

EMC VNX FAST VP

VNX5100, VNX5300, VNX5500, VNX5700, & VNX7500
A Detailed Review

Abstract

This white paper discusses EMC® Fully Automated Storage Tiering for Virtual Pools (FAST VP) technology and describes its features and implementation. Details on how to use the product in Unisphere™ are discussed, and usage guidance and major customer benefits are also included.

December 2013

Copyright © 2013 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided “as is.” EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

VMware is a 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 h8058.4

Table of Contents

Executive summary	4
Audience.....	4
Introduction	5
Storage tiers.....	6
Extreme Performance Tier drives: Flash drives	6
Performance Tier drives: SAS drives.....	6
Capacity Tier drives: NL-SAS drives.....	7
FAST VP operations	8
Storage pools.....	8
FAST VP algorithm	9
Statistics collection.....	9
Analysis	10
Relocation.....	10
Managing FAST VP at the storage pool level.....	11
Per-Tier RAID Configuration.....	12
Automated scheduler.....	13
Manual Relocation	14
FAST VP LUN management.....	14
Tiering policies.....	15
Common uses for Highest Available Tiering policy.....	17
Using FAST VP for File	18
Management.....	18
Best practices for VNX for File.....	21
General Guidance and recommendations	22
FAST VP and FAST Cache.....	22
What drive mix is right for my I/O profile?	23
Conclusion	24
References	24

Executive summary

Fully Automated Storage Tiering for Virtual Pools (FAST VP) can lower the total cost of ownership (TCO) and increase performance by intelligently managing data placement according to activity level. When FAST VP is implemented, the storage system measures, analyzes, and implements a dynamic storage-tiering policy much faster and more efficiently than a human analyst could ever achieve.

Storage provisioning can be repetitive and time consuming and, when estimates are calculated incorrectly, it can produce uncertain results. It is not always obvious how to match capacity to the performance requirements of a workload's data. Even when a match is achieved, requirements change, and a storage system's provisioning requires constant adjustments.

Storage tiering allows a storage pool to use varying levels of drives. LUNs use the storage capacity needed from the pool on the devices with the required performance characteristics. FAST VP uses I/O statistics at a 1 GB slice granularity (known as *sub-LUN tiering*). The relative activity level of each slice is used to determine the need to promote to higher tiers of storage. Relocation is initiated at the user's discretion through either manual initiation or an automated scheduler. FAST VP removes the need for manual, resource intensive LUN migrations while still providing the performance levels required by the most active dataset.

FAST VP is a licensed feature available on the EMC® VNX5300™ and larger systems. FAST VP licenses are available as part of a FAST Suite of licenses that offers complementary licenses for technologies such as FAST Cache, Analyzer, and Quality of Service Manager.

This white paper discusses the EMC FAST VP technology and describes its features, functions, and management.

Audience

This white paper is intended for EMC customers, partners, and employees who are considering using the FAST VP product. Some familiarity with EMC midrange storage systems is assumed. Users should be familiar with the material discussed in the white papers *Introduction to EMC VNX Series Storage Systems* and *EMC VNX Virtual Provisioning*.

Introduction

Data has a lifecycle. As data progresses through its lifecycle, it experiences varying levels of activity. When data is created, it is typically heavily used. As it ages, it is accessed less often. This is often referred to as being temporal in nature. FAST VP is a simple and elegant solution for dynamically matching storage requirements with changes in the frequency of data access. FAST VP segregates disk drives into three tiers:

- Extreme Performance Tier —Flash drives
- **Performance Tier** —Serial Attach SCSI (SAS) drives
- **Capacity Tier** —Near-Line SAS (NL-SAS) drives

You can use FAST VP to aggressively reduce TCO and to increase performance. By going from all-SAS drives to a mix of Flash, SAS, and NL-SAS drives, you can address the performance requirements and still reduce the drive count. In some cases, an almost two-thirds reduction in drive count is achieved. In other cases, performance throughput can double by adding less than 10 percent of a pool's total capacity in Flash drives.

FAST VP has been proven highly effective for a number of applications. Tests in OLTP environments with Oracle¹ or Microsoft SQL Server² show that users can lower their capital expenditure by 15 percent to 38 percent, reduce power and cooling costs (by over 40 percent), and still increase performance by using FAST VP instead of a homogeneous drive deployment. More details regarding these benefits are located in the [What drive mix is right for my I/O profile?](#) section of this paper.

You can use FAST VP in combination with other performance optimization software, such as FAST Cache. A common strategy is to use FAST VP to gain TCO benefits while using FAST Cache to boost overall system performance. There are other scenarios where it makes sense to use FAST VP for both purposes. This paper discusses considerations for the best deployment of these technologies.

The VNX series of storage systems delivers even more value over previous systems by providing a unified approach to auto-tiering for file and block data. Now, file data served by VNX Data Movers can also use virtual pools and most of the advanced data services as block data. This provides compelling value for users who want to optimize the use of high-performing drives across their environment.

The key enhancements available in the new EMC® VNX™ Operating Environment (OE) for block release 5.32, file version 7.1 are:

- [New default FAST VP policy: Start High, then Auto-Tier](#)
- [Per-Tier RAID Selection: Additional RAID configurations](#)

¹ *Leveraging Fully Automated Storage Tiering (FAST) with Oracle Database Applications* EMC white paper

² *EMC Tiered Storage for Microsoft SQL Server 2008—Enabled by EMC Unified Storage and EMC Fully Automated Storage Tiering (FAST)* EMC white paper

- [Pool Rebalancing upon Expansion](#)
- [Ongoing Load-Balance within tiers \(with FAST VP license\)](#)

Note: EMC® VNX™ Operating Environment (OE) for block release 5.32, file version 7.1 supports VNX platforms only.

Storage tiers

FAST VP can leverage two or all three storage tiers in a single pool. Each tier offers unique advantages in performance and cost.

FAST VP differentiates tiers by drive type. However, it does not take rotational speed into consideration. EMC strongly encourages you to avoid mixing rotational speeds per drive type in a given pool. If multiple rotational-speed drives exist in the array, multiple pools should be implemented as well.

Extreme Performance Tier drives: Flash drives

Flash technology is revolutionizing the external-disk storage system market. Flash drives are built on solid-state drive (SSD) technology. As such, they have no moving parts. The absence of moving parts makes these drives highly energy-efficient, and eliminates rotational latencies. Therefore, migrating data from spinning disks to Flash drives can boost performance and create significant energy savings.

Tests show that adding a small (single-digit) percentage of Flash capacity to your storage, while using intelligent tiering products (such as FAST VP and FAST Cache), can deliver double-digit percentage gains in throughput and response time performance in some applications. Flash drives can deliver an order of magnitude better performance than traditional spinning disks when the workload is IOPS-intensive and response-time sensitive. They are particularly effective when small random-read I/Os with a high degree of parallelism are part of the I/O profile, as they are in many transactional database applications. On the other hand, bandwidth-intensive applications perform only slightly better on Flash drives than on spinning drives.

Flash drives have a higher per-gigabyte cost than traditional spinning drives and lower per I/O cost. To receive the best return, you should use Flash drives for data that requires fast response times and high IOPS. A good way to optimize the use of these high-performing resources is to allow FAST VP to migrate “hot” data to Flash drives at a sub-LUN level.

Performance Tier drives: SAS drives

Traditional spinning drives offer high levels of performance, reliability, and capacity. These drives are based on industry-standardized, enterprise-level, mechanical hard-drive technology that stores digital data on a series of rapidly rotating magnetic platters.

The Performance Tier includes 10K and 15K rpm spinning drives, which are available on all EMC midrange storage systems. They have been the performance medium of choice for many years, and have the highest availability of any mechanical storage device. These drives continue to serve as a valuable storage tier, offering high all-around performance, including consistent response times, high throughput, and good bandwidth, at a mid-level price point.

As noted above, users should choose between either 10K or 15K drives in a given pool.

Capacity Tier drives: NL-SAS drives

Capacity drives are designed for maximum capacity at a modest performance level. They have a slower rotational speed than Performance Tier drives. NL-SAS drives for the VNX series have a 7.2K rotational speed.

Using capacity drives can significantly reduce energy use and free up more expensive, higher-performance capacity in higher storage tiers. Studies have shown that 60 percent to 80 percent of the capacity of many applications has little I/O activity. Capacity drives can cost about four times less than performance drives on a per-gigabyte basis, and a small fraction of the cost of Flash drives. They consume up to 96 percent less power per TB than performance drives. This offers a compelling opportunity for TCO improvement considering both purchase cost and operational efficiency.

However, the reduced rotational speed is a trade-off for significantly larger capacity. For example, the current Capacity Tier drive offering is 2 TB, compared to the 600 GB Performance Tier drives and 200 GB Flash drives. These Capacity Tier drives offer roughly half the IOPS/drive of Performance Tier drives. Future drive offerings will have larger capacities, but the relative difference between disks of different tiers is expected to remain approximately the same.

Table 1. Feature tradeoffs for Flash, Performance, and Capacity Tier drives

	Extreme Performance (Flash)	Performance (SAS)	Capacity (NL-SAS)
Performance	<ul style="list-style-type: none"> High IOPS/GB and low latency User response time between 1 to 5 ms Multi-access response time less than 10 ms 	<ul style="list-style-type: none"> High bandwidth with contending workloads User response time is ~5 ms Multi-access response time is between 10 to 50 ms 	<ul style="list-style-type: none"> Low IOPS/GB User response time between 7 to 10 ms Multi-access response time up to 100 ms Leverages storage array SP cache for sequential and large block access
Strengths	<ul style="list-style-type: none"> Provide extremely fast access for reads Execute multiple sequential streams better than SAS 	<ul style="list-style-type: none"> Sequential reads leverage read-ahead Sequential writes leverage system optimizations favoring disks Read/write mixes provide predictable performance 	<ul style="list-style-type: none"> Large I/O is serviced efficiently Sequential reads leverage read-ahead Sequential writes leverage system optimizations favoring disks
Limitations	<ul style="list-style-type: none"> Writes slower than read Heavy concurrent writes affect read rates Single-threaded large sequential I/O equivalent to SAS 	<ul style="list-style-type: none"> Uncached writes are slower than reads 	<ul style="list-style-type: none"> Long response times for heavy-write loads Not as good as SAS/FC at handling multiple streams

FAST VP operations

FAST VP operates by periodically relocating the most active data up to the highest available tier (typically the Extreme Performance or Performance Tier). To ensure sufficient space in the higher tiers, FAST VP relocates less active data to lower tiers (Performance or Capacity Tiers) when new data needs to be promoted. FAST VP works at a granularity of 1 GB. Each 1 GB block of data is referred to as a “slice.” FAST VP relocates data by moving the entire slice to highest available storage tier.

Storage pools

Storage pools are the framework that allows FAST VP to fully use each of the storage tiers discussed. Heterogeneous pools are made up of more than one type of drive. LUNs can then be created within the pool. These pool LUNs are no longer bound to a

single storage tier; instead, they can be spread across different storage tiers within the same pool.

You also have the flexibility of creating homogeneous pools, which are composed by only one type of drive and still benefit from some FAST VP capabilities.

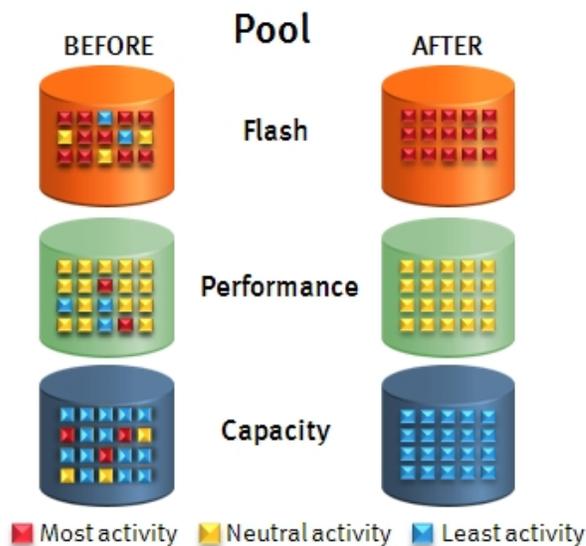


Figure 1. Heterogeneous storage pool concept

LUNs must reside in a pool to be eligible for FAST VP relocation. Pools support thick LUNs and thin LUNs. Thick LUNs are high-performing LUNs that use logical block addressing (LBA) on the physical capacity assigned from the pool. Thin LUNs use a capacity-on-demand model for allocating drive capacity. Thin LUN capacity usage is tracked at a finer granularity than thick LUNs to maximize capacity optimizations. FAST VP is supported on both thick LUNs and thin LUNs.

RAID groups are by definition homogeneous and therefore are not eligible for sub-LUN tiering. LUNs in RAID groups can be migrated to pools using LUN Migration. For a more in-depth discussion of pools, please see the white paper *EMC VNX Virtual Provisioning - Applied Technology*.

FAST VP algorithm

FAST VP uses three strategies to identify and move the correct slices to the correct tiers: statistics collection, analysis, and relocation.

Statistics collection

A slice of data is considered *hotter* (more activity) or *colder* (less activity) than another slice of data based on the relative activity level of those slices. Activity level is determined simply by counting the number of I/Os, which are reads and writes bound for each slice. FAST VP maintains a cumulative I/O count and weights each I/O by how recently it arrived. This weighting deteriorates over time. New I/O is given full weight. After approximately 24 hours, the same I/O will carry only about half-weight.

Over time the relative weighting continues to decrease. Statistics collection happens continuously in the background on all pool LUNs.

Analysis

Once per hour, the collected data is analyzed. This analysis produces a rank ordering of each slice within the pool. The ranking progresses from the hottest slices to the coldest. This ranking is relative to the pool. A hot slice in one pool may be cold by another pool's ranking. There is no system-level threshold for activity level. The user can influence the ranking of a LUN and its component slices by changing the tiering policy, in which case the tiering policy takes precedence over the activity level.

Relocation

During user-defined relocation windows, slices are promoted according to the rank ordering performed in the analysis stage. During relocation, FAST VP prioritizes relocating slices to higher tiers. Slices are only relocated to lower tiers if the space they occupy is required for a higher priority slice. In this way, FAST attempts to ensure the maximum utility from the highest tiers of storage. As data is added to the pool, it is initially distributed across the tiers and then moved up to the higher tiers if space is available. Ten percent space is maintained in each of the tiers to absorb new allocations that are defined as "Highest Available Tier" between relocation cycles. Lower tier spindles are used as capacity demand grows.

The following two features are new enhancements that are included in the EMC® VNX™ Operating Environment (OE) for block release 5.32, file version 7.1:

Pool Rebalancing upon Expansion

When a storage pool is expanded, the sudden introduction of new empty disks combined with relatively full existing disks causes a data imbalance. This imbalance is resolved by automating a one-time data relocation, referred to as rebalancing. This rebalance relocates slices within the tier of storage that has been expanded, to achieve best performance. Rebalancing occurs both with and without the FAST VP enabler installed.

Ongoing Load-Balance within tiers (With FAST VP license)

In addition to relocating slices across tiers based on relative slice temperature, FAST VP can now also relocate slices within a tier to achieve maximum pool performance gain. Some disks within a tier may be heavily used while other disks in the same tier may be underused. To improve performance, data slices may be relocated within a tier to balance the load. This is accomplished by augmenting FAST VP data relocation to also analyze and move data within a storage tier. This new capability is referred to as load balancing, and occurs within the standard relocation window discussed below.

Managing FAST VP at the storage pool level

You view and manage FAST VP properties at the pool level. [Figure 2](#) shows the tiering information for a specific pool.

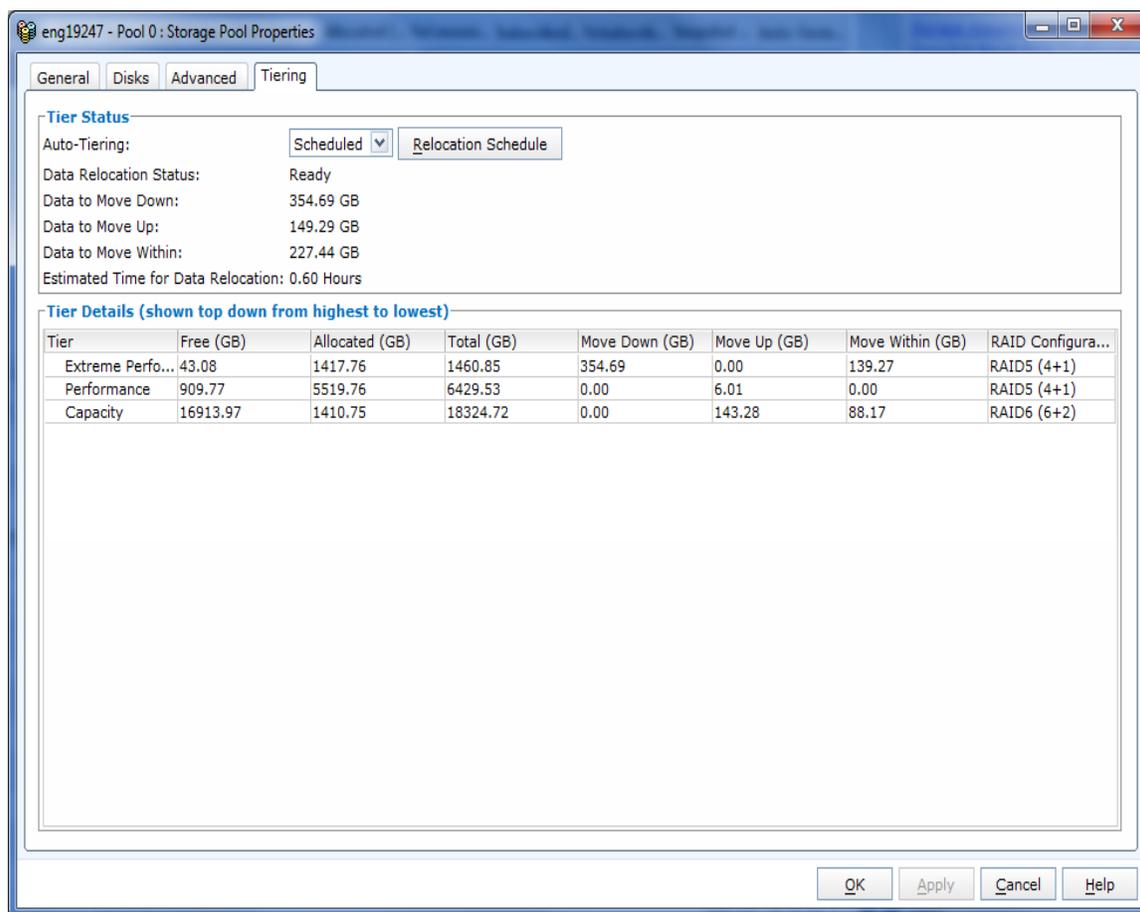


Figure 2. Storage Pool Properties window

The Tier Status section of the window shows FAST VP relocation information specific to the pool selected. For each pool, the Auto-Tiering option can be set to either Scheduled or Manual. Users can also connect to the array-wide relocation schedule using the Relocation Schedule button, which is discussed in the [Automated scheduler](#) section. Data Relocation Status displays what state the pool is in with regards to FAST VP. The Move Down and Move Up figures represent the amount of data that will be relocated in the next scheduled window, followed by the total amount of time needed to complete the relocation. “Data to Move Within” displays the total amount of data to be relocated within a tier.

The Tier Details section displays the data distribution per tier. This panel shows all tiers of storage residing in the pool. Each tier then displays the free, allocated, and total capacities; the amount of data to be moved down and up; the amount of data to move within a tier; and RAID Configuration per tier.

Per-Tier RAID Configuration

Another new enhancement in the EMC® VNX™ Operating Environment (OE) for block release 5.32, file version 7.1 is the ability to select RAID protection according to tier. During pool creation you can choose the appropriate RAID type according to drive type. Keep in mind that each tier has a single RAID type and once the RAID configuration is set for that tier in the pool, you cannot change it.

Figure 3. Per-Tier RAID Selection

Another RAID enhancement coming in this release is the option for more efficient RAID configurations. Users have the following options in pools (new options noted with an asterisk):

Table 2. RAID Configuration Options

RAID Type	Preferred Drive Count Options
RAID 1/0	4+4
RAID 5	4+1, 8+1*
RAID 6	6+2, 14+2*

RAID 5 (8+1) and RAID 6 (14+2) provide 50% savings over the current options, because of the higher data:parity ratio. The tradeoff for higher data:parity ratio translates into larger fault domains and potentially longer rebuild times. This is especially true for RAID 5, with only a single parity drive. Users are advised to choose carefully between (4+1) and (8+1), according to whether robustness or efficiency is a

higher priority. For RAID 6, with 2 parity drives, robustness is less likely to be an issue.

For best practice recommendations, refer to the *EMC Unified Storage Best Practices for Performance and Availability –Common Platform and Block Storage* white paper on EMC Online Support.

Automated scheduler

The scheduler launched from the Pool Properties dialog box's Relocation Schedule button is shown in [Figure 4](#). You can schedule relocations to occur automatically. EMC recommends to set the relocation window either at high rate and run shorter windows (typically off-peak hours), or low rate and run during production time to minimize any potential performance impact that the relocations may cause.

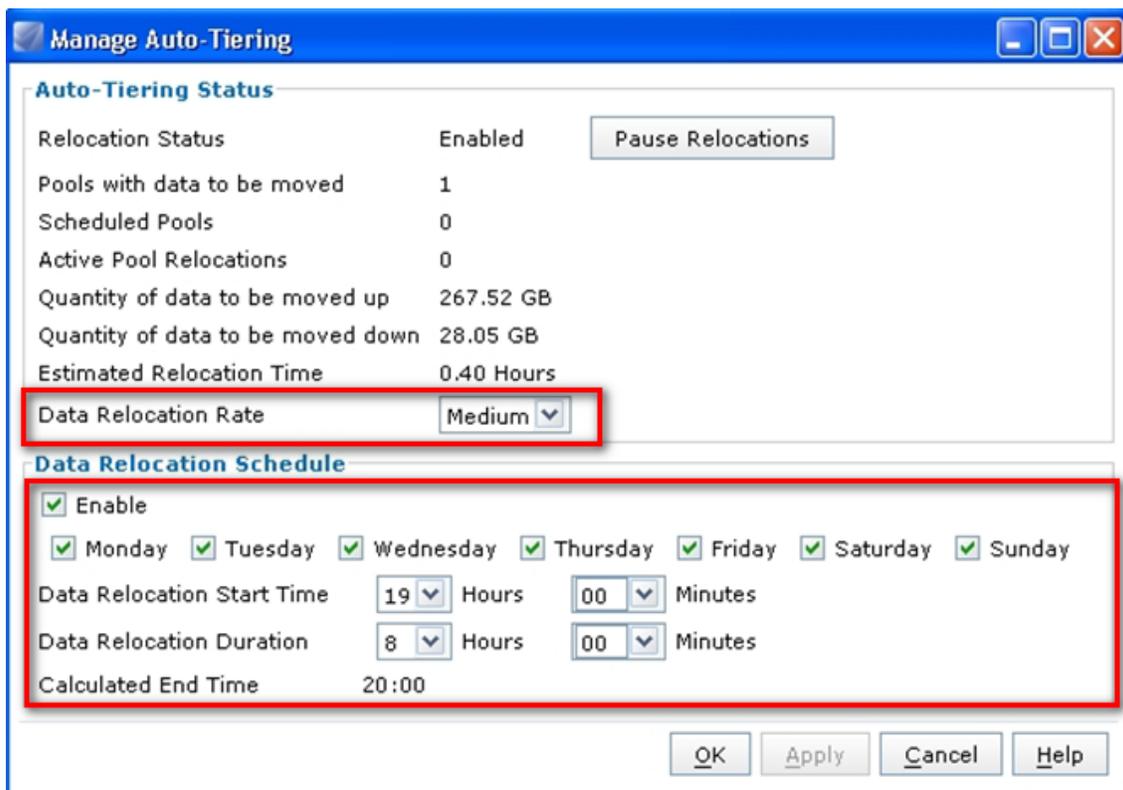


Figure 4. Manage Auto-Tiering window

The Data Relocation Schedule shown in [Figure 4](#) initiates relocations every 24 hours for a duration of eight hours. You can select the days, the start time and the duration on which the relocation schedule should run.

In this example, relocations run seven days a week, which is the default setting. From this status window, you can also control the data relocation rate. The default rate is

set to Medium in order to avoid significant impact to host I/O. This rate relocates data up to approximately 400 GB per hour³.

If all of the data from the latest analysis is relocated within the window, it is possible for the next iteration of analysis and relocation to occur. Any data not relocated during the window is simply re-evaluated as normal with each subsequent analysis. If relocation cycles are incomplete when the window closes, that relocation plan is discontinued, and a new plan is generated when the next relocation window opens.

Manual Relocation

Unlike automatic scheduling, manual relocation is user-initiated. FAST VP performs analysis on all statistics gathered, independent of its default hourly analysis and prior to beginning the relocation.

Although the automatic scheduler is an array-wide setting, manual relocation is enacted at the pool level only. Common situations when users may want to initiate a manual relocation on a specific pool include:

- When new LUNs are added to the Pool and the new priority structure needs to be realized immediately
- When adding a new tier to a pool.
- As part of a script for a finer-grained relocation schedule

FAST VP LUN management

Some FAST VP properties are managed at the LUN level. [Figure 5](#) shows the tiering information for a single LUN.

³ This rate depends on system type, array utilization, and other tasks competing for array resources. High utilization rates may reduce this relocation rate.

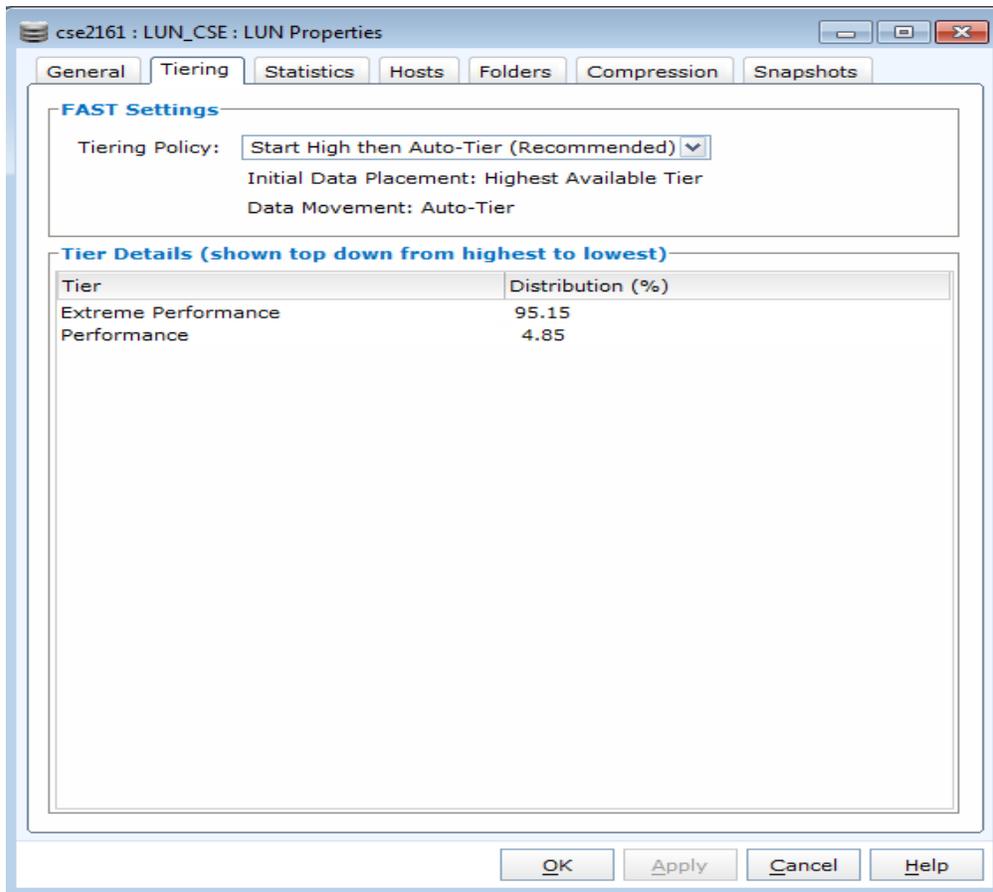


Figure 5. LUN Properties window

The Tier Details section displays the current distribution of slices within the LUN. The Tiering Policy section displays the available options for tiering policy.

Tiering policies

As its name implies, FAST VP is a completely automated feature and implements a set of user defined policies to ensure it is working to meet the data service levels required for the business. FAST VP tiering policies determine how new allocations and ongoing relocations should apply to individual LUNs in a storage pool by using the following options:

- Start High then Auto-tier (New default Policy)
- Highest available tier
- Auto-tier
- Lowest available tier
- No data movement

Start High then Auto-tier (New default policy)

Start High then Auto-tier is the default setting for all pool LUNs upon their creation. Initial data placement is on the highest available tier and then data movement is subsequently based on the activity level of the data. This tiering policy maximizes initial performance and takes full advantage of the most expensive and fastest drives first, while providing subsequent TCO by allowing less active data to be tiered down, making room for more active data in the highest tier.

When there is a pool with multiple tiers, the Start High then Auto-tier design is capable of relocation data to the highest available tier regardless of the drive type combination Flash, SAS or NL SAS. Also when adding a new tier to a pool the tiering policy remains the same and there is no need to manually change it.

Highest available tier

The highest available tier setting should be selected for those LUNs which, although not always the most active, require high levels of performance whenever they are accessed. FAST VP will prioritize slices of a LUN with *highest available tier* selected above all other settings.

Slices of LUNs set to highest tier are rank ordered with each other according to activity. Therefore, in cases where the sum total of LUN capacity set to highest tier is greater than the capacity of the pool's highest tier, the busiest slices occupy that capacity.

Auto-tier

FAST VP relocates slices of LUNs based solely on their activity level after all slices with the highest/lowest available tier have been relocated. LUNS specified with the highest tier will have precedence over LUNS set to Auto-tier.

Lowest available tier

Lowest available tier should be selected for LUNs that are not performance-sensitive or response-time sensitive. FAST VP will maintain slices of these LUNs on the lowest storage tier available regardless of activity level.

No data movement

No data movement may only be selected after a LUN has been created. Once the no data movement selection is made, FAST VP still relocates within a tier, but does not move LUNs up or down from their current position. Statistics are still collected on these slices for use if and when the tiering policy is changed.

The tiering policy chosen also affects the initial placement of a LUN's slices within the available tiers. Initial placement with the pool set to auto-tier results in the data being distributed across all storage tiers available within the pool. The distribution is based

on available capacity in the pool. If 70 percent of a pool's free capacity resides in the lowest tier, then 70 percent of the new slices is most likely placed in that tier.

LUNs set to highest available tier will have their slices placed on the highest tier that has capacity available. LUNs set to lowest available tier will have their slices placed on the lowest tier that has capacity available.

LUNs with the tiering policy set to no data movement will use the initial placement policy of the setting preceding the change to no data movement. For example, a LUN that was previously set to highest tier but is currently set to no data movement still takes its initial allocations from the highest tier possible.

Common uses for Highest Available Tiering policy

Sensitive applications with relative less frequent access

When a pool consists of LUNs with stringent response time demands and relatively less frequent data access, it is not uncommon for users to set certain LUNs in the pool to *highest available tier*. That way, the data is assured of remaining on the highest tier when it is subsequently accessed.

For example if an office has a very important report that is being accessed only once a week, such as every Monday morning, and it contains extremely important information that affects the total production of the office. In this case, you want to ensure the highest performance available even when data is not hot enough to be promoted due to its infrequent access.

Large scale migrations

The general recommendation is to simply turn off tiering until the migration completes. Only if it's critical that data is being tiered during or immediately following the migration, would the recommendation of using the Highest Available tiering policy apply.

When you start the migration process, it is best to fill the highest tiers of the pool first. This is especially important for live migrations. Using the Auto-tier setting would place some data in the Capacity tier. At this point, FAST VP has not yet run an analysis on the new data, so it cannot distinguish between hot and cold data. Therefore, with the Auto-tier setting, some of the busiest data may be placed in the Capacity tier.

In these cases, the target pool LUNs can be set to highest tier. That way, all data is initially allocated to the highest tiers in the pool. As the higher tiers fill and capacity from the Capacity (NL-SAS) tier starts to be allocated, you can stop the migration and run a manual FAST VP relocation.

Assuming an analysis has had sufficient time to run, relocation will order the slices by rank and move data appropriately. In addition, since the relocation will attempt to free 10 percent of the highest tiers, there is more capacity for new slice allocations in those tiers.

You continue this iterative process while you migrate more data to the pool, and then run FAST VP relocations when most of the new data is being allocated to the Capacity tier. Once all of the data is migrated into the pool, you can make any tiering preferences that you see fit.

Using FAST VP for File

In the VNX Operating Environment for File 7, file data is supported on LUNs created in pools with FAST VP configured on both the VNX Unified and VNX gateways with EMC Symmetrix® systems.

Management

The process for implementing FAST VP for File begins by provisioning LUNs from a storage pool with mixed tiers that are placed in the File Storage Group. Rescanning the storage systems from the System tab in the Unisphere software starts a diskmark operation that makes the LUNs available to VNX for File storage. The rescan automatically creates a pool for file using the same name as the corresponding pool for block. Additionally, it creates a disk volume in a 1:1 mapping for each LUN that was added to the File Storage Group. A file system can then be created from the pool for file on the disk volumes. The FAST VP policy that was applied to the LUNs presented to the VNX for File will operate as it does for any other LUN in the system, dynamically migrating data between storage tiers in the pool.

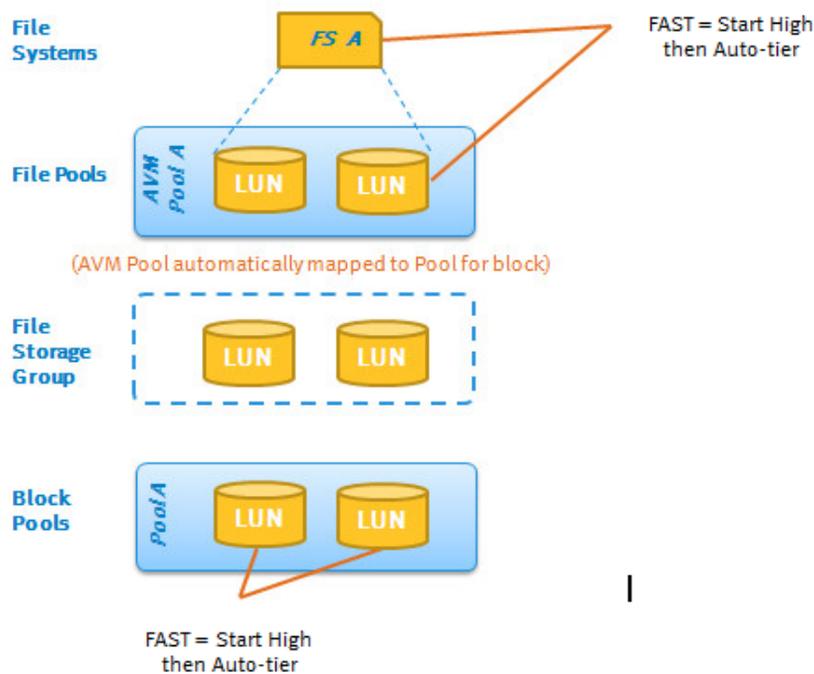


Figure 6. FAST VP for File

FAST VP for File is supported in the Unisphere software and the CLI, beginning with Release 31. All relevant Unisphere configuration wizards support a FAST VP configuration except for the VNX Provisioning Wizard. FAST VP properties can be seen within the properties pages of pools for File (see [Figure 7](#)), and property pages for volumes and file systems (see [Figure 8](#)). They can only be modified through the block pool or LUN areas of the Unisphere software. On the File System Properties page, the FAST VP tiering policy is listed in the Advanced Data Services section along with thin, compression, or mirrored if enabled. If Advanced Data Services is not enabled, the FAST VP tiering policy does not appear. For more information on the thin and compression features, refer to the *EMC VNX Virtual Provisioning* and *EMC VNX Deduplication and Compression* white papers.

Disk type options of mapped disk volumes are:

- LUNs in a storage pool with a single disk type
 - Extreme Performance (Flash drives)
 - Performance (10K and 15k rpm SAS drives)
 - Capacity (7.2K rpm NL-SAS)
- LUNs in a storage pool with multiple disk types (used with FAST VP)
 - Mixed
- LUNs that are mirrored (mirrored means remote mirrored through MirrorView™ or RecoverPoint)
 - Mirrored_mixed
 - Mirrored_performance
 - Mirrored_capacity
 - Mirrored_Extreme Performance

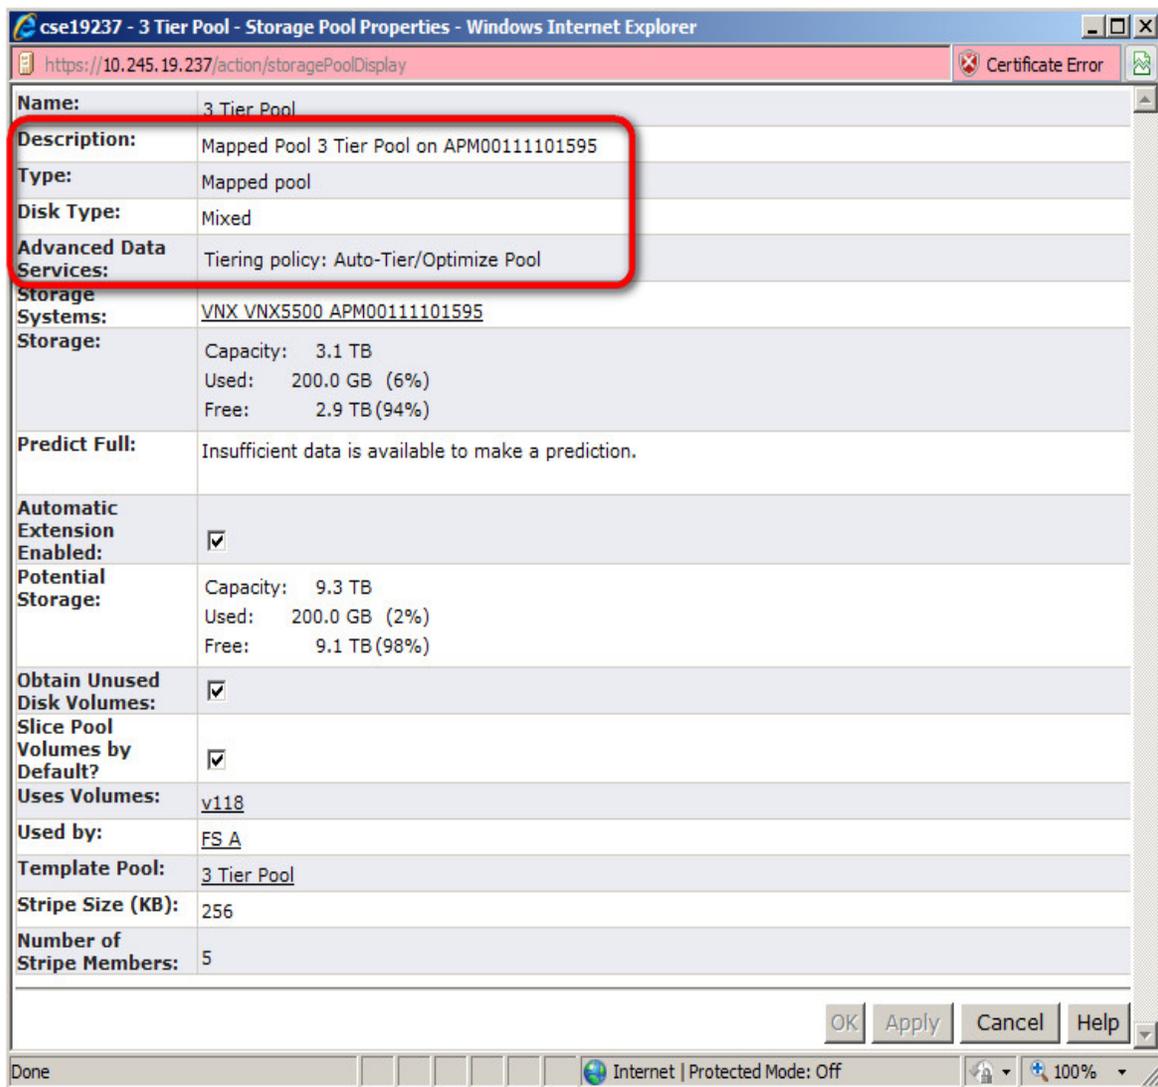


Figure 7. File Storage Pool Properties window

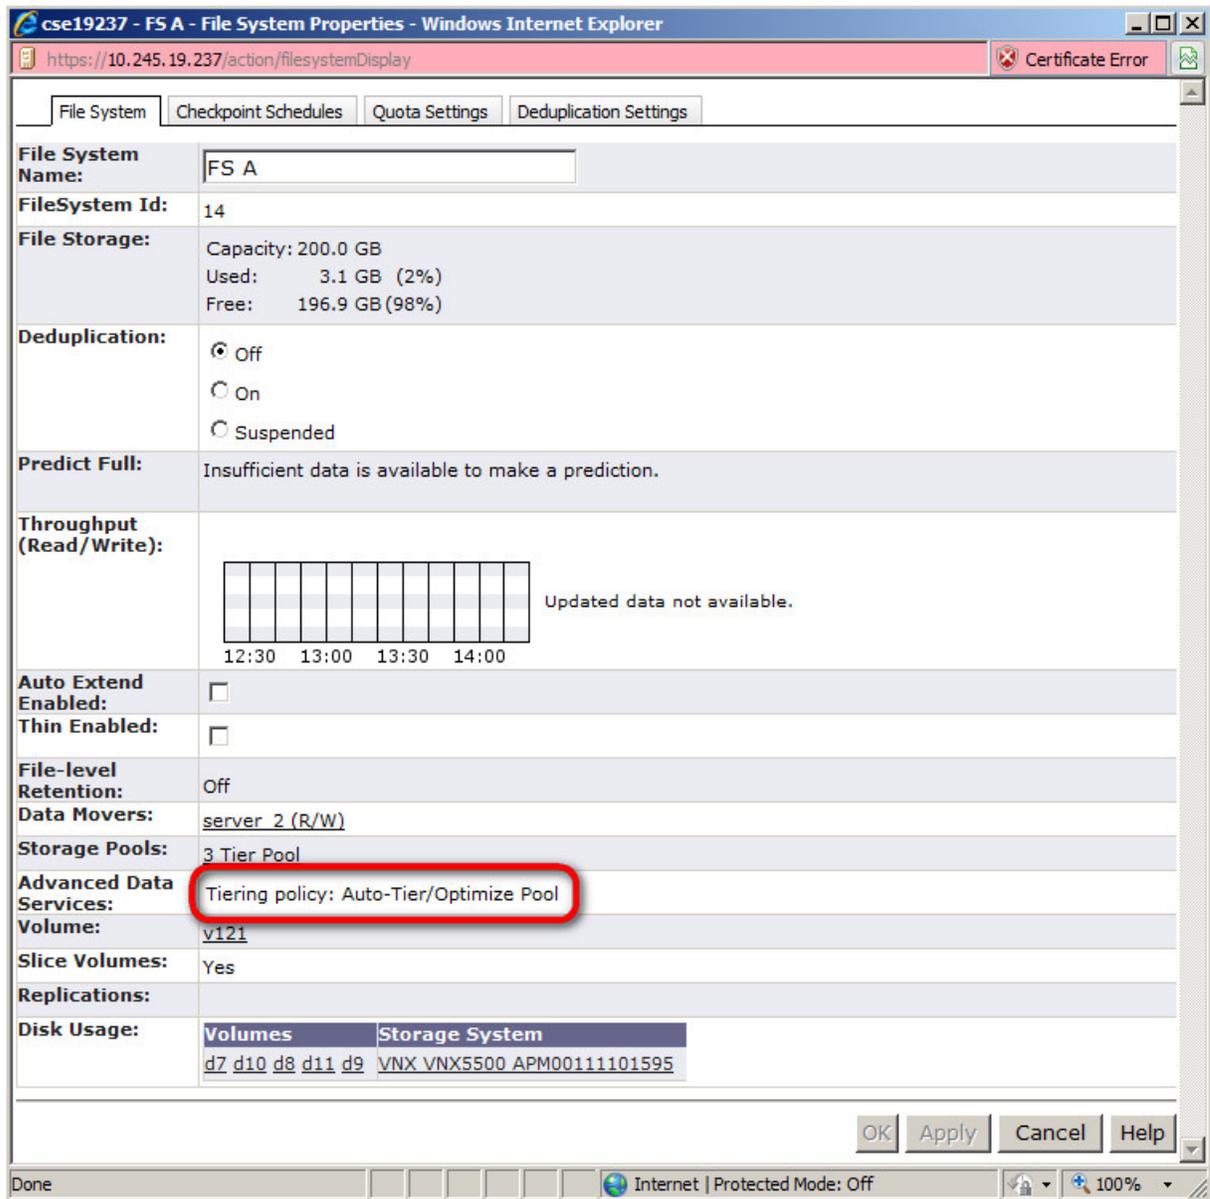


Figure 8. File System Properties window

Best practices for VNX for File

- The entire pool should be allocated to file systems.
- The entire pool should use thick LUNs only.
- Recommendations for thick LUNs:
 - 1 thick LUN per every 4 physical drives.
 - Minimum of 10 LUNs per pool and even multiples of 10 to facilitate striping and SP balancing.
 - Balance SP ownership.

- All LUNs in the pool should have the same tiering policy. This is necessary to support slice volumes.
- EMC does not recommend that you use block thin provisioning or compression on VNX LUNs used for file system allocations. Instead, you should use the File side thin provisioning and File side deduplication (which includes compression).

Note: VNX file configurations will not allow mixing of mirrored and non-mirrored types in a pool. If you try to do this, the disk mark will fail.

- Where the highest levels of performance are required, maintain the use of RAID Group LUNs for File.

General Guidance and recommendations

FAST VP and FAST Cache

FAST Cache and FAST VP are part of the FAST suite. These two products are sold together, work together and complement each other. FAST Cache allows the storage system to provide Flash drive class performance to the most heavily accessed chunks of data across the entire system. FAST Cache absorbs I/O bursts from applications, thereby reducing the load on back-end hard disks. This improves the performance of the storage solution. For more details on this feature, refer to the *EMC CLARiiON, Celerra Unified, and VNX FAST Cache white paper* available on EMC Online Support.

Table 2. Comparison between the FAST VP and FAST Cache features

FAST Cache	FAST VP
Enables Flash drives to be used to extend the existing caching capacity of the storage system.	Leverages pools to provide sub-LUN tiering, enabling the utilization of multiple tiers of storage simultaneously.
Has finer granularity – 64 KB.	Less granular compared to FAST Cache – 1 GB.
Copies data from HDDs to Flash drives when they are accessed frequently.	Moves data between different storage tiers based on a weighted average of access statistics collected over a period of time.
Adapts continuously to changes in workload.	Uses a relocation process to periodically make storage tiering adjustments. Default setting is one 8-hour relocation per day.
Is designed primarily to improve performance.	While it can improve performance, it is primarily designed to improve ease of use and reduce TCO.

You can use the FAST Cache and the FAST VP sub-LUN tiering features together to yield high performance and improved TCO for the storage system. As an example, in scenarios where limited Flash drives are available, the Flash drives can be used to create FAST Cache, and the FAST VP can be used on a two-tier, Performance and Capacity pool. From a performance point of view, FAST Cache dynamically provides performance benefits to any bursty data while FAST VP moves warmer data to Performance drives and colder data to Capacity drives. From a TCO perspective, FAST

Cache with a small number of Flash drives serves the data that is accessed most frequently, while FAST VP can optimize disk utilization and efficiency.

As a general rule, use FAST Cache in cases where storage system performance needs to be improved immediately for burst-prone data. Use FAST VP to optimize storage system TCO, moving data to the appropriate storage tier based on sustained data access and demands over time. FAST Cache focuses on improving performance while FAST VP focuses on improving TCO. Both features are complementary to each other and help in improving performance and TCO.

The FAST Cache feature is storage-tier-aware and works with FAST VP to make sure that the storage system resources are not wasted by unnecessarily copying data to FAST Cache if it is already on a Flash drive. If FAST VP moves a chunk of data to the Extreme Performance Tier (which consists of Flash drives), FAST Cache will not promote that chunk of data into FAST Cache, even if FAST Cache criteria is met for promotion. This ensures that the storage system resources are not wasted in copying data from one Flash drive to another.

A general recommendation for the initial deployment of Flash drives in a storage system is to use them for FAST Cache. In almost all cases, FAST Cache with a 64 K granularity offers the industry's best optimization of flash.

What drive mix is right for my I/O profile?

As previously mentioned, it is common for a small percentage of overall capacity to be responsible for most of the I/O activity. This is known as skew. Analysis of an I/O profile may indicate that 85 percent of the I/Os to a volume only involve 15 percent of the capacity. The resulting active capacity is called the working set. Software like FAST VP and FAST Cache keep the working set on the highest-performing drives.

It is common for OLTP environments to yield working sets of 20 percent or less of their total capacity. These profiles hit the sweet spot for FAST and FAST Cache. The white papers *Leveraging Fully Automated Storage Tiering (FAST) with Oracle Database Applications* and *EMC Tiered Storage for Microsoft SQL Server 2008—Enabled by EMC Unified Storage and EMC Fully Automated Storage Tiering (FAST)* on the EMC Online Support website discuss performance and TCO benefits for several mixes of drive types.

At a minimum, the capacity across the Performance Tier and Extreme Performance Tier (and/or FAST Cache) should accommodate the working set. However, capacity is not the only consideration. The spindle count of these tiers needs to be sized to handle the I/O load of the working set. Basic techniques for sizing disk layouts based on IOPS and bandwidth are available in the *EMC VNX Fundamentals for Performance and Availability white paper* on the EMC Online Support website.

Performance Tier drives are versatile in handling a wide spectrum of I/O profiles. Therefore, EMC highly recommends that you include Performance Tier drives in each pool. FAST Cache can be an effective tool for handling a large percentage of activity, but inevitably, there will be I/Os that have not been promoted or are cache misses. In these cases, Performance Tier drives offer good performance for those I/Os.

Performance Tier drives also facilitate faster promotion of data into FAST Cache by quickly providing promoted 64 KB chunks to FAST Cache. This minimizes FAST Cache warm-up time as some data gets hot and other data goes cold. Lastly, if FAST Cache is ever in a degraded state due to a faulty drive, the FAST Cache will become read only. If the I/O profile has a significant component of random writes, these are best served from Performance Tier drives as opposed to Capacity drives.

Capacity drives can be used to optimize TCO. This often equates to comprising 60 percent to 80 percent of the pool's capacity. Of course, there are profiles with low IOPS/GB and or sequential workload that may result in the use of a higher percentage of Capacity Tier drives.

EMC Professional Services and qualified partners can be engaged to assist with properly sizing tiers and pools to maximize investment. They have the tools and expertise to make very specific recommendations for tier composition based on an existing I/O profile.

Conclusion

Through the use of FAST VP, users can remove complexity and management overhead from their environments. FAST VP utilizes Flash, Performance, and Capacity drives (or any combination thereof) within a single pool. LUNs within the pool can then leverage the advantages of each drive type at the 1 GB slice granularity. This sub-LUN-level tiering ensures that the most active dataset resides on the best-performing drive tier available, while maintaining infrequently used data on lower-cost, high-capacity drives.

Relocations can occur without user interaction on a predetermined schedule, making FAST VP a truly automated offering. In the event that relocation is required on-demand, you can invoke FAST VP relocation on an individual pool by using the Unisphere software.

Both FAST VP and FAST Cache work by placing data segments on the most appropriate storage tier based on their usage pattern. These two solutions are complementary because they work on different granularity levels and time tables. Implementing both FAST VP and FAST Cache can significantly improve performance and reduce cost in the environment.

References

The following white papers are available on the EMC Online Support website:

- *EMC Unified Storage Best Practices for Performance and Availability – Common Platform and Block – Applied Best Practices*
- *EMC VNX Virtual Provisioning*
- *EMC Storage System Fundamentals for Performance and Availability*

- *EMC CLARiiON, Celerra Unified, and VNX FAST Cache*
- *EMC Unisphere: Unified Storage Management Solution*
- *An Introduction to EMC CLARiiON and Celerra Unified Platform Storage Device Technology*
- *EMC Tiered Storage for Microsoft SQL Server 2008—Enabled by EMC Unified Storage and EMC Fully Automated Storage Tiering (FAST)*
- *Leveraging Fully Automated Storage Tiering (FAST) with Oracle Database Applications*