaction)	Designed to simulate
SQL01	20k user DB	250 GB	>80%	Hot database
SQL02	20k user DB	250 GB	<5%	Cold database
SQL03	30k user DB	350 GB	10%	Warm database
SQL04	50k user DB	550 GB	<5%	Cold database

Physical device configuration

Table 8 and Table 9 show the thin pool configuration details for the environment.

Note The thin device configuration is designed for performance, not for capacity. Therefore, the unused space in the FC, SATA, and Flash drive pools can be used for other applications in a real customer’s environment.

Table 8. Baseline testing with FC thin pools

Thin pool	Tier	Data device size (GB)	Thin device (GB)
SQL_FC_DATA	Fibre Channel	50	1843.2
SQL_TempDB_log_FC	N/A	50	400

Table 9. FAST VP testing with Flash and SATA thin pools

Thin pool	Tier	Data device size (GB)	Thin device size (GB)
SQL_FLASH_DATA	Flash drive	17	1843.2
SQL_SATA	SATA	100	
SQL_TempDB_log_FC	N/A	50	400

SQL server logical layout

Each virtual machine has one user database. The four user databases share two thin devices created on VMAX from separate disk groups. Each database has the same schema and store location but has a different database size. Each virtual machine has two drives to store the user database (on drive R:), and log and tempdb (on drive T:). Drive R: uses the SQL_FC_DATA thin pool in the baseline test and SQL_FLASH_DATA + SQL_SATA tiered thin pool in the FAST VP test, while Drive T: uses the SQL_TempDB_log_FC thin pool in both test scenarios. The FAST VP policy is implemented to monitor the workloads of the thin device that stores the data files only.

Figure 2 shows the overview of relationships of file groups, files, and provisioned meta thin devices to the virtual machine. The user database of SQL01 is implemented to be comprised of several filegroups where each filegroup mapped is to one or more files located on two drives from different meta thin devices.

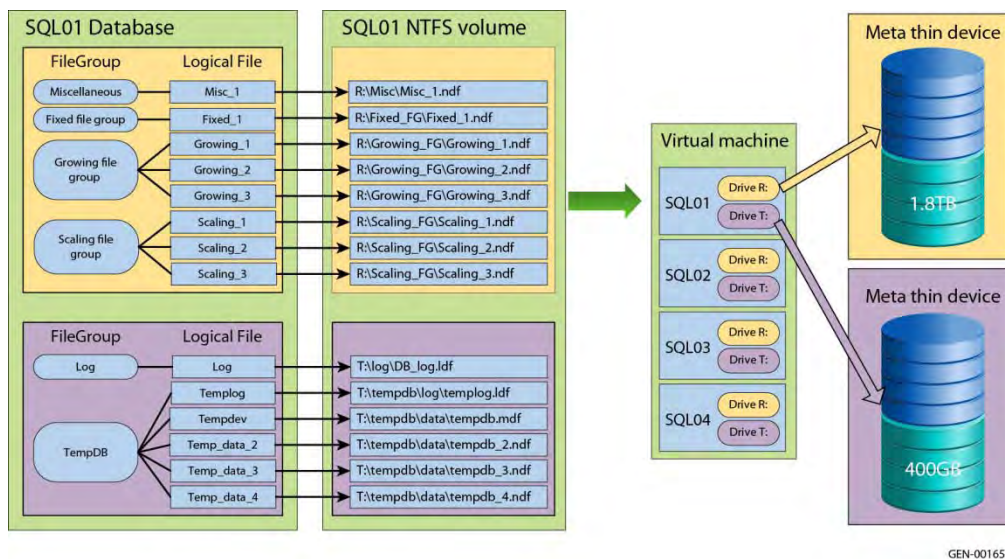


Figure 2. SQL server logical layout

SQL storage design without FAST VP

Using the SAN environment to deploy large databases has many benefits, such as leveraging RAID protection against disk failure and device level replication (snapshot or clone) for fault tolerance.

However, deployment can be very complex and maintenance costs can be high when a DBA wants to use different storage tiers to improve I/O performance.

In order to allocate data in different devices (SATA, FC, and Flash) to satisfy the performance requirement, the DBA needs to first identify the hot and idle tables when the application is running, then lay them on different storage types. Database partitioning is usually needed to further isolate the hot spots and less accessed data in the same table.

When the database activity pattern changes over time, the DBA may need to repeat the previous process to relocate the database in order to maintain performance.

The DBA also needs to consider the pre-allocating space to account for data growth and to prevent the database from running out of space.

SQL storage design with FAST VP

By using FAST VP on VMAX, configuration and deployment of storage tiering and Virtual Provisioning can be made easy to operate.

The SQL database can be configured on thin VMAX devices without going through a detailed storage layout analysis. A FAST VP policy can then be configured to automatically tier the data on the appropriate tier at the right time.

This solution validates that FAST VP can simplify storage design with automated data movement to reduce the total cost of acquisition and maintenance while improving the application performance in general.

The storage design for this solution is as follows:

- One thin device is provisioned to the ESX host to create a single data store.
- This thin device is bound to a single pool that is added as tiers for the FAST VP policy on the VMAX.
- Thin VMDK volumes carved out of the data store are presented to the four SQL virtual machines as storage for database files.

In this solution, the four SQL servers are allocated 1,690 GB of data from a single data store with 1,870 GB of formatted space available. Table 10 shows the virtual machine disk allocation details. Figure 3 shows the process of creating the thin pool, creating the thin device, provisioning to the two ESX servers, cutting the thin device with a different percentage (the pie chart shows the actual allocation percentage of the thin device), and then provisioning to the four virtual machines.

Table 10. SQL virtual machine disk allocation

Data store (VMDK) size (GB)	SQL server instance	Virtual machine allocation (GB)	Percentage
1870	SQL01	320	18.9%
	SQL02	320	18.9%
	SQL03	400	23.7%
	SQL04	650	38.5%

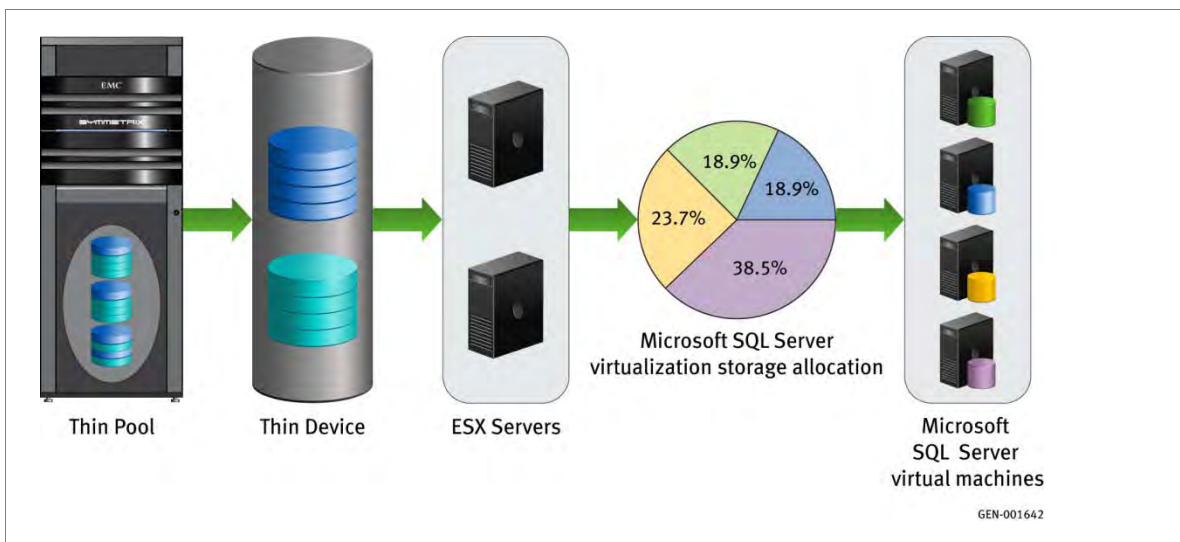


Figure 3. Thin device provisioning and SQL virtual disk allocation

In this solution, FAST VP is leveraged to provide automated tiering for the SQL server workload. This is done by creating a FAST VP policy to specify how much data is on Flash drive and how much is on SATA drives. Performance data is collected at regular intervals and promotes the hot blocks of data up to the Flash tier and demote the less active blocks down to the SATA tier as required. This can greatly reduce both the DBA and storage administrator's resources (time and effort) in maintaining performance in the environment.

FAST VP

Introduction to FAST VP

FAST VP builds on and extends the capabilities of Virtual Provisioning to provide a new form of FAST for Symmetrix at a sub-LUN level. The key capability of FAST VP is its ability to automatically plan and implement nondisruptive changes to storage allocations and to optimize the use of a storage configuration comprising different drive technologies. The goals of the optimizations performed by FAST VP are to take advantage of the low response times of the Flash drives to the greatest possible extent and in general, to provide the most cost-effective use of a mixed drive type configuration.

How FAST VP works

FAST VP in a VMAX environment provides an easy way to employ the storage service specializations of an array configuration with a mixture of drive types. FAST VP offers a simple and cost-effective way to provide optimal performance of a given mixed configuration: by automatically tiering storage to the changing application needs. This places the right data in the right place at the right time.

FAST VP manages the storage allocations of thin devices. The storage used by a thin device managed by FAST VP is normally from a set of thin pools (pools of data devices). The thin device must belong to a storage group that is associated with a FAST VP policy. The thin pools that FAST VP uses are organized into FAST VP tiers, which is a collection of thin pools. The association between a storage group and the FAST VP tiers is specified using a FAST VP policy. A storage group can be associated with only one FAST VP policy (while a FAST VP policy can be associated with more than one storage group). For each FAST VP tier, a maximum configurable allocation percentage is defined in the FAST VP policy.

A FAST VP policy is associated with between one and three FAST VP tiers (but there is no benefit if the configuration is limited to one tier). A given FAST VP Tier may be used by multiple FAST VP Policies.

FAST VP configuration

This section describes the configuration tasks for setting up FAST VP in the test environment.

Step	Action
1	Set FAST VP Performance Time and Move Time windows.
2	Create a virtual pool tier.
3	Create a FAST VP policy.
4	Associate the storage group.

Management tools

The following management options are available to configure FAST VP:

- Solutions Enabler Command Line Interface (SYMCLI)
- SMC

Set Policy time windows

Set two time windows before FAST VP policy works:

- Performance Time Window
- Move Time Window

Performance is sampled in the **Performance Time Window** and the data promotion and demotion are automatically processed in the **Move Time Window** by the FAST VP policy as shown in Figure 4 and Figure 5.

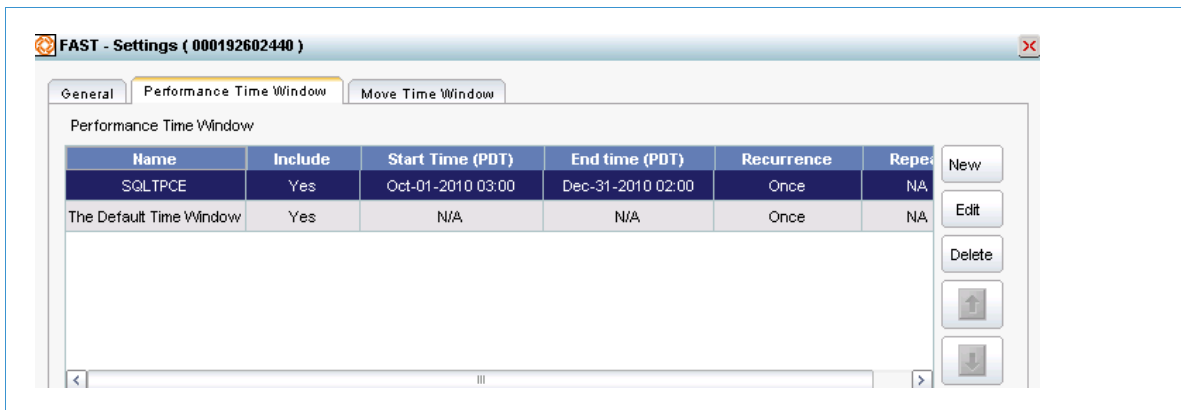


Figure 4. Performance Time Window

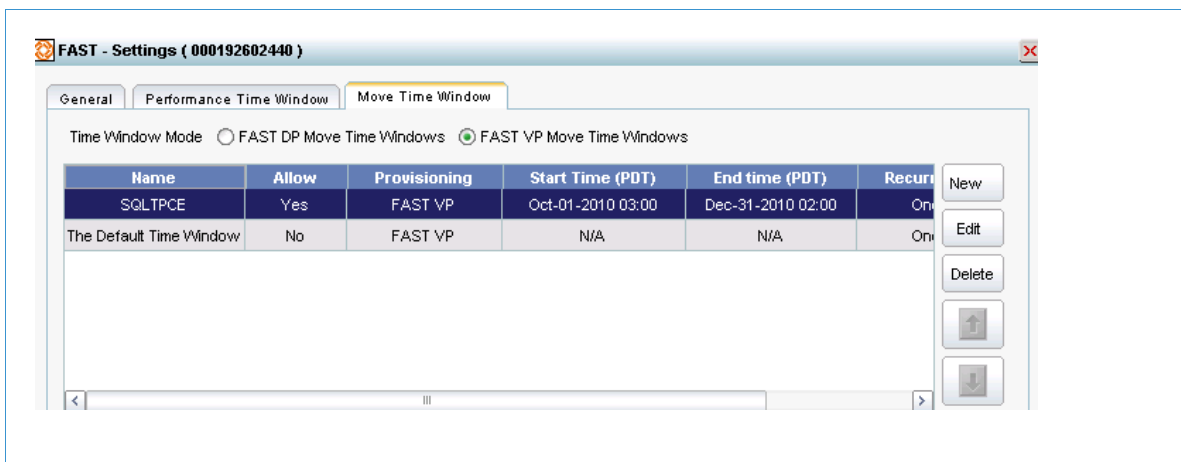


Figure 5. Move Time Window

Create a virtual pool tier

The thin pools to be managed by FAST VP need to be organized as tiers. A thin tier is a set of thin pools. To add a thin pool to a tier, right-click on thin tiers under FAST, select **Create Tier** and then add the required pools to the tier. The **Create Dialog** window is displayed as shown in Figure 6.

The screenshot shows the 'Tier - Create Dialog' window with the following configuration:

- Tier Type: Virtual Pool
- Tier Name: SATA
- Technology: SATA
- Protection: RAID-1

Capacity Numbers for new Tier:

- Symmetrix Capacity: 547200 GB
- Logical Enabled Capacity: 215000GB
- Logical Free Capacity: 15763GB
- Logical Used Capacity: 5737GB

Selection	Virtual Pool Name	In Tier	Emulation	Config	Technology	Total (GB)	
<input checked="" type="checkbox"/>	SQL_SATA	Yes	FBA	RAID-1	SATA	12500	12
<input checked="" type="checkbox"/>	DREM1	No	FBA	RAID-1	SATA	9000	90
<input type="checkbox"/>	DREM2	No	FBA	RAID-1	SATA	9000	90
<input type="checkbox"/>	DREM3	No	FBA	RAID-1	SATA	9000	90
<input type="checkbox"/>	DREM4	No	FBA	RAID-1	SATA	9000	90

Buttons: Select All, Deselect All, OK, Cancel, Help

Figure 6. Create a virtual pool tier

Create a policy

A FAST VP policy specifies a policy of thin tiers. Policies of this type only affect the thin devices in an associated storage group. Each policy can have up to three tiers. The administrator determines the percentage of the storage group's data that resides on a particular tier. This percentage, with the range from 1 percent to 100 percent, is the maximum amount (%) of associated storage group's logical capacity that the FAST controller can allocate to the tier. The total capacities for a policy must be greater than or equal to 100. The **Create FAST Policy** window is shown in Figure 7.

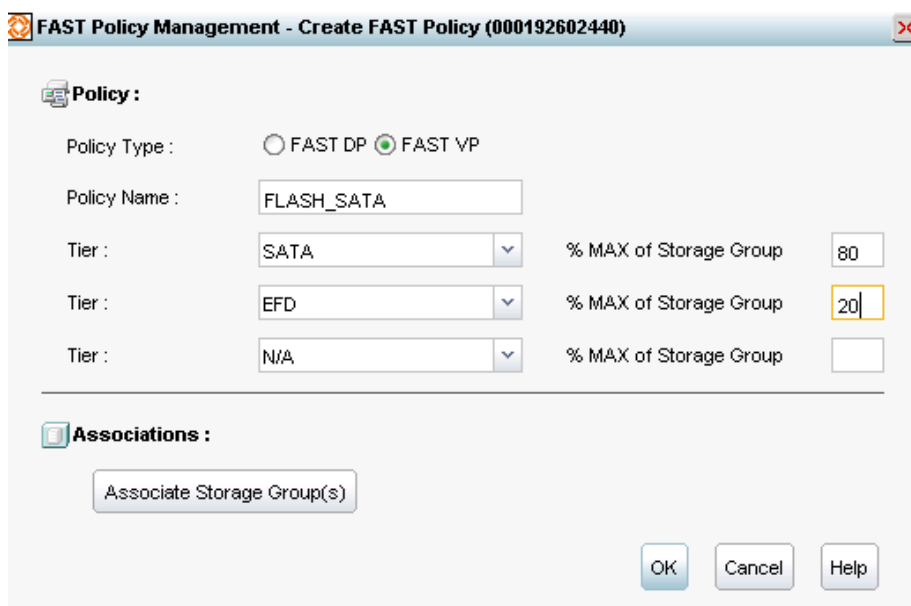


Figure 7. Create FAST Policy

Associate a storage group

Users can associate the storage group containing the thin LUNs to the FAST VP policy. Consider the following items before associating a storage group with a policy.

- The storage group is not already associated with another policy.
- The storage group does not contain a device that is part of another storage group already associated with another policy.
- The policy has at least one tier.
- The storage group only contains meta heads; meta members are not allowed.
- The Symmetrix array has less than the maximum number of allowed associations (8,192).

Users can also adjust the priority of the association priority. Storage groups associated with a policy are assigned a priority value that determines the order in which the tier serves them during conflicts. Possible values range from 1 (the highest) to 3 (the lowest). The **Storage Group Association** window is shown in Figure 8.

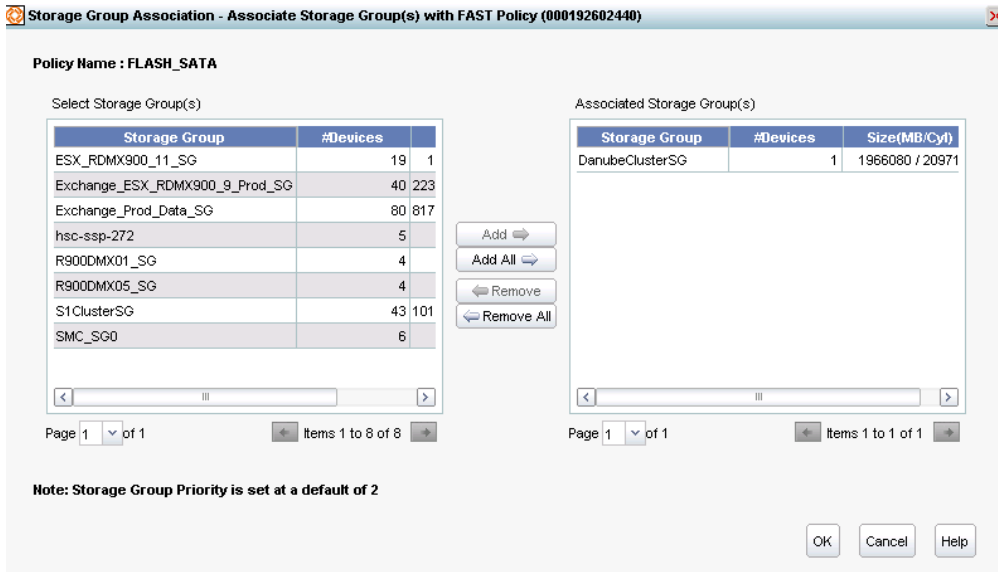


Figure 8. Associate a storage group

FAST VP test and validation

FAST VP test introduction

A performance baseline test was performed for an OLTP workload (TPC-E-like environment) using 64 FC disks. The same data was moved to a FAST VP two-tiered environment with four Flash and 28 SATA drives. The goal was to maintain or improve the SQL server performance and achieve cost savings by using fewer disks, and consume less energy by saving power and cooling. The ease of management for the FAST VP environment was highlighted during the testing process:

- Four SQL servers were configured with all data on the 64 FC disks as baseline.
- The same data was moved to the two-tiered (four Flash drives and 28 SATA drives) storage pool for FAST VP testing.
- The environment is designed to support average disk latency under 20 ms on each virtual machine.
- The environment was designed to process more than 1,000 TPS.
- The environment was designed to support a workload of about 10,000 IOPS.

Test application I/O characteristics

Figure 9 shows the OLTP application workload view. Most portions of the IOPS are Random Read Hit and Random Read Miss. A Random Read Hit is a non-sequential read to the Symmetrix system where it finds the requested data already in cache. Therefore, RRHs run at memory speeds, not disk speeds. A Random Read Miss is a non-sequential read to the Symmetrix system when requested data is not in cache and must be retrieved from physical disks. Based on the observation, the Read Miss and Read Hit ratio were very close, which indicated that the I/O characteristics of the TPC-E-like application were fairly random in the workload.




Figure 9. I/O characteristics of the OLTP application

Baseline performance statistics before FAST VP

The OLTP application was originally configured with all data for the four databases on the FC thin pool. The baseline performance test was performed when all OLTP databases were residing on the FC thin pool. A performance baseline was defined to represent the OLTP environment before applying the FAST VP policies as follows:

- This configuration represented the database performance characteristics before enabling the FAST VP.
- Running the simulated user load with this configuration showed that the ESX server had no CPU or memory constraints, and the client application emulated the varying workload.
- The system processed 1,504 TPS.
- The drives supported a workload of 10,781 IOPS.

Note For the detailed baseline performance statistics, see the *Additional information* section of this white paper.

Migrating into multi-tier thin pool

The data on FC was successfully moved to the two-tier storage pool (with Flash and SATA drives) by specifying the FAST VP policy in Figure 10.

FAST Policy Management - Modify FAST Policy (000192602440)

Policy :

Policy Type : **FAST VP**

Policy Name : **FLASH_SATA**

Tier : EFD % MAX of Storage Group 20

Tier : SATA % MAX of Storage Group 80

Tier : N/A % MAX of Storage Group

Associations :

Associate Storage Group(s)

OK Cancel Help

Figure 10. Modify FAST Policy

Implementing FAST VP on a SATA and Flash multi-tier thin pool

The application workload skew was observed to be 80/20. By implementing a utilization policy of 20% to Flash drives and 80% to SATA drives in the thin pool, this solution successfully used the FAST VP technology to identify and address the hot spots for the four SQL server database workloads.

Storage tiering

Figure 11 shows the IOPS of the SQL virtual machines (host IOPS) when FAST VP automated the data movement to Flash and SATA thin tiered storage from FC. The hot data was moved to the Flash thin tier while the cold data was moved to the SATA thin tier. The stabilization duration was around 3.5 hours.

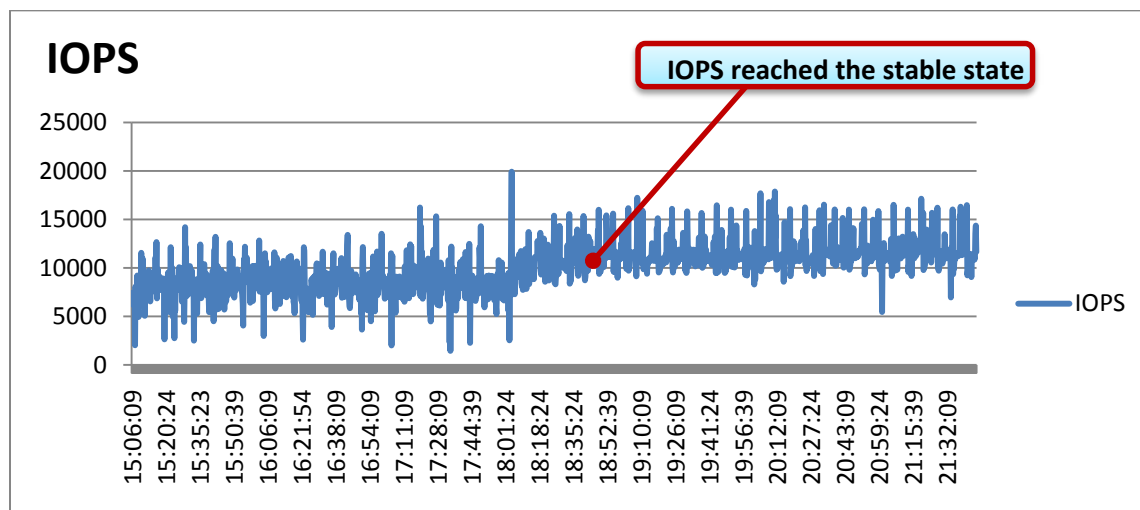


Figure 11. IOPS of the SQL server workload's virtual machine

Performance statistics after FAST VP data relocation on the Flash and SATA tiered storage

Performance improvement of more than 5 percent was observed on tiered storage after migrating the data from the FC thin pool into Flash and SATA tiered storage pools:

- The OLTP application resided on a combination of Flash and SATA thin pools.
- Running the same simulated user load as before, the ESX server CPU utilization and memory remained similar to before.
- The system processed with 1,582 TPS.
- Disks supported a workload of 11,613 IOPS.

ESX server performance statistics on tiered storage with FAST VP enabled

Figure 12 displays the ESX server CPU utilization with FAST VP enabled. The database checkpoint event every 7 minutes on average of the aggregated workload of the four instances caused the CPU utilization dips.




Figure 12. ESX server CPU utilization with FAST VP enabled

Figure 13 and Figure 14 display the thin device IOPS on the Flash and SATA tiered storage.

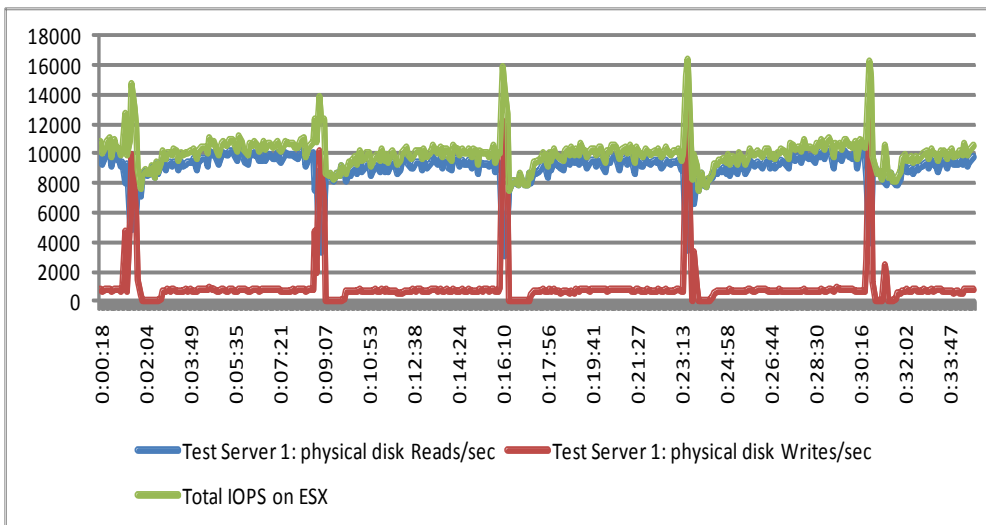


Figure 13. IOPS of the tiered storage on ESX Server 1




Figure 14. IOPS of the tiered storage on ESX Server 2

Table 11 shows the ESX server performance when the workload is on the Flash and SATA thin pool device. Comparing with the baseline performance, with more IOPS being handled by the tiered storage, CPU utilization of each ESX server increases slightly (ESX Server 1 is from 46 percent to 52 percent on average, and ESX Server 2 is from 9 percent to 10 percent on average) and IOPS (total disk Read/sec and Write/sec) and TPS increase too (see the *Performance comparison* section for details).

Table 11. ESX server performance on the thin pool device

	Physical CPU utilization			Disk Read/sec			Disk Write/sec		
	Min	Max	Avg	Min	Max	Avg	Min	Max	Avg
VMware ESX server									
ESX Server 1	28.38	59.57	51.88	3020.86	10394.1	9061.5	12.97	12849.05	1014.7
ESX Server 2	7.49	16.51	10.10	500.87	1730.16	1210.2	0.57	6819.68	130.60

SQL server performance statistics on tiered storage with FAST VP enabled

Table 12 shows the SQL Server performance when the workload is on the Flash drive and SATA thin pool device.

With more IOPS being handled by the tiered storage, CPU utilization of the four SQL virtual machines increased accordingly, and the Read and Write latency improved dramatically for the most active database (SQL01) and moderately for the less active databases. The write latency was improved for the least active database, only the Read latency for the least active database showed a slight increase, because the majority of that database resided on the SATA drive.

Table 12. SQL server performance on the thin pool device

SQL server results	SQL1	SQL2	SQL3	SQL4	Total
TPS	1395	55	106	26	1582
CPU utilization (%)	81.9	3.68	5.64	2.32	N/A
Disk results	SQL 1	SQL 2	SQL 3	SQL 4	Total
Tempdb and log device IOPS	192	13	24	6	235
IOPS	9575	463	961	379	11378
Average Reads/sec	8642	410	867	352	N/A
Average Writes/sec	933	53	94	27	N/A
Read latency (ms)	2	6	9	16	N/A
Write latency (ms)	2	2	3	1	N/A
Average latency (ms)	2	6	9	15	N/A
Tempdb and log device latency (ms)	1	1	1	1	N/A

Performance comparison

Table 13 and Figure 15 show the performance and capacity comparison on FC, and SATA and Flash drive tiered storage.

Table 13. Performance comparison

	IOPS	TPS	Formatted capacity
64 FC	10559	1504	14400 GB
4 Flash drives and 28 SATA	11378	1582	14936 GB
Change percentage	+8%	+5%	+4%




Figure 15. Performance and capacity comparison

Cost comparison

This solution used a much smaller number of drives with the FAST VP multi-tier configuration (Flash drive and SATA drive) when compared to the standard single-tiered, FC-based environment. The FAST VP environment provides notably lower TCO while improving the performance of SQL server.

The cost of disk drives and its power savings are calculated in Table 14 and Figure 16.

Table 14. Disk drives saving percentage

Costs	Saving percentage
Disk drives	32 percent
Disk cost /GB	34 percent
Disk cost/IOPS	32 percent
Annual power consumption	60 percent
Total drive cost over 5 years	35 percent

Note Table 14 represents relative data, based on the list prices of all hardware components involved.

The annual power consumption is calculated using the EMC Power Calculator at <http://powercalculator.emc.com>. This calculation is only limited to the power that is consumed on running the amount of disks in this test.




Figure 16. Cost saving percentage

VAAI validation

VAAI test introduction

VMware vStorage APIs for Array Integration (VAAI) are a new feature of vSphere 4.1. VAAI can offload I/O-intensive tasks such as initializing and copying large files, or a set of files, to the array. VAAI also provides a way to accelerate frequent virtual machine metadata updates that occur when deploying virtual desktop environments at scale. These features can result in a faster deployment of virtual machine disk file (.VMDK file) initialization, virtual machine clones, and virtual desktop environments.

VAAI enabled SUB-LUN cloning

In this solution, both cloning of a virtual machine and storage vMotion were tested. Table 15 demonstrates how VAAI can reduce the duration involved with the process of virtual machine cloning and virtual machine data migration using VMware Storage vMotion.

Table 15. Processing time comparison with VAAI disabled and enabled

Process	Virtual machine capacity	Processing time with VAAI disabled (minutes)	Processing time with VAAI enabled (minutes)
Virtual machine cloning	75 GB	5:49	5:09
	110 GB	8:46	5:57
Storage vMotion	75 GB	3:25	0:57
	110 GB	9:38	3:10

VAAI test conclusion

VMware VAAI offloaded the traffic from LANs to within the array, eliminated the host-based traffic when cloning or performing Storage vMotion.

VSI integration validation

VSI introduction

VSI for vSphere Client version 4.0 provides the Storage Viewer functionality. It extends the vSphere Client to facilitate the discovery and identification of EMC Symmetrix, CLARiiON®, and Celerra® storage devices that are allocated to VMware ESX/ESXi hosts and virtual machines.

VSI validation

This solution validated the functionality of VSI 4.0 as shown in Figure 17 and Figure 18. With the integration of vCenter 4.1, it can help to easily discover Symmetrix Array information, and identify device information, such as the storage information, device name, ID, type, RAID type, meta thin device info, and so on. Also, combining with the PowerPath/VE remote tool installed in the vCenter it can display the path info such as path number and policy.

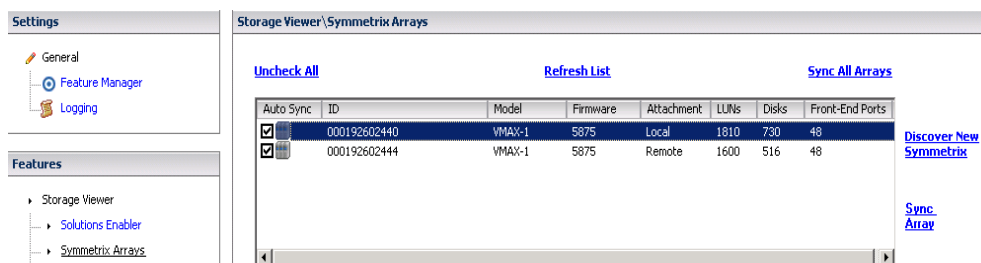


Figure 17. VSI Storage Viewer: Array information

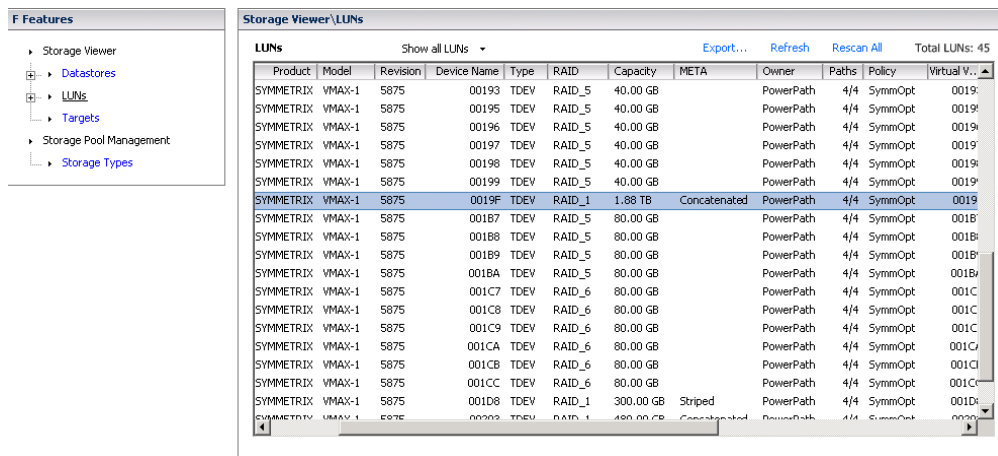


Figure 18. VSI Storage Viewer: LUN information

VSI test conclusion

VSI provided the Storage Viewer functionality to view data store, LUN, and targets information in vCenter.

Pool management via SMC

Pool LUN expansion and shrinkage management via SMC

This solution demonstrates an easy storage management using SMC, thin LUN pool expansion, and shrinking functionality in particular. Also the solution demonstrates how to set alert watermarks to monitor the pool utilization.

Shrink thin pool

As shown in Figure 19 and Figure 20, the data device can be disabled and then drained from the thin pool in the SMC through Device Pool Management > DATA Device Management > Deactivate Device. The drained data is reallocated to the other active data devices with a Round-Robin write-balance algorithm. The disabled data device can be deleted safely afterwards. Extra devices can be easily removed from a pool to ensure there is no wasted space.

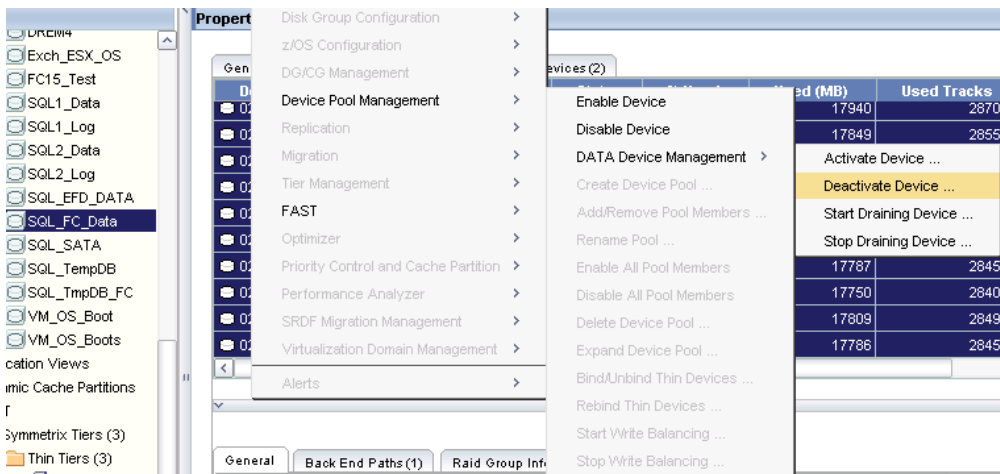


Figure 19. Shrink thin pool: Deactivate Device

Properties - 000192602440 Local/SQL_FC_Data

Bound Thin Devices (2)								
Dev	Emulation	Dev Config	State	% Used	Used (MB)	Used Tracks	Free (MB)	
<input type="radio"/> 029C	FBA	2-Way Mir	Enabled	40	20647	330348	30554	
<input type="radio"/> 029D	FBA	2-Way Mir	Enabled	40	20681	330900	30519	
<input type="radio"/> 029E	FBA	2-Way Mir	Disabled	0	0	0	51200	
<input type="radio"/> 029F	FBA	2-Way Mir	Disabled	0	0	0	51200	
<input type="radio"/> 02A0	FBA	2-Way Mir	Disabled	0	0	0	51200	
<input type="radio"/> 02A1	FBA	2-Way Mir	Disabled	0	0	0	51200	
<input type="radio"/> 02A2	FBA	2-Way Mir	Disabled	0	0	0	51200	
<input type="radio"/> 02A3	FBA	2-Way Mir	Draining	16	8110	129756	43091	
<input type="radio"/> 02A4	FBA	2-Way Mir	Draining	1	557	8904	50644	
<input type="radio"/> 02A5	FBA	2-Way Mir	Draining	1	677	10824	50524	

Figure 20. Shrink thin pool: Data device state

Expand thin pool

As shown in Figure 21 and Figure 22 the thin pool can be expanded by percentage or actual size in GB through Device Pool Management > Expand Device Pool in the SMC. Users can choose to start write balancing automatically or they can manually perform the operation afterwards.

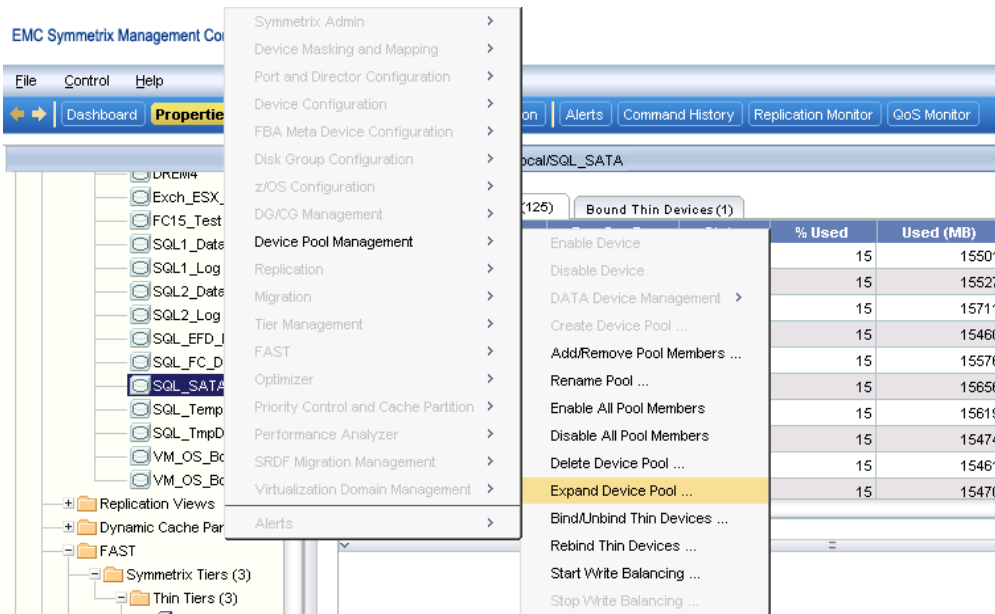


Figure 21. Expand device pool

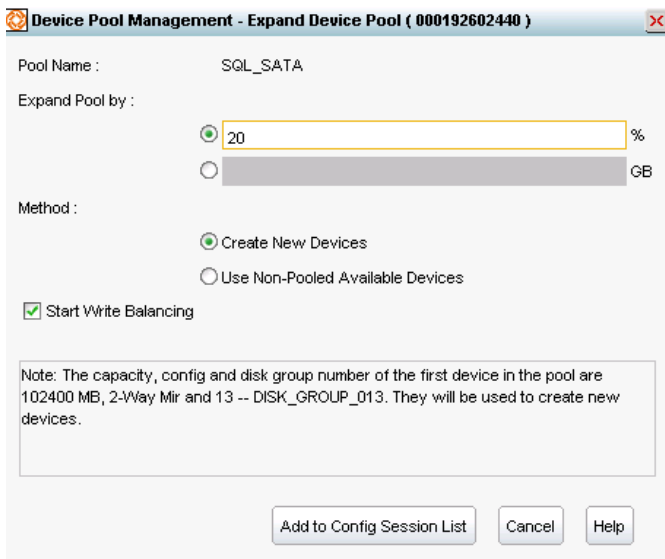


Figure 22. Expand thin pool: Capacity setting

Thin pool monitor - Set pool utilization threshold via SMC

As shown in Figure 23 and Figure 24, the thin pool utilization can be monitored by setting a pool utilization threshold through **Tasks > Alert Configure** in the SMC. Users can disable specific pool monitoring, and set the WARNING, CRITICAL and FATAL watermarks of the thin pool utilization. When the thin pool utilization reaches or exceeds the preset threshold the alert will be triggered and display in the **Alert** window.

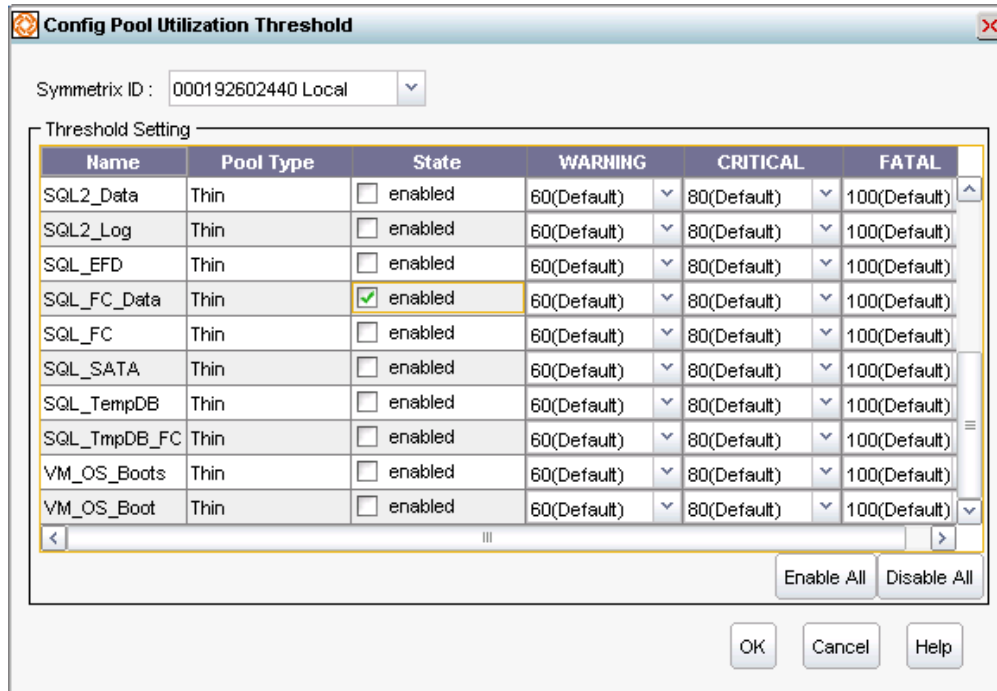


Figure 23. Set pool utilization threshold




Figure 24. The WARNING alert is triggered

Pool LUN management via SMC

Pool expansion and shrinkage through the SMC were easy to operate and can be done dynamically while the application is running with minimal impact on the application. Also users can set the alerts of different levels to monitor the pool utilization.

Conclusion

Summary

This solution compares two architectures for hosting high-performance SQL databases in a virtualized environment. The results demonstrate that a FAST VP implemented tiered storage configuration with Flash and SATA drives can improve performance while significantly reducing storage-related costs.

Findings

This white paper concludes the following findings:

- Moving from a single-tiered, static storage configuration to a multi-tiered, dynamic FAST VP configuration can greatly reduce TCO while improving performance:
 - Three areas of cost and time savings:
 - Acquisition costs: 32 percent reduction than the costs in an all-FC configuration
 - Power costs: 60 percent reduction than the costs in an all-FC configuration
 - Management time: greatly reduced
 - Two areas of SQL server performance improvement:
 - SQL storage system throughput measured in IOPS increased by 8 percent
 - SQL transaction throughput measured in TPS increased by 5 percent

Easy to configure and mostly automated, FAST VP with its multi-tier storage is proven to be a better solution to support the tested OLTP environment than the typical FC-based configuration.

- VMware VAAI moves the LAN operation to the array, and shows a three-fold performance improvement in virtual machine cloning and Storage vMotion.
- EMC VSI provides the Storage Viewer functionality to view data store, LUN, and targets information in vCenter.
- Storage pools expansion and shrinkage through the SMC are easy to operate and can be done dynamically while the application is running with minimal impact on the application.

References

White papers

For white papers that provide similar solutions, see the following:

- *New Features in EMC Engenuity 5875 for Open Systems Environments*
- *EMC Tiered Storage for Microsoft SQL Server 2008 - Enabled by EMC Unified Storage and EMC Fully Automated Storage Tiering (FAST) - An Architectural Overview*
- *EMC Tiered Storage for Microsoft SQL Server 2008 - Enabled by EMC Symmetrix V-Max with FAST - A Detailed Review*

Other documentation

For additional information on related topics, see the following:

- *EMC Symmetrix Management Console (SMC) online help (integrated with the Symmetrix VMAX)*

Additional information

Overview

This section provides the detailed baseline performance statistics.

Baseline ESX server performance statistics

Figure 25 displays the two ESX server CPU utilizations during the baseline performance test.




Figure 25. ESX server CPU utilization of the baseline test

Figure 26 and Figure 27 display the thin device IOPS during the baseline performance test.

Each workload triggered a database checkpoint event every 7 minutes on average. Therefore, there were one or more data-flushing operations on the back-end devices, which caused a spike in write I/O activity every 7 minutes.

Because the four SQL servers used the same data store from the same thin device, the device IOPS summation of the two ESX servers was the same as the one of the four SQL server virtual machines.




Figure 26. IOPS of the FC thin device on ESX Server 1




Figure 27. IOPS of the FC thin device on ESX Server 2

Table 16 shows the ESX server performance when the workload is on the FC-only thin pool device. The baseline performance of the ESX servers indicates that the CPU utilization is normal and there is no obvious performance bottleneck, and the varying workloads of the four SQL servers hosted by two ESX servers contribute to the total IOPS of the thin device.

Table 16. ESX Server performance

	Physical CPU UT (%)			Disk Reads/sec			Disk Writes/sec		
	Min	Max	Avg	Min	Max	Avg	Min	Max	Avg
VMware ESX Server									
ESX Server 1	20.49	56.30	45.53	2490.23	9626.20	7642.90	26.70	7793.24	886.00
ESX Server 2	7.58	20.55	9.42	763.35	1858.52	1293.70	0.16	6462.48	123.34

Baseline SQL server performance statistics

Table 17 shows the SQL server performance when the workload is on the FC-only thin pool device.

Table 17. SQL server performance

Performance value	SQL1	SQL2	SQL3	SQL4	Total
TPS	1312	55	109	28	1504
CPU UT (%)	73	3.1	4.9	1.5	N/A
Disk results	SQL1	SQL2	SQL3	SQL4	Total
Tempdb and log device IOPS	177	13	25	7	222
IOPS	8659	452	1017	431	10559
Average Reads/sec	7797	398	913	401	N/A
Average Writes/sec	862	54	104	30	N/A
Read latency (ms)	4	7	6	7	N/A
Write latency (ms)	4	3	3	2	N/A
Average latency (ms)	4	7	6	7	N/A
Tempdb and log device write latency(ms)	1	1	1	1	N/A