

STORAGE TIERING FOR MICROSOFT SQL SERVER AND EMC SYMMETRIX VMAX WITH ENGINUITY 5875

Abstract

This white paper examines the application of managing Microsoft SQL Server in a storage environment utilizing multiple storage configurations to optimize performance cost-effectiveness of storage allocations. Practical examples are provided of automated data movement between storage tiers to meet performance requirements for SQL Server databases and cost-effectiveness of EMC® Symmetrix VMAX™ resources.

January 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.

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

Part Number h8125

Table of Contents

Executive summary	4
Audience.....	5
Product and features overview	6
Symmetrix VMAX and the Enginuity Operating Environment	6
Enterprise Flash Drives	7
Enhanced Virtual LUN (VLUN) migrations.....	8
Symmetrix Fully Automated Storage Tiering.....	9
Symmetrix Fully Automated Storage Tiering with Virtual Pools (FAST VP)	10
How FAST is configured	11
Symmetrix Management Console	12
Skewed data access.....	13
Storage tiering mechanics	14
Working with automated storage tiering and SQL Server databases	14
Fully Automated Storage Tiering (FAST VP).....	14
Test methodology and configuration	14
Configuring FAST VP for the environment.....	18
Monitoring workload performance.....	22
Identification of database workloads.....	23
Effect of FAST VP.....	24
Assessing the benefits of FAST VP	28
Conclusion	30

Executive summary

The EMC® Symmetrix VMAX™ series with Enginuity™ incorporates a scalable fabric interconnect design that allows the storage array to seamlessly grow from an entry-level configuration to a 2 PB system. Symmetrix® VMAX provides predictable, self-optimizing performance and enables organizations to scale out on demand in Private Cloud environments. It automates storage operations to exceed business requirements in virtualized environments, with management tools that integrate with virtualized servers and reduce administration time in private cloud infrastructures. Customers are able to achieve “always on” availability with maximum security, fully nondisruptive operations and multisite migration, recovery, and restart to prevent application downtime. EMC Symmetrix VMAX delivers enhanced capability and flexibility for deploying Microsoft SQL Server databases throughout the entire range of business applications, from mission-critical applications to test and development. In order to support this wide range of performance and reliability at minimum cost, Symmetrix VMAX arrays support multiple drive technologies including Enterprise Flash Drives (EFD), Fibre Channel drives (FC), and SATA drives. In addition, various RAID protection mechanisms are allowed that affect the performance, availability, and economic impact of a given Microsoft SQL Server system on a Symmetrix VMAX storage array.

Organizations are under increasing pressure to improve performance as well as the return on investment (ROI) for storage technologies, and ensuring that data is placed on the most appropriate storage tier is a key requirement. It is common to find applications utilizing the services of the SQL Server environment, having differing access patterns to different files, directing a significant workload to a relatively small subset of the data stored within the database. Additionally, ranges of data within a given data file may become more heavily accessed as a result of query execution, or specification of query range parameters. These phenomena are referred to as data access skewing. Optimizing the placement of these various components of the database on the storage tier that matches their needs best will help increase performance, reduce the overall number of drives, and improve the total cost of ownership (TCO) and ROI for storage deployments.

As organizations increase deployment of multiple drive types and protection types in their high-end storage arrays, storage and database administrators now have the challenge of selecting the correct storage configuration for each application. Often, a single storage tier is selected for all the data in a given database, effectively placing both active and idle data portions on fast FC drives. This approach can be expensive and inefficient, because infrequently accessed data will reside unnecessarily on high-performance drives.

Alternatively, making use of high-density, low-cost SATA drives for the less active data, FC drives for the medium active, and EFD for the very active enables efficient use of storage resources, and reduces the overall cost and number of drives necessary. This, in turn, also helps to reduce energy requirements and floor space, allowing the business to grow more rapidly.

To achieve this “tiered” approach with Microsoft SQL Server databases, administrators had been able to use Symmetrix Enhanced Virtual LUN Technology to move full logical devices between drive types and RAID protections seamlessly inside the storage array.

Symmetrix Virtual LUN technology does not require application downtime. It preserves the device IDs, which means there is no need to change file system mount points, volume manager settings, database file locations, or scripts. It also maintains any TimeFinder® or SRDF® business continuity operations even as the data migration takes place.

The manual and often time-consuming approach to full volume storage tiering was automated using Fully Automated Storage Tiering or FAST. FAST uses policies to manage sets of logical devices and available storage tiers. Based on the policy guidance and the actual workload profile over time, the FAST controller will recommend and even execute automatically the movement of the managed devices between the storage tiers.

Extending the value of FAST technology, Symmetrix VMAX with Enginuity 5875 now introduces Fully Automated Storage Tiering for Virtual Pools (FAST VP). The enhanced functionality allows for the definition of storage tiers within a Symmetrix VMAX array to service the storage allocations for a set of thin devices. Thin device extents are now migrated between the defined tiers based on a user-defined FAST policy, as compared to the full non-thin devices.

This white paper describes a tiered storage architecture approach for SQL Server databases, and the way in which devices can be moved nondisruptively, using either Enhanced Virtual LUN or FAST VP technologies, in order to allow FAST VP to dynamically relocate storage extents to the most appropriate storage tier.

Audience

This white paper is intended for Microsoft SQL Server database administrators, storage administrators and architects, customers, and EMC field personnel who want to understand the implementation of storage tiering in a Symmetrix VMAX and the new functionality made available with Enginuity 5875.

Product and features overview

The following is a brief summary of EMC Symmetrix technologies that must be understood in order to be able to deploy a storage tiering strategy with the Microsoft SQL Server database environment.

Symmetrix VMAX and the Enginuity Operating Environment

Symmetrix VMAX along with its Enginuity operating environment is EMC's solution to organizations' need for high-end storage performance and scalability. With the release of Enginuity 5875, Symmetrix VMAX systems now deliver additional software capabilities that improve capacity utilization, ease of use, business continuity, and security. These features provide significant advantage to customer deployments in a virtualized environment, where ease of management, and protection of virtual machine assets and data assets are required.

Symmetrix VMAX arrays extend the scalability of previous generations of Symmetrix technology, by providing a superior level of scalability, and support for a broad new range of drive technologies as detailed in Figure 1.

Additionally, Symmetrix VMAX offers the ultimate in flexibility, including the ability to incrementally increase back-end performance by adding VMAX Engines and storage bays. Each high-availability VMAX Engine controls eight redundant Fibre Channel loops that support up to either 240 or 360 drives depending upon configuration. Subsequently, each high-availability VMAX Engine provides front-end as well as back-end connectivity, providing enhanced scalability.



- 1 to 8 high-availability VMAX Engines
- Up to 2.1 PB usable capacity
- Up to 128 FC FE ports
- Up to 64 FICON FE ports
- Up to 64 GigE/10 GigE/iSCSI FE ports
- Up to 1 TB global memory
- 48 to 2,400 drives
- Enterprise Flash Drives 200/400 GB
- FC drives 146/300/450/600 GB 15k rpm
- FC drives 300/450/600 GB 10k rpm
- SATA drives 1/2 TB 7.2k rpm

Figure 1. Symmetrix VMAX hardware scalability

Symmetrix VMAX systems also maintain customer expectations for high-end storage in terms of availability. High-end availability is more than just redundancy; it means nondisruptive operations and upgrades, and being “always online.” Beyond previous Symmetrix generations, Symmetrix VMAX arrays provide:

- Nondisruptive expansion of capacity and performance at a lower price point
- Sophisticated migration for multiple storage tiers within the array
- The power to maintain service levels and functionality as consolidation grows
- Simplified control for provisioning in complex environments

Enterprise Flash Drives

With over two years’ worth of support for Flash drives, EMC has created an ultra-high-performance storage tier that removes previous performance limitations imposed by magnetic disk drives.

For years, the most demanding enterprise applications have been limited by the performance of magnetic disk media. Performance in storage arrays has been constrained by the physical limitations of hard disk drives. Enterprise Flash Drives deliver unprecedented performance and response times, which are benefits well suited for demanding SQL Server configurations.

Enterprise Flash Drives dramatically increase performance for latency-sensitive databases like SQL Server. Enterprise Flash Drives, also known as solid state drives

(SSD), contain no moving parts, which removes much of the storage latency delay associated with traditional magnetic disk drives. A Symmetrix VMAX with Enterprise Flash Drives can deliver single-millisecond application response times and up to 30 times more I/O operations per second (IOPS) than traditional Fibre Channel hard disk drives (HDD). Additionally, because there are no mechanical components, Flash drives consume significantly less energy than hard disk drives. When replacing a larger number of HDDs with a lesser number of Enterprise Flash Drives, energy consumption can be reduced by up to 98 percent for a given IOPS workload.

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 SQL Server deployments. The practice of underutilizing a hard disk drive for increased performance is commonly referred to as short-stroking. EFDs can increase SQL Server database performance and eliminate the need to short-stroke drives, thus keeping storage footprint and power consumption to a minimum and reducing TCO.

Enhanced Virtual LUN (VLUN) migrations

The Enhanced Virtual LUN Technology migration feature introduced with Symmetrix VMAX with Enginuity 5874 offered SQL Server DBAs the ability to transparently migrate full, non-thin LUNs used by SQL Server databases between different storage tiers. The storage tiers are defined based on a number of attributes that may include a particular storage device technology such as high-performance EFDs, Fibre Channel drives, or high-capacity, low-cost SATA drives. Migrations also allow for a change in RAID protection, which may also be an attribute of a storage tier.

Symmetrix VMAX with Enginuity 5875 introduces enhanced functionality that allows for migration of Virtual Provisioning™ (VP) extents between different thin pools. This new functionality now tracks extent allocation for a thin device, and provides the facility to migrate all the extents comprising the storage allocations to a user-specified target thin pool. Thin device migrations allow for changes in RAID protection and in technology type as provided by migrations of non-thin devices.

Enhanced Virtual LUN is fully integrated with Symmetrix replication technology and maintains consistency and source/target device relationships in replications such as SRDF, TimeFinder/Clone, TimeFinder/Snap, or Open Replicator. The migration does not require swap or Dynamic Relocation Volumes (DRV) space, and is nondisruptive to the attached SQL Server environment and other internal Symmetrix applications such as TimeFinder and SRDF. All migration combinations of drive types and protection types are valid except for unprotected volumes.

The SQL Server LUNs can be migrated to either unallocated space (also referred to as unconfigured space) or to configured space, which is defined as existing Symmetrix volumes that are not currently assigned to a server. The data on the original source volumes is cleared using instant initialization (iVTOC) once the migration has been completed.

The migration is completely transparent to both SQL Server and the underlying Microsoft Windows Server operating system since the operation is executed against the Symmetrix device; thus all attributes of the LUNs are not changed and SQL Server operations are uninterrupted. Furthermore, in SRDF environments, the migration does not require customers to re-establish their disaster recovery protection after the migration.

Enhanced Virtual LUN migration helps customers to implement an Information Lifecycle Management (ILM) strategy for their databases, such as the movement of the entire or partial database, specific files, or filegroups between storage tiers. It also allows adjustments in service levels and performance requirements to application data. For example, customers often provision storage for a particular application before clear performance requirements are known. LUN migration may be utilized at a later time, once the requirements are better understood, to implement any adjustment to increase user experience and ROI using the correct storage class.

Symmetrix Fully Automated Storage Tiering

Introduced in the Enginuity 5874 Q4 service release and supported with Enginuity 5875, EMC Symmetrix VMAX Fully Automated Storage Tiering (FAST) is Symmetrix software that utilizes intelligent algorithms to continuously analyze device I/O activity and generate plans for moving and swapping full devices, or in the case of FAST VP, storage extents, for the purposes of allocating or re-allocating application data across different storage tiers within a Symmetrix array.

FAST proactively monitors workloads at both the Symmetrix device (LUN) and extent level in order to identify “busy” storage allocations that would benefit from being moved to higher-performing tiers such as EFD. FAST will also identify less “busy” allocations that could be relocated to higher-capacity, more cost-effective storage such as SATA drives without altering performance.

Time windows can be defined to specify when FAST should collect performance statistics (upon which the analysis to determine the appropriate storage tier), and when FAST should perform the configuration changes necessary to move storage allocations between the defined storage tiers. Movement is based on user-defined storage tiers and FAST policies.

The primary benefits of FAST include:

- Automating the process of identifying storage allocations that can benefit from higher-performance tiers (for example, EFD) and/or that can be kept on higher-capacity, less-expensive drives without impacting performance (for example, SATA)
- Improving application performance at the same cost, or providing the same application performance at lower cost. Cost is defined as space, energy, acquisition, management, and operational expense
- Optimizing and prioritizing business applications, allowing customers to dynamically allocate resources within a single array

- Delivering greater flexibility in meeting different price/performance ratios throughout the lifecycle of the information stored

The management and operation of FAST and FAST VP can be conducted using either the Symmetrix Management Console (SMC) or the Solutions Enabler Command Line Interface (SYMCLI). Additionally, detailed performance trending, forecasting, alerts, and resource utilization are provided through Symmetrix Performance Analyzer (SPA). And if so desired, EMC Ionix™ ControlCenter® provides the capability for advanced reporting and analysis that can be used for chargeback and capacity planning.

EMC Symmetrix VMAX FAST and FAST VP automate the identification of data volumes for the purposes of relocating application data across different performance/capacity tiers within an array.

FAST, also known as standard FAST, operates on standard, non-thin, Symmetrix devices. Data movements executed between tiers are performed at the full volume level.

FAST VP operates on Virtual Provisioning thin devices. As such, data movements executed can be performed at the sub-LUN level, and a single thin device may have extents allocated across multiple thin pools within the array.

Because FAST and FAST VP support different device types – standard and thin, respectively – they both can operate simultaneously within a single array. Aside from some shared configuration parameters, the management and operation of each can be considered separately.

The main focus of this document is to discuss the implementation and management of FAST VP in Virtual Provisioning environments. For more information on FAST, in standard provisioning environments, please refer to *Storage Tiering for Microsoft SQL Server and EMC Symmetrix VMAX with Enginuity 5874*. The implementation of FAST and full device migrations detailed in the paper applies for full device FAST in Enginuity 5875.

Symmetrix Fully Automated Storage Tiering with Virtual Pools (FAST VP)

The implementation of FAST VP with Symmetrix VMAX utilizes three main elements. These are:

- Storage tier – A shared resource with common technologies
- FAST policy – A policy that manages data placement and movement across Symmetrix tiers to achieve service levels and for one or more storage groups
- Storage group – A logical grouping of devices for common management

Figure 2 shows the relationship between FAST managed objects.




Figure 2. FAST managed objects

Each of the three managed objects can be created and managed by using either SMC or SYMCLI.

How FAST is configured

FAST is configured by defining three distinct objects:

- *Storage groups.* A storage group is a logical grouping of up to 4,096 Symmetrix devices. Storage groups are shared between FAST and Auto-provisioning Groups; however, a Symmetrix device may only belong to one storage group that is under FAST control.
- *Storage tier.* Storage tiers are a combination of a drive technology (for example, EFD, FC 15k rpm, or SATA) and a RAID protection type (for example, RAID 1, RAID 5 3+1, or RAID 6 6+2). There are two types of storage tiers – standard and thin.

Disk group provisioning tiers

The storage tier type used by FAST is called a DP tier. It is defined by combining one or more physical disk groups, of the same technology type, and a RAID protection type.

Virtual provisioning tiers

For FAST VP, the storage tier is called a VP tier. When defined, VP tiers contain between one and four thin storage pools – each thin pool must contain data devices of the same RAID protection type, and be configured on the same drive technology.

VP tier characteristics

A thin storage pool must contain data devices configured as the same RAID protection on a single drive technology. In the case of Fibre Channel and SATA drives, the rotational speed of the drives must also match. However, thin pools containing data devices configured on rotating drives of different sizes and speeds may be combined in a single VP tier.

A thin pool may only belong to one VP tier. There is no overlapping of pools between tiers allowed.

All VP tiers are considered to be static, meaning that thin pools must be explicitly added to the tier. However, if data devices are added to an existing pool, the additional capacity will automatically be made available within the VP tier.

A Symmetrix VMAX storage array will support up to 256 Symmetrix tiers. Each Symmetrix tier name may contain up to 32 alpha-numeric characters, hyphens (-), and underscores (_). Tier names are not case-sensitive.

- *FAST policies.* A FAST policy groups between one and three tiers and assigns an upper usage limit for each storage tier. The upper limit specifies how much capacity of a storage group associated with the policy can reside on that particular storage tier.

FAST policies may only contain storage tiers of one type – DP or VP. The first tier added to a policy will determine the type of tiers that can subsequently be added.

For policies containing VP tiers, the upper capacity usage limit for each storage tier is specified as a percentage of the allocated capacity of the associated storage group. The limit defines the maximum amount of allocated capacity of the storage group that can exist on a particular VP tier.

The usage limit for each tier must be between 1 percent and 100 percent. When combined, the upper usage limit for all thin storage tiers in the policy must total at least 100 percent, but may be greater than 100 percent.

Creating a policy with a total upper usage limit greater than 100 percent allows flexibility with the configuration of a storage group whereby data may be moved between tiers without necessarily having to relocate a corresponding amount of other data within the same storage group.

Multiple FAST policies may reuse the same tier, allowing different usage limits to be applied to different storage groups for the same tier.

A Symmetrix VMAX storage array will support up to 256 FAST policies. Each FAST policy name may be up to 32 alpha-numeric characters, hyphens (-), and underscores (_). Policy names are not case-sensitive.

FAST VP can only be configured in an Automatic mode in which the system will continually gather statistics, analyze, and execute moves and swaps to maintain optimal configuration based on policy, or in a “user approval” mode (User Approved) in which all configuration change plans made by FAST must be approved for a FAST suggested plan to be executed.

Symmetrix Management Console

Many large enterprise data centers have stringent change control processes that ensure reliable execution of any modification to their IT infrastructure. Often changes are implemented using scripts that are fully documented, have been thoroughly reviewed, and can be consistently executed by all storage administrators. An alternative to using Solutions Enabler scripting is to use the Symmetrix Management Console (SMC).

SMC contains a number of wizards to simplify the configuration of multi-step tasks. In subsequent sections, individual steps are executed to display SMC functionality; however, SMC provides a wizard to initially configure Symmetrix FAST VP environments as shown in Figure 3, and is the recommended method to appropriately configure initial deployments.

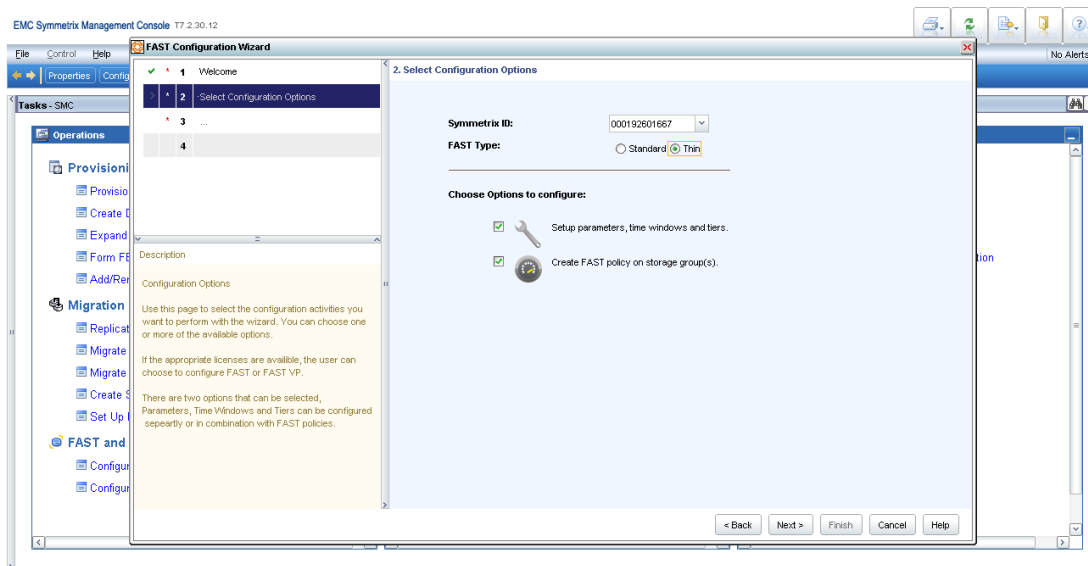


Figure 3. FAST Configuration Wizard

Skewed data access

Skewed data access is a phenomenon that is present in most databases. Skewing is when a small number of SQL Server LUNs, or the data files resident on those LUNs, receive a large percentage of the I/O from the applications. This skewed type of activity can be transient, that is to say lasting only a short period of time, or persistent and lasting much longer periods of time. Skewed, persistent access to storage allocations makes those allocations good candidates to place on a higher-performing storage tier, if they are not already on the highest-performing tier.

Skewing by its nature means that there are SQL Server data structures that are receiving lower activity than others. If the storage allocations servicing those structures are persistently receiving fewer requests than others, then they may be good candidates to move a lower-performing storage tier.

Microsoft SQL Server database environments are created by defining one or more filegroups (the PRIMARY filegroup will always exist). Each filegroup is subsequently defined to exist on one or more data files located on NTFS volumes. Due to the proportional fill mechanism used by SQL Server, data files within a given filegroup will generate almost identical I/O patterns. This highly correlated workload also leads to contention, and is the primary motivation for the best practice recommendation to distribute these files across volumes and LUNs. It is a commonplace occurrence, however, to have storage allocations for these multiple LUNs be provided from within the same storage pool. Thus, while the workload is distributed across different LUNs

and NTFS volumes at the host level, the workload is generated to the same physical spindles supporting the thin pool.

Storage tiering mechanics

The goal of an effective storage tiering approach in a multi-tiered storage configuration is to place the right data on the right storage at the right time. Access patterns to data and index pages are an important consideration with respect to the efficiencies of FAST VP. Persistent access to a given data range will result in those storage allocations used by that data range to be moved to a higher-performing tier. Such mechanisms can be assisted by defining SQL Server table partitions, where data files may be co-resident within a single thin LUN, but where high access may be limited to a lesser number of data files.

A given range of data within a SQL Server data file may be highly active and highly accessed when data is created on it in the first instance. But over time, its usage may drop to where it could be deployed on a storage tier that has lower-cost and lower-performance characteristics. Alternatively, co-located storage devices within a given storage tier may be contending for I/O resources. One workload may represent a much more demanding I/O profile. Identification of this demand, and mapping this to available storage tiers within the array, may not only improve the performance of the higher workload, but in so doing, may provide benefits to the remaining portions of the database as the available resources are able to satisfy the reduced I/O demand.

FAST VP operates at a sub-LUN level, and can dynamically migrate those portions of data files to be placed on the most appropriate storage tier based on their utilization level. This continual, dynamic process can readily adapt the storage allocations based on the changing nature of the data access patterns.

Working with automated storage tiering and SQL Server databases

To validate the mechanisms and show the efficacy of the FAST VP solution for Microsoft SQL Server environments a test environment was defined to simulate a user environment. A moderate to high user workload was generated against a SQL Server 2008 R2 environment. Storage allocations were initially provided in a manner that is consistent with current provisioning practices deployed within the SQL Server user community. Storage tiers were defined within the environment, and a FAST policy was defined for the SQL Server database.

Fully Automated Storage Tiering (FAST VP)

The following section discusses implementation of FAST VP with thin devices.

Test methodology and configuration

For all testing, Microsoft Windows Server 2008 R2 and Microsoft SQL Server 2008 R2 were utilized. The primary user database was implemented to be comprised of several

filegroups where each filegroup mapped to one or more files located over eight thin LUNs, as shown in Figure 4.

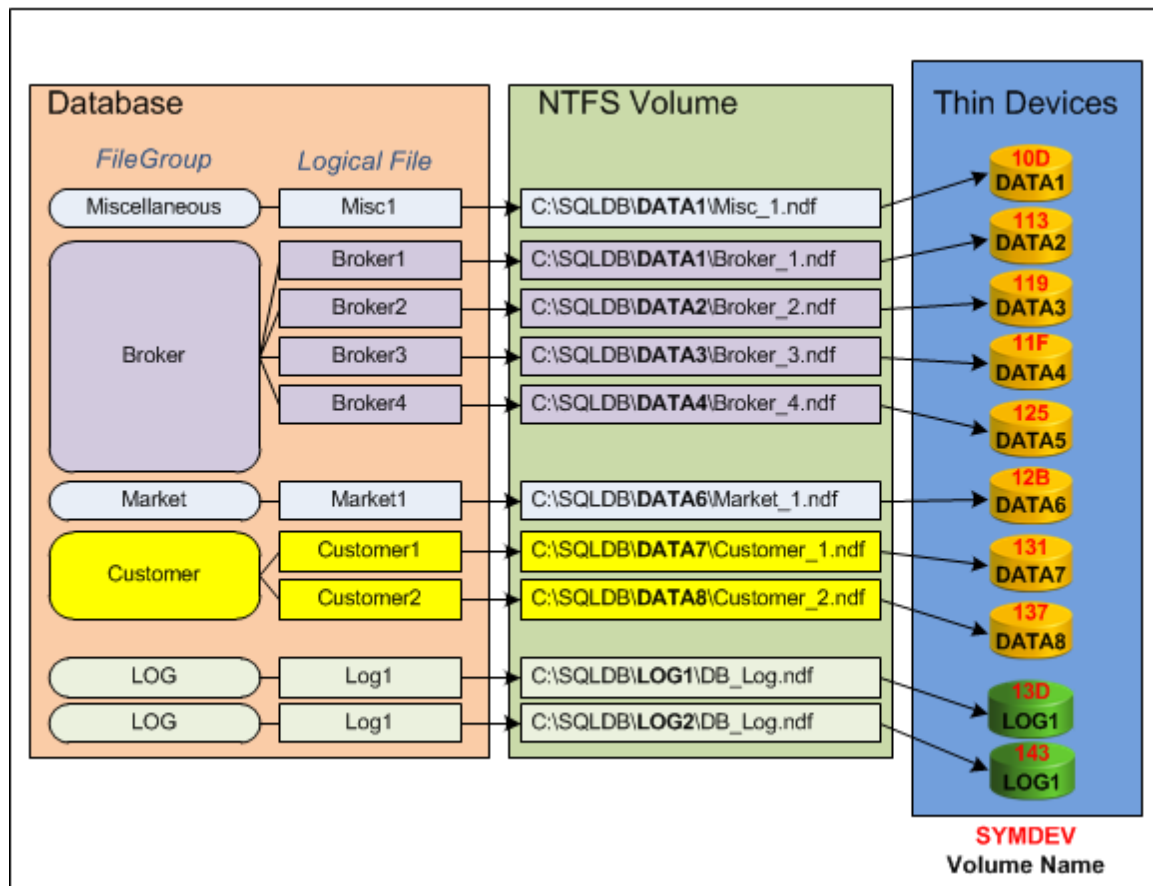


Figure 4. Overview of relationships of filegroups, files, and physical drives

Using the information presented in Figure 4, it is possible to map the locations of the SQL Server logical structures, such as tables and indexes, to physical locations. Given the proportional fill algorithm used by SQL Server, each logical entity, such as the Broker filegroup, will have allocations from each of the constituent files. Subsequently, all I/O activity generated against tables and indexes defined within the Broker filegroup will be serviced by the physical files. As a result this I/O will be directed at the LUNs that provide the storage for the files and then finally be serviced by the data devices configured in the thin pool. In addition to the displayed data devices, two thin LUNs were configured for TEMPDB usage.

All LUNs were initially bound to a single thin pool, named SQL_FC, that was configured with 48 data devices in a RAID 5 3+1 configuration. The 48 data devices were 60 GB each, and were provisioned from a disk group of 96 physical 450 GB 15k rpm drives. This resulted in at least two data devices on each physical disk drive for the thin pool. Figure 5 displays the relationship of thin LUNs to data devices and to the physical disk drives ultimately used for the initial Fibre Channel disk pool.

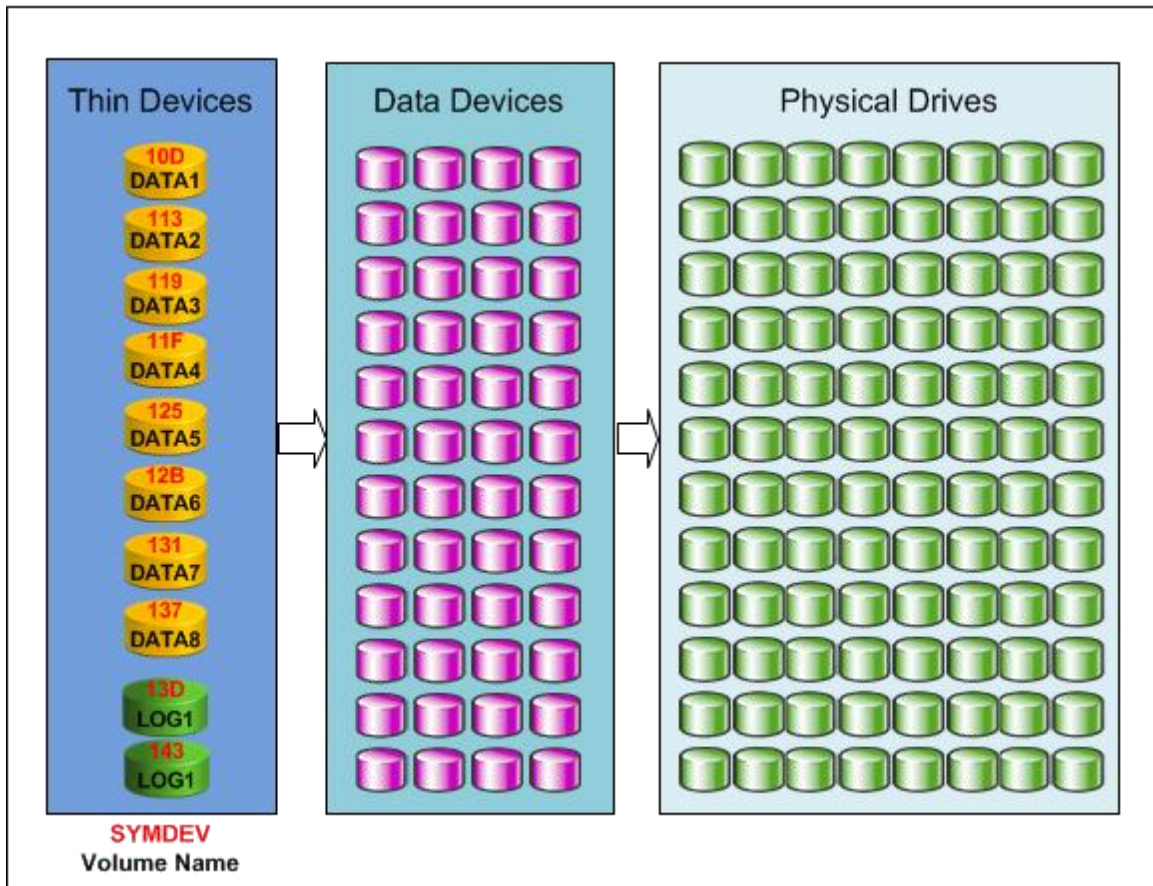


Figure 5. Relationship of thin LUNs to data devices and physical drives

Each thin LUN was defined as approximately 360 GB in size, and was configured as a four-member striped thin metavolume. Each LUN contained a single NTFS volume, and each NTFS volume was mounted at a mount point located under the directory C:\SQLDB.

Actual storage allocations for the thin devices were based on the actual utilization, and as EMC best practices were followed (Quick NTFS volume format and SQL Server Instant File Initialization) the actual allocations were initially minimal.

The *symcfg* command can be used to detail the current allocations of storage extents for thin devices. This command was used to detail the thin device allocations after the configuration was defined, and the database was restored. Again, when initially configured, all storage allocation was provided by the SQL_FC thin pool.

```
C:> symcfg -sid 1667 show -pool SQL_FC -thin -detail

Symmetrix ID: 000192601667

Symmetrix ID           : 000192601667
Pool Name              : SQL_FC
Pool Type              : Thin
Dev Emulation          : FBA
Dev Configuration     : RAID-5 (3+1)
Pool State             : Enabled
```

```

# of Devices in Pool      : 48
# of Enabled Devices in Pool : 48
Max. Subscription Percent : None
Rebalance Variance      : 1%
Max devs per rebalance scan : 256

```

<... content removed from output ...>

Pool Bound Thin Devices(12):

```

{
-----

```

Sym Dev	Total Tracks	Pool Subs (%)	Pool Allocated		Total Written		Status
			Tracks	(%)	Tracks	(%)	
010D	5898240	13	1908	0	1592	0	Bound
0113	5898240	13	2442444	41	2299789	39	Bound
0119	5898240	13	2442432	41	2299067	39	Bound
011F	5898240	13	2439228	41	2296789	39	Bound
0125	5898240	13	2438844	41	2296242	39	Bound
012B	5898240	13	251364	4	247130	4	Bound
0131	5898240	13	1026384	17	965385	16	Bound
0137	5898240	13	1025832	17	965388	16	Bound
013D	5898240	13	1921620	33	1921388	33	Bound
0143	5898240	13	1921620	33	1921388	33	Bound
0149	5898240	13	1560	0	1387	0	Bound
014F	5898240	13	1560	0	1387	0	Bound
Tracks	70778880	150	15914796	34	15216932	32	

```

-----
}

```

The storage allocations can then be mapped to the SQL Server database file components as shown in Figure 4 on page 15 and highlight the space efficiency provided by allocation on demand for thin devices used with SQL Server database files. For example, the Broker files are located on devices 113, 119, 11F, and 125. These thin devices each have only 39 percent allocation of the entire LUN size as a result of SQL Server proportional fill mechanisms and the on-demand allocation of storage on thin devices.

The simulated environment was that of a TPC-E-like environment. As such, it represents the workload as anticipated of an online brokerage environment. Simulated users perform a range of processing tasks, such as stock transactions, and look-ups, as well as the generation of larger reports. The workload generated a significant read/write requirement against the various portions of the database.

The Windows Perfmon instrumentation was set to a 60-second interval and the results before, during, and after FAST VP movements were collected and analyzed. The results are presented in the following sections.

The tests outlined in this white paper were performed on the following configuration:

Configuration aspect	Description
Storage controller	Symmetrix VMAX 1667
Microcode	Enginuity 5875
Microsoft Windows	Microsoft Windows Server 2008 R2
SQL Server	Microsoft SQL Server 2008 R2
Volumes	12 x 6-member thin striped metavolume
Symmetrix EFD Tier	4 x 400 GB Flash drives
Symmetrix FC Tier	96 x 450 GB 15k rpm Fibre Channel drives
Symmetrix SATA Tier	38 x 1 TB SATA drives

Configuring FAST VP for the environment

When initially configured, all LUNs were bound to a single thin pool (FC_SQL), resulting in all allocations being provided by this single thin pool. All thin pools were configured to utilize a RAID 5 3+1 protection type. The Symmetrix VMAX array contained a number of different technologies including EFD, 450 GB 15k rpm drives, and 1 TB SATA drives.

In the tested environment, multiple storage tiers were defined, as shown in Figure 6. The tiers created in this configuration varied in terms of the underlying technology, where the tier “EFD_5R3” was defined as RAID 5 3+1 protection on EFD, “FC_5R3” was defined as being a RAID 5 3+1 protection scheme on Fibre Channel 450 GB 15k rpm drives, and “SATA_5R3” was defined as being a RAID 5 3+1 protection scheme on 1 TB 5,400 rpm drives. These three tiers were subsequently used to define a policy for the SQL Server environment.

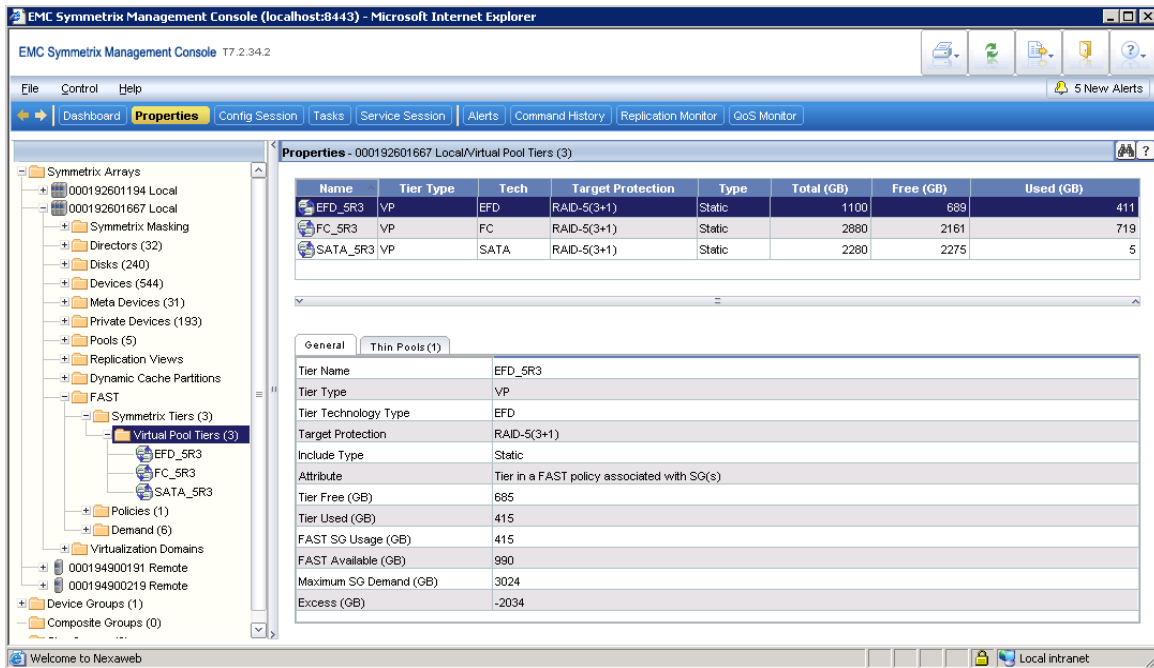


Figure 6. Defining FAST VP storage tiers within Symmetrix Management Console

The storage tiers were applied to create a policy named “SQLPRD_POLICY”. This policy defined the applicable storage tiers, and applied capacity percentages to the various tiers in SMC. The percentages define how much of the storage space currently allocated by the storage group defined in the policy can be utilized from each of the tiers. Thus, the values as shown in Figure 7 indicate that when applied to a storage group, 100 percent of the storage allocation for the group may come from the FC_5R3 storage tier, 30 percent may come from the EFD_5R3 storage tier, and 100 percent may be allocated from the SATA_5R3 storage tier. The total allocation in this instance is 230 percent, and will allow the FAST policy to allocate storage from the most appropriate tier based on its statistical sampling.

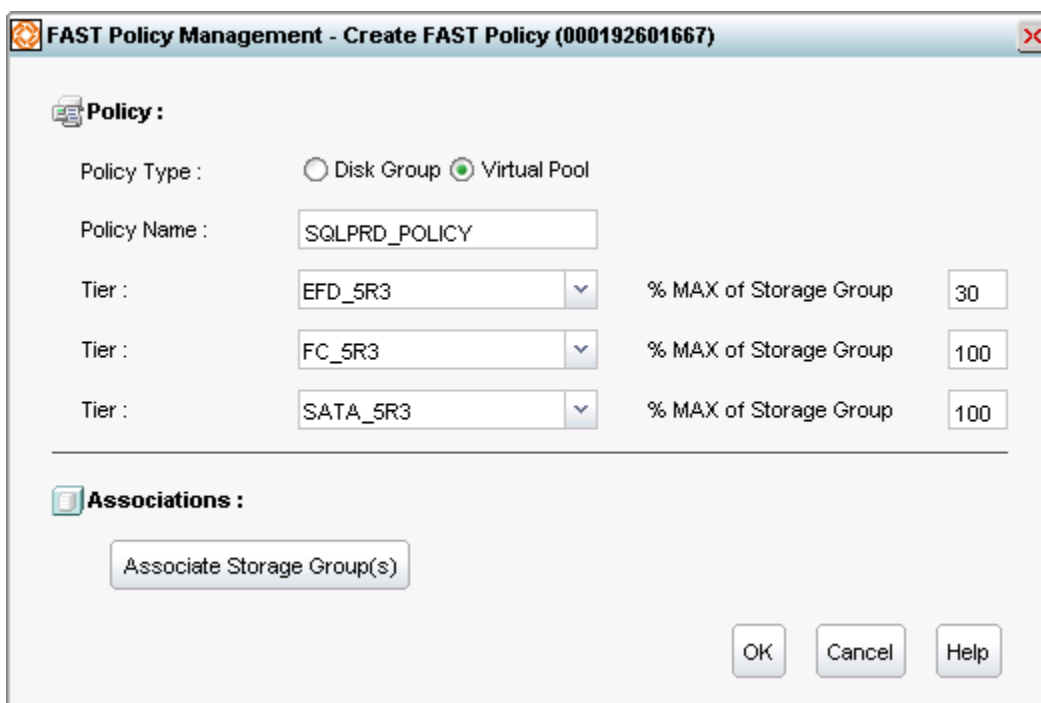


Figure 7. FAST VP policy definition within Symmetrix Management Console

The association of a storage group to a defined tier binds the policy to the devices contained within the storage group. The storage group is defined when implementing Auto-provisioning Groups, and defines the storage devices that will be available to a host when bound to a view. For the tested workload the storage group name was “SQLPRD_Thin_Devs” and defined all storage devices that were presented to the SQL Server host. This storage group naming is typically conducted when initially provisioning storage to the SQL Server host during initial deployment. In Figure 8 the storage group is displayed with its component thin devices.

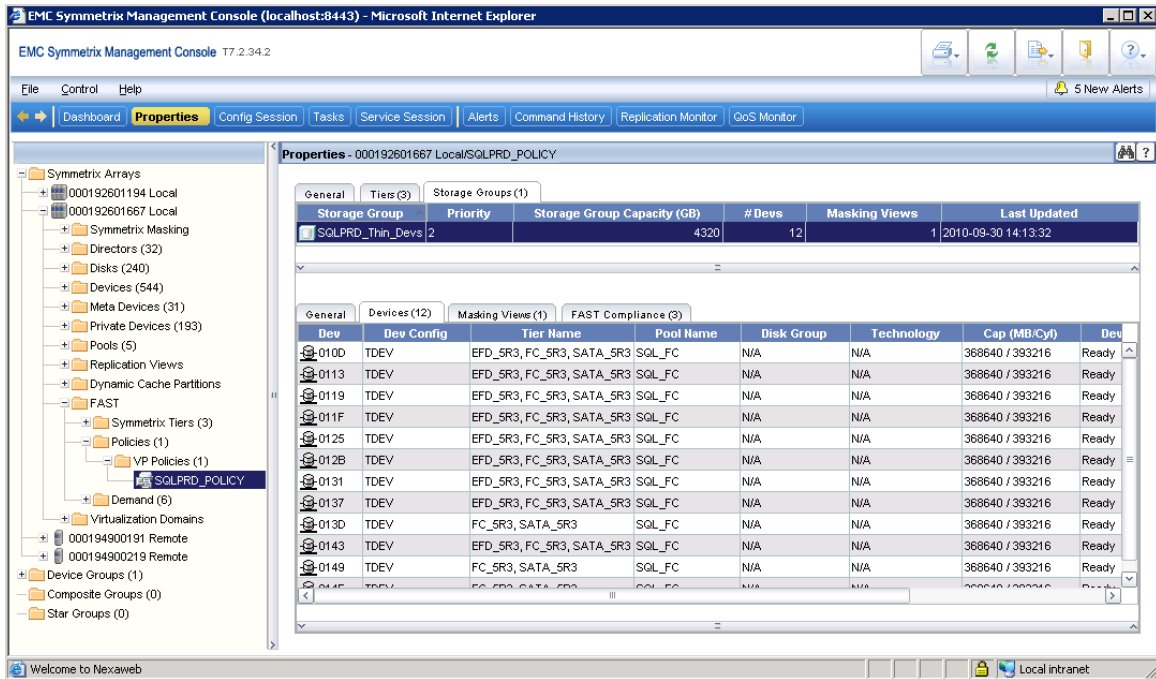


Figure 8. Allocating a storage group to a policy in Symmetrix Management Console

To complete the implementation of the FAST VP mechanisms, it is necessary to set appropriate time periods for the policy engine to collect statistics for workload analysis. Statistics collection is defined within the Symmetrix Optimizer environment, and may also be set through SMC during execution of the FAST Configuration Wizard, as shown in Figure 9.

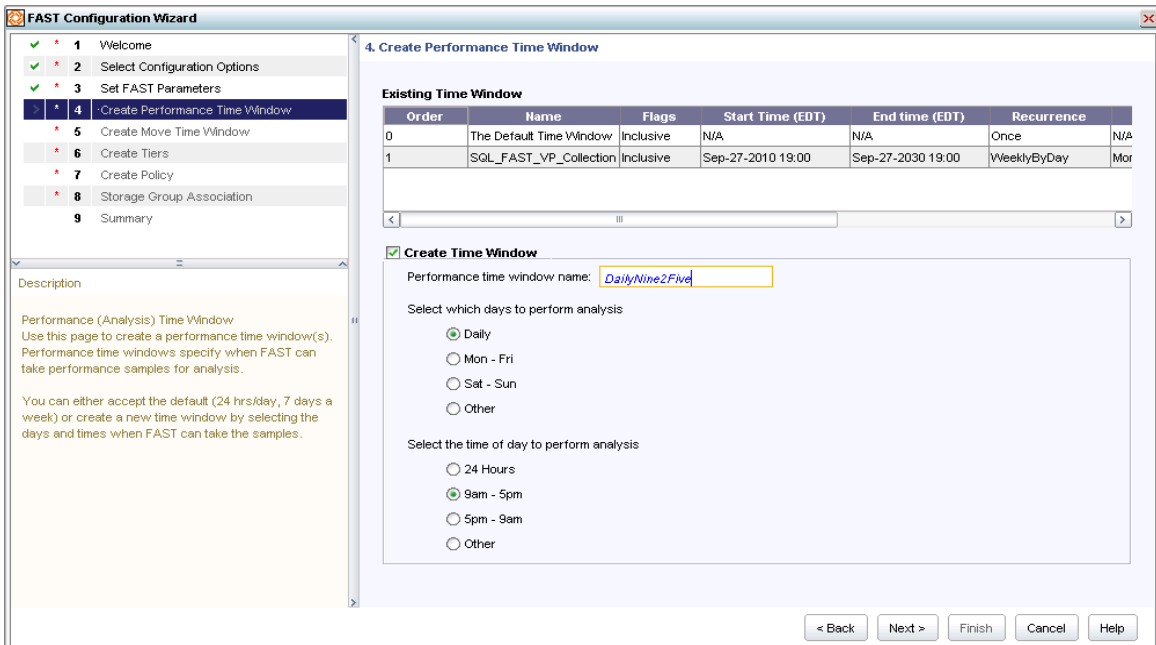


Figure 9. Symmetrix FAST Configuration Wizard performance time window

When implementing FAST VP, all movement operations are fully automatic, and there is no user approval mode. All storage movements relating to FAST VP policies are executed during the periods specified in the Movement Time Window. Figure 10 displays the configuration settings available through the Symmetrix Management Console FAST Configuration Wizard. Multiple time windows may be specified for movements to occur.

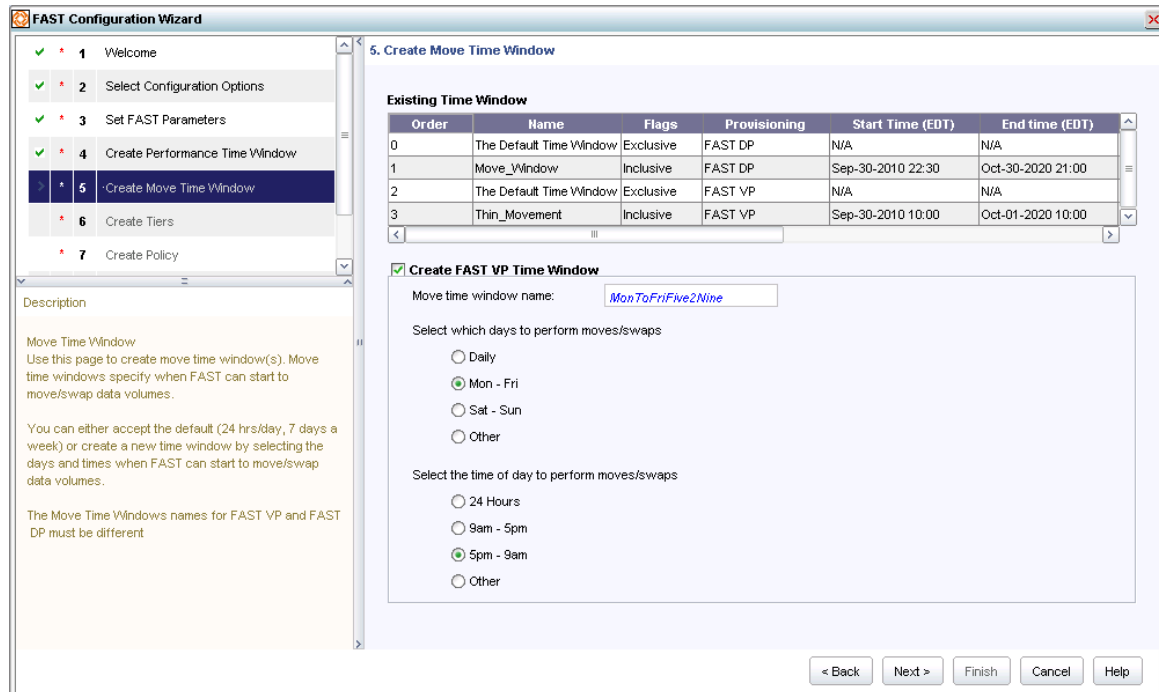


Figure 10. Symmetrix FAST Configuration Wizard movement time window

Having defined the storage group association to the storage tiers through the policy, and scheduling the application time periods within SMC, the environment was configured appropriately.

Monitoring workload performance

Microsoft SQL Server database administrators and Windows Server administrators will typically utilize Windows performance counters to monitor overall system performance. The instrumentation provided by the Windows performance counters represents a valid performance profile of the behavior of the storage devices presented to the host. Additionally, SQL Server implements a Performance Warehouse implementation that can be used to specifically monitor those portions of the environment utilized by a SQL Server database. The SQL Server Performance Warehouse can be a valuable tool for database administrators to identify specific performance and transaction counters.

Only Windows Performance Monitor counters were collected during testing. When implementing FAST VP, it is important to consider the storage movement. Unlike standard FAST, where an entire LUN is moved to an alternate tier, FAST VP only moves portions of the thin devices defined within the policy. As a result, some components

may be located on any of the tiers and will exhibit the performance characteristics of the tier on which they are located. As workloads change on the storage allocations, FAST VP will transparently migrate the data to the most appropriate tier.

Identification of database workloads

The user workload was executed for an extended period of time to ensure that the workload reached a constant state and that the initial performance statistics time collection interval as defined within Symmetrix Management Console, for FAST VP, had passed. Windows Performance Monitor counters were being collected over the time interval. In Figure 11 overviews of the Disk Reads and Writes per Second are displayed, covering a 4-hour time period when the initial workload was executed.

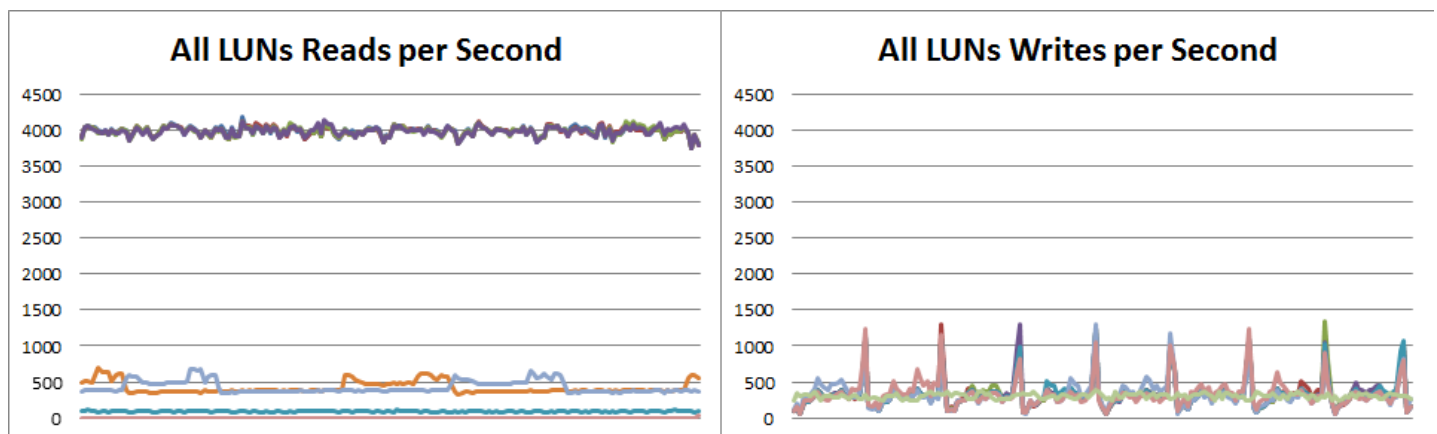


Figure 11. Windows performance counters for reads and writes IOPs

During the initial workload period, Read operations per second were averaging 4,000 for the Broker FileGroup. Other file groups produced significantly less workload. Write operations did not approach the levels seen with reads but did exceed 1,000 IOPS during checkpoint intervals.

Latencies for the LUNs were also recorded, and are displayed in Figure 12 and show that read operation latencies for the most heavily accessed LUNs (Broker FileGroup) were approaching 20 milliseconds. Write latencies were lower with the most heavily accessed LUNs maintaining an 11 ms latency.




Figure 12. Windows performance counters for read and write latencies

It is clear therefore that there are differing performance characteristics for the various data files. The Broker files have a significant workload generated against them as compared to the other files and resulting LUNs.

Effect of FAST VP

Migrations of storage extents in the LUNs defined for the FAST VP storage group will begin executing after the initial statistics collection period and within a valid movement time window. In the tested configuration, the movement time window was set for a 24-hour period, seven days a week. The limiting factor was the initial statistics collection time interval, which was set to 7 hours. After this time period had elapsed, the movements began.

The size of the movements will be based on the utilization and access patterns to the various storage allocations for the LUNs. Details on the sizes and mechanisms used are covered in the *Implementing Fully Automated Storage Tiering with Virtual Pools (FAST VP) for EMC Symmetrix VMAX Series Arrays Technical Notes*, available on Powerlink®.

As storage extents are moved between the various tiers defined within the policy, performance of the thin devices will begin to exhibit the characteristics of the tier to which the extents have been moved. The resulting workload profile can be seen in Figure 13, which covers the initial statistics collection time interval as well as the time period where the majority of migrations are occurring and then the time period where the workload reaches a stable point.




Figure 13. Read and write workload before, during, and after migrations

The read workload exhibits the largest change in profile as extents are migrated to the various storage tiers. In the tested environment, the reads per second increased from a rate of around 4,000 operations per second for the Broker FileGroup to a level just below 5,000 operations per second. Other SQL Server FileGroups can also be seen to have an increase in throughput, but to a lesser degree than the Broker FileGroup.

As with the changes in read and writes throughputs, the latency numbers were also positively affected, and these results can be seen in Figure 14. It can be seen that there is little impact to the latencies during the period of migration of extents, and that as a result of the migrations, the latency figures for both read and write operations improve.

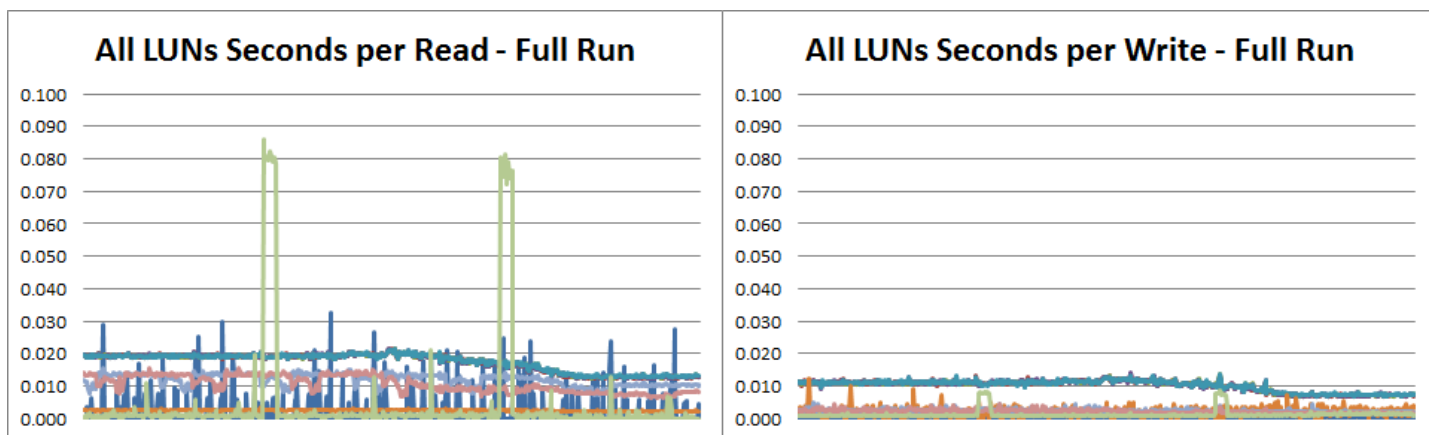


Figure 14. Read and write latencies before, during, and after migrations

Using Solutions Enabler commands, it is possible to monitor the migrations between the tiers defined within the policy. The symfast command can be utilized to show the current allocations across the tiers defined. In the tested configuration, the command `symfast -sid 1667 list -association -demand -mb` was used to collect information on the tiers being used. The command provides details on the current allocation in the format shown next.

```

Symmetrix Id      : 000192601667
Policy Name       : SQL_PROD_POLICY
Storage Group    : SQLPRD_Thin_Devs
Priority          : 2

```

```
Tiers (3)
```

```
{
-----
                                     Logical Capacities (MB)
-----
Name          Type  Target  Max SG   Max SG   FAST SG   Growth
              Prot  Percent Demand  Usage
-----
EFD_5R3      THIN R5 (3+1)  30    298403    0    +298403
FC_5R3       THIN R5 (3+1) 100   994675   994675 +0
SATA_5R3     THIN R5 (3+1) 100   994675    0    +994675

      Total
-----
                                     2287753  994675
-----
}
```

The output from this command was collected over time to monitor the change in allocations. The results from this output are summarized in Figure 15 and show that extent migrations occurred between the Fibre Channel and EFD tiers in a highly correlated manner.

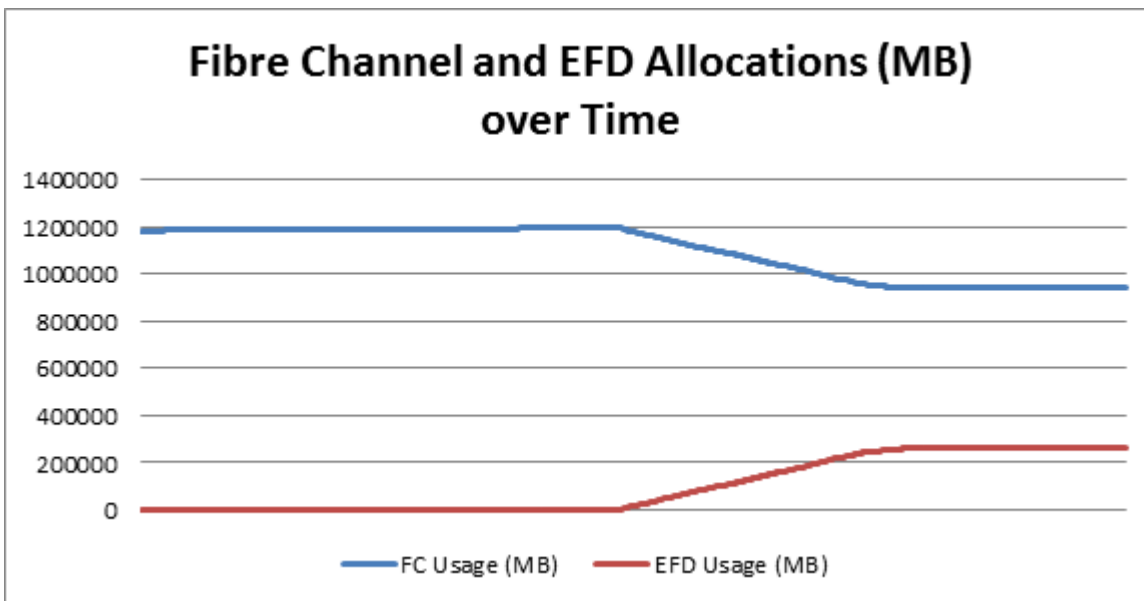


Figure 15. Summary of thin pool allocations over time

The FAST VP policy in place will be monitoring all thin LUNs allocated to the storage group, and statistical sampling will be occurring for all LUNs. Migrations between the tiers defined in the policy will be based on the level of activity. It should be expected

that the most heavily accessed storage allocation, or those that contain the most heavily accessed data, will be identified for migrations to a higher-performing storage tier. In Figure 16 one of the metavolumes used for a Broker file is used to display the storage allocations during a similar time period as previously shown for all tiers.

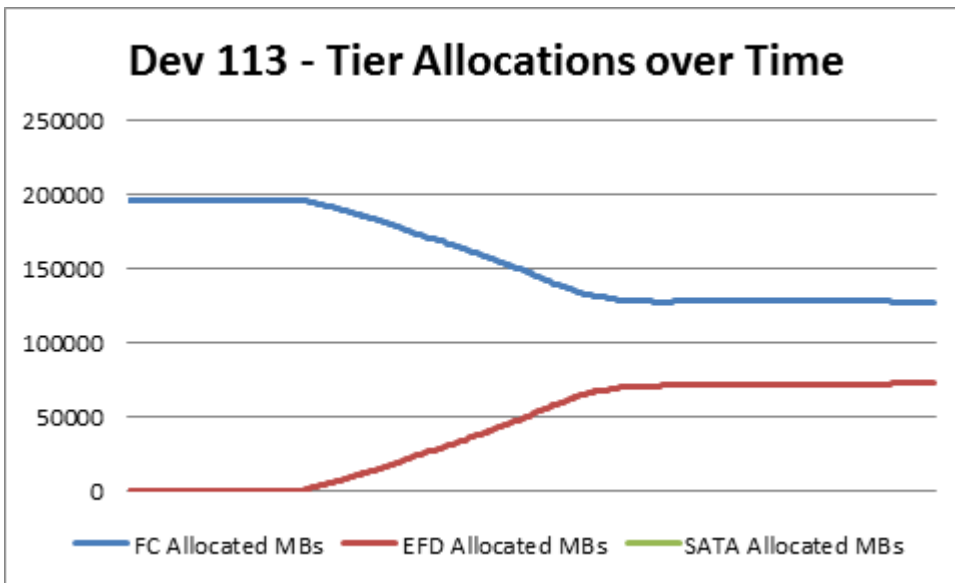


Figure 16. Storage allocations for a LUN used by a Broker file

Storage allocations also occurred against the SATA storage tier; however this was not to the same degree as the migration into the EFD tier. SATA tier allocations are shown in Figure 17, and comprise a significantly smaller allocation rate and overall size than the Fibre Channel and EFD pools.

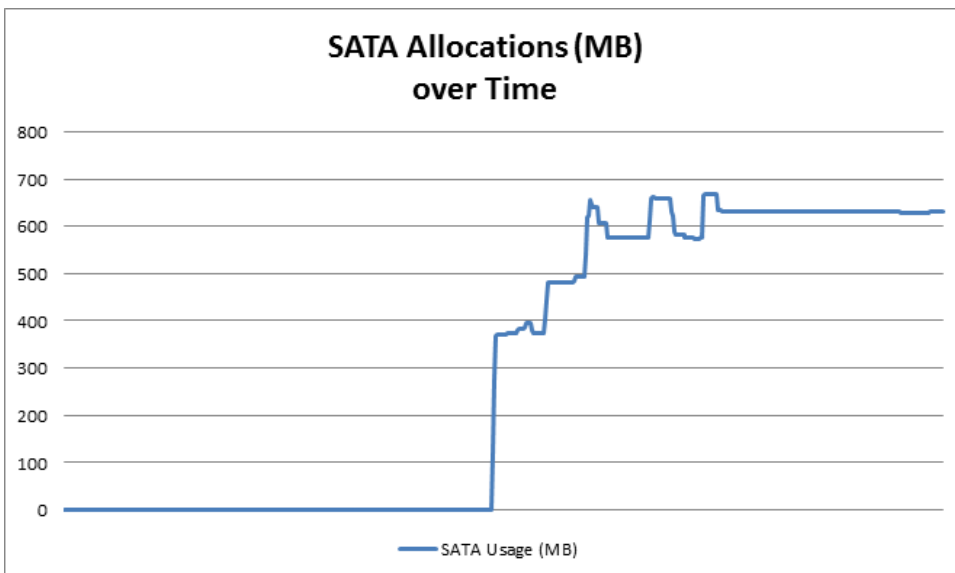


Figure 17. Storage allocations for the SATA tier

The main cause of the minimal usage of the SATA tier is based on the workload generated by the workload used during the testing. The workload generates a

constant workload across all data files comprising the database storage. This constant workload will limit the ability to move extents to a higher density storage pool, since demotions will require no workload to be generated across the storage extents.

The only demotions that could occur are for NTFS metadata structures that generate no workload access over time. It is these structures that tend to be demoted to the SATA tier. The storage allocations for these structures is relatively small, and account for some of the storage allocations on the LUNs. Additionally, two LUNs were provided to store TEMPDB in the environment. These LUNs generated no workload during the testing, and were fully migrated to the SATA pool.

It should be noted that the LUNs were thin devices, and were appropriately configured using best practice implementation for Windows. These best practices will result in minimal storage allocations occurring on the thin devices themselves. Thus these storage allocations are insignificant as compared to the approximate 1.3 TB allocated for the database files and transaction logs.

Assessing the benefits of FAST VP

The migrations automatically implemented by the FAST VP Policy engine improved overall performance of the total environment by migrating heavily accessed storage allocations to a higher-performing storage tier. In the example configuration, a large proportion of the more active storage allocations of the database (Broke FileGroup) were migrated to the EFD storage tier, resulting in improved workload throughput and reduction in I/O latency time for these LUNs. The migration of these heavily accessed storage allocations also resulted in a reduction of contention within the original storage pool, thereby improving response times for the remaining devices.

Post-migration I/O throughput rates were significantly better than pre-migration rates, and were executing with significantly reduced latency numbers. Figure 18 shows that overall system read requests per second increased over the duration of the testing. Reads increased from a rate of around 17,000 to around 21,000 reads/sec, while providing improved latency as seconds/read rates dropped from around 12 ms per read to an average of 8 ms per read.

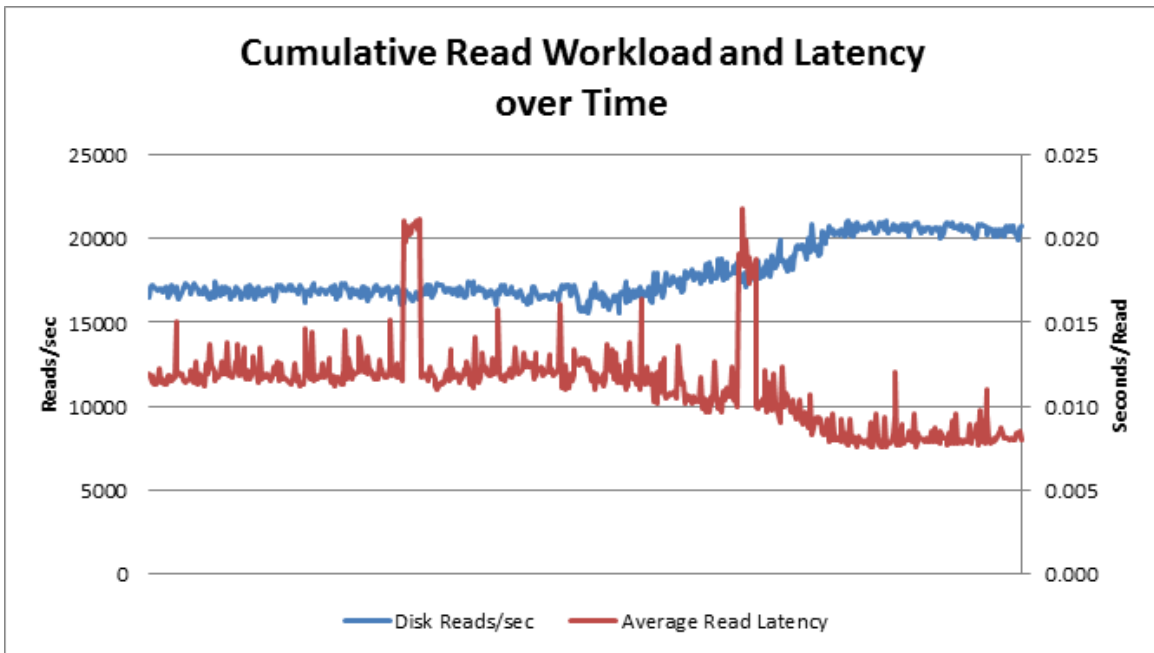


Figure 18. Read I/O rates for data volumes - Post-migration

As a direct result of higher throughput and lower latencies, the operation of the SQL Server environment improved across all major attributes. Figure 19 details the relative improvement in performance for the environment. The performance improvements were obtained after only 28 percent of the total storage allocations across all LUNs had been migrated to an alternate storage tier.

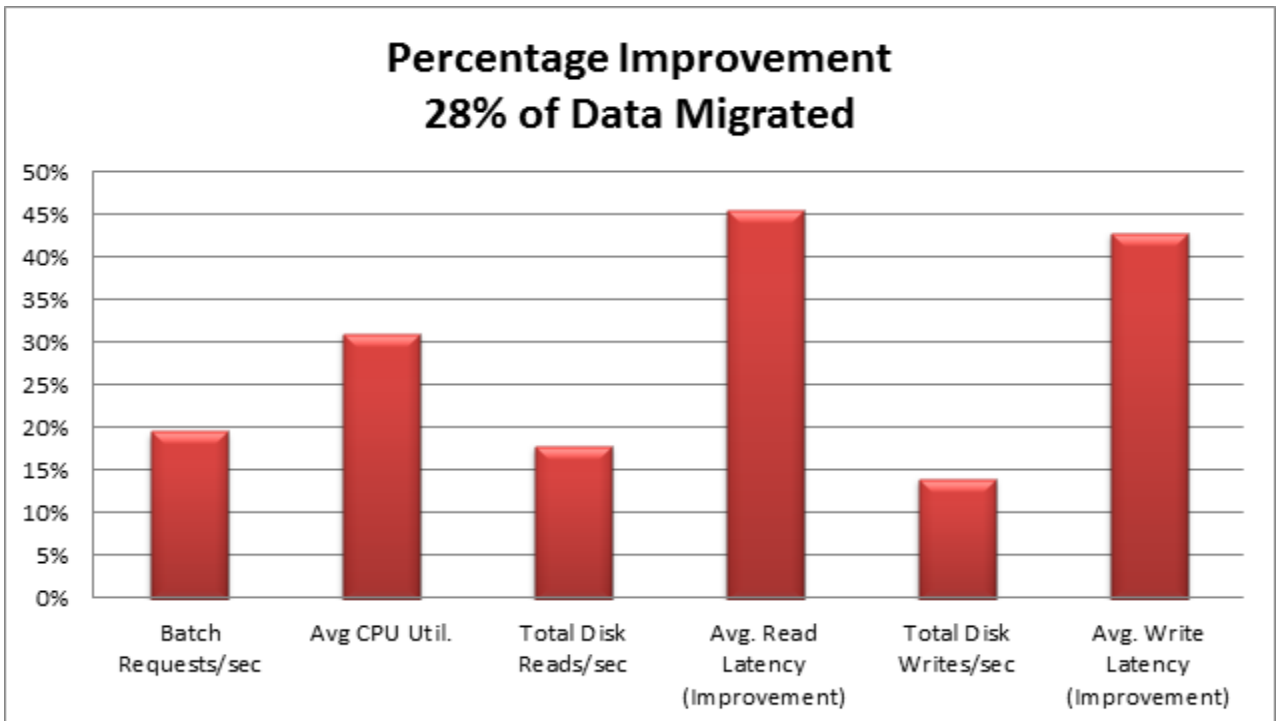


Figure 19. Comparison of improvement for major metrics

The movement of storage extents to a lower, more cost-effective tier was limited by the utilization of the storage devices. SQL Server ensures that there is a uniform distribution of I/O across all data files. This style of activity ensures that disk resources are equally accessed. SQL partitioning, and refinements in the allocation of data and indexes to discrete data files, may improve the overall efficiency of the system.

Conclusion

Microsoft SQL Server database environments often exhibit a skewed workload across files comprising differing filegroups. In many implementations, the LUNs used as storage devices are allocated from the same storage pools, which will often lead to contention for physical resources. In enterprise-class storage arrays, tiered storage configurations are implemented to provide differing performance and cost-effective storage solutions.

Implementing FAST VP can improve performance across the array. The sub-LUN blocks migrated to higher-performance storage can significantly improve application performance, while also improving performance of the remaining non-migrated data due to diminished contention. Moving lightly accessed storage allocations from a higher-performing storage tier to a more appropriate storage tier helps improve capacity utilization. The cost savings from tiered storage implementations can be realized in lower acquisition costs and long-term operational savings, especially in lower long-term power usage. Symmetrix VMAX with Enginuity 5875 enables a further refinement in operational efficiencies to ensure that data placement is optimized based on workload requirements.