

EMC UNITY: FAST TECHNOLOGY OVERVIEW

Overview

ABSTRACT

This white paper is an introduction to the EMC® FAST™ technology for the EMC Unity™ family of storage systems. It describes the background concepts, major components, and implementation steps for EMC FAST technology, which includes FAST VP and FAST Cache. Guidelines and other useful information such as benefits will also be included.

December, 2016

The information in this publication is provided “as is.” Dell Inc. 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 software described in this publication requires an applicable software license.

Copyright © 2016 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA [12/16] [White Paper] [H15086.2]

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

TABLE OF CONTENTS

EXECUTIVE SUMMARY	4
Audience	4
Terminology.....	4
FAST VP	6
Introduction	6
FAST VP Licensing	6
Using FAST VP	6
FAST VP Tiering Policies	8
The FAST VP Algorithm	10
Managing FAST VP.....	12
FAST CACHE	18
Introduction	18
FAST Cache Licensing.....	18
FAST Cache Components	18
Theory of Operation	18
FAST Cache Management.....	22
Supported Drives and Configurations.....	29
Failure Handling	29
COMPARISON OF FAST CACHE AND FAST VP	29
Interoperability	30
CONCLUSION	31
REFERENCES.....	33

EXECUTIVE SUMMARY

EMC FAST™ software allows the EMC Unity™ product family to leverage high performance Flash drives. FAST software consists of Fully Automated Storage Tiering for Virtual Pools (FAST VP) and FAST Cache. These two features work in tandem to utilize the storage within the system as efficiently as possible. Each of these software features ensures that the most active data is serviced from Flash.

When FAST VP is enabled, the FAST VP software will measure and record performance statistics on each slice within a Pool. Later, FAST VP analyzes this data and makes decisions to move data across multiple tiers in a Pool to maximize Pool performance and efficiently utilize the space within the Pool. Slices that are highly accessed are automatically moved to the higher tiers within a Pool, while slices with less activity move to lower tiers within a Pool. Data already residing on Flash within a Pool will not promote to FAST Cache, allowing more data within the system to take advantage of Flash.

FAST Cache is a high capacity secondary cache which sits below System Cache, and above the drives in the system. FAST Cache complements the System Cache by utilizing Flash drives to service highly active data being requested from drive. Frequently accessed data located on spinning drives in the Pool is copied into FAST Cache for higher performance and lower response time. As FAST Cache copies data onto Flash in 64 KB chunks, FAST Cache is very efficient at utilizing Flash. It is not uncommon for an entire active data set within a small capacity to reside in FAST Cache. With large maximum capacities on the higher end Unity models, FAST Cache can be configured to handle most of the I/O on a system.

FAST VP and FAST Cache are licensed features available on EMC Unity.

AUDIENCE

This white paper is intended for EMC customers, partners, and employees who are considering the use of the FAST VP and FAST Cache features in the EMC Unity family of storage systems. It assumes familiarity with EMC Unity and EMC's management software.

TERMINOLOGY

Chunk – A piece of data located within a particular 64 KB region.

DRAM memory – DRAM based memory used by the storage system to store data.

FAST Cache clean page – A 64 KB FAST Cache page which is in use and contains data that is an exact copy of the corresponding page within the storage resource.

FAST Cache copy – The process of copying a FAST Cache dirty page to the corresponding storage resource.

FAST Cache dirty page – A 64 KB FAST Cache page which is in use and contains a newer copy of the data than the corresponding page within the storage resource. These pages will be synchronized when the FAST Cache page is cleaned.

FAST Cache flush – The process of freeing a FAST Cache page for a new promotion by first copying the contents of a FAST Cache page to its storage resource, then freeing the page.

FAST Cache hit – The instance when the data being requested or updated is contained within FAST Cache.

FAST Cache miss – The instance when the data being requested or updated is not contained within FAST Cache.

Fast Cache page – A single allocation unit (page) located within FAST Cache used to store data. This page is 64 KB in size.

FAST Cache promotion – The process of copying data from a storage resource into FAST Cache.

Flash drive (Solid State Drive - SSD) – A Flash based storage device used to store data.

Hard Disk Drive (HDD) – A storage device based on spinning platters used to store data.

Locality of reference – Used when describing a data access pattern where adjacent blocks of a data set are accessed frequently.

Logical Block Address (LBA) – The addressing scheme used when accessing a particular block of data within a storage resource.

Memory map – A FAST Cache component which tracks the current contents and locations of data within FAST Cache. A copy of the memory map is stored within system memory and on drive.

Pool – A set of drives that provide specific storage characteristics for the resources that use them, such as LUNs, VMware Datastores, and File Systems.

Rebalance – A FAST VP process automatically started when unused drives are used to increase the capacity a tier of storage within a Pool. The process relocates data to correct the imbalance of data placement across the tier.

Slice – A 256 MB unit of capacity which is allocated to storage resources to provide space to store data.

Slice Relocation – A physical movement of a 256 MB slice of data within a tier or across tiers within a Pool.

Storage Resource – A LUN, VMware Datastore, VVol, or File System.

System Cache (DRAM Cache) – EMC Unity software component which leverages DRAM memory to improve host read and write performance.

System Cache hit – The instance when a host I/O can be serviced with the contents of System Cache.

System Cache miss – The instance when a host I/O cannot be serviced with the contents of System Cache.

Temperature – The weighted average of a 256 MB Slice's activity level over time.

Tier – A label used to describe the various categories of media used within a Pool. In a physical system, the tier directly relates to the drive types used within the Pool. The available tiers are the Extreme Performance Tier (Flash drives), the Performance Tier (SAS drives), and the Capacity Tier (NL-SAS drives). For the virtual Unity system, the storage tier of a Virtual Drive must be entered manually and should be chosen to match the underlying characteristics of the Virtual Disk.

FAST VP

INTRODUCTION

When reviewing the access patterns for data within a system, most access patterns show a basic trend. Typically the data is most heavily accessed near the time it was created, and the activity level decreases as the data ages. This trending is also referred to as the lifecycle of the data. EMC Unity Fully Automated Storage Tiering for Virtual Pools (FAST VP) monitors the data access patterns within Pools on the system, and dynamically matches the performance requirements of the data with drives that provide that level of performance. FAST VP classifies drives into three categories, called tiers. These tiers are:

- **Extreme Performance Tier** – Comprised of Flash drives
- **Performance Tier** – Comprised of Serial Attached SCSI (SAS) drives
- **Capacity Tier** - Comprised of Near-Line SAS (NL-SAS) drives

FAST VP helps to reduce the Total Cost of Ownership (TCO) by maintaining performance while efficiently utilizing the configuration of a Pool. Instead of creating a Pool with one type of drive, mixing Flash, SAS, and NL-SAS drives can help reduce the cost of a configuration by reducing drive counts and leveraging larger capacity drives. Data requiring the highest level of performance is tiered to Flash, while data with less activity resides on SAS or NL-SAS drives.

EMC Unity has a unified approach to create storage resources on the system. Block LUNs, File Systems, and VMware Datastores can all exist within a single Pool, and can all benefit from using FAST VP. In system configurations with minimal amounts of Flash, FAST VP will efficiently utilize the Flash drives for active data, regardless of the resource type. For efficiency, FAST VP also leverages low cost spinning drives for less active data. The access patterns for all data within a Pool are compared against each other, and the most active data is placed on the highest performing drives while adhering to the storage resource's tiering policy. Tiering policies are explained later in this document.

FAST VP LICENSING

In Unity, FAST VP is supported on Unity Hybrid systems and the UnityVSA. For EMC Unity Hybrid systems, FAST VP is included with the EMC Unity Essentials Software package, which is included with all Unity Hybrid systems. For UnityVSA, FAST VP is included with the base software license. When these licenses are installed, the following FAST VP related features are available:

- The ability to create a Pool with multiple drive types
- The ability to set tiering policies on Block LUNs, File Systems, and VMware Datastores.
- The ability to access the FAST VP tab within the Pool properties window or a storage resource properties window

USING FAST VP

POOLS

A Pool is a set of drives on which storage resources are created. Pools can be created on all EMC Unity systems, including a Unity All Flash system, Unity Hybrid system, or the UnityVSA. In the Unity All Flash and Hybrid systems, Pools are comprised of physical drives found within the system. In UnityVSA, Pools are created on Virtual Disks, which have been provided from the VMware ESXi host the UnityVSA is deployed on. Pools can contain just a few drives, or contain all drives within a system. FAST VP is the software feature which helps to efficiently utilize the resources found within a Pool. A Pool can either contain a single drive type, or contain a mix of drive types.

A Pool which only contains a single drive type is called a single tier Pool, also known as a Homogenous Pool. A single tier Pool can contain either all Flash drives, all SAS drives, or all NL-SAS drives. Single tier Pools provide predictable performance, as all drives within the Pool are of the same type. All data contained within a single tier Pool have the same performance potential, regardless of the data's age. Single tier Pools are best utilized when data access is uniform across large address ranges.

A Pool which contains a mixture of drive types is called a multi-tiered Pool, also known as a Heterogeneous Pool or Hybrid Pool. A multi-tiered Pool can contain any combination of Flash, SAS, and NL-SAS drives. For example, a Pool may contain all three drive types, or may contain only Flash and NL-SAS drives, or only SAS and NL-SAS drives. In a multi-tiered Pool, data from storage resources are spread across the tiers by FAST VP in 256 MB slices. FAST VP monitors the usage of each tier, and relocates data slices within the Pool based on the temperature of the slices and capacity.

FAST VP TIERS

When a Pool is created, each drive type within the Pool is assigned to a specific tier. Each tier provides a different level of performance due to the drive type and number of drives within the tier. Each tier within a Pool can also contain a different capacity of storage, allowing you to customize how much storage of each drive type is allocated to the Pool. FAST VP has three categories in which it segregates different drives into. These tiers are:

- **Extreme Performance Tier** – Comprised of Flash drives
- **Performance Tier** – Comprised of Serial Attached SCSI (SAS) drives
- **Capacity Tier** - Comprised of Near-Line SAS (NL-SAS) drives

FAST VP differentiates each of these tiers by drive type, and not rotational speed. EMC suggests not mixing drives with different rotational speeds within a tier of a Pool. For example, do not mix 10K RPM and 15K RPM SAS drives within the same Pool. It is suggested to allocate these drives to different Pools. It is OK however to mix 10K RPM SAS drives in a Pool with 7.2K RPM NL-SAS drives as they are different drive types and will exist in different tiers.

FAST VP leverages all tiers within a Pool, as each tier provides unique advantages regarding performance and cost.

POOL TIER RAID CONFIGURATIONS

When creating a Pool, you can customize the RAID protection for each of the tiers being configured on a per-tier basis. A single RAID protection is chosen for each tier, and once the RAID protection is selected and the Pool is created, the RAID protection cannot be changed. When expanding a Pool to include a new drive type, you will be able to select a RAID protection at that time. Table 1 shows the different RAID configurations that are supported for each tier. When creating a Pool, drives must be added in multiples of the chosen RAID configuration.

Table 1. Supported RAID Types and Drive Configurations.

RAID Type	Default Configuration	Supported Configurations
RAID 1/0	4+4	1+1*, 2+2, 3+3, 4+4
RAID 5	4+1	4+1, 8+1, 12+1
RAID 6	6+2	4+2, 6+2, 8+2, 10+2, 12+2, 14+2

*RAID 1/0 (1+1) is RAID 1

When deciding on which RAID Configuration to use, consider the performance, capacity, and protection levels each configuration provides. RAID 1/0 is suggested for applications with large amounts of random writes, as there is no parity write penalty in this RAID type. RAID 5 is preferred when cost and performance are a concern. RAID 6 provides the maximum level of protection against drive faults of all the supported RAID types. When considering a RAID configuration which includes a large number of drives, (12+1, 12+2, 14+2), consider the tradeoffs that the larger drive counts contain, such as the fault domain and potentially long rebuild times.

THE EXTREME PERFORMANCE TIER

The Extreme Performance Tier is most often utilized for storing data that is response time sensitive and requires a high level of performance. Flash drives are used in an Extreme Performance Tier as they have low latency and outperform the other supported drive types on the system. As Flash drives do not use spinning platters, seek latency seen with rotating drives is entirely avoided, which makes Flash a great choice for random access workloads.

When creating a Pool, it is suggested to add Flash drives to the configuration. Even a small amount of Flash capacity added to a Pool can be leveraged by FAST VP to increase the overall performance of the system. For the best return on investment, utilize Flash drives to store hot data on storage resources requiring fast response times and high IOPs. FAST VP will optimize the entire Pool's resources and automatically relocate less active data to other tiers as needed.

The Extreme Performance Tier can be created using SAS Flash 2, SAS Flash 3 or SAS Flash 4 Flash drives. While it is possible to mix different size SAS Flash drives within the same Pool, it is not recommended. In the Unity OE version 4.1 and later, SAS Flash 3 drives can be used in hybrid (mixed drive type) pools. SAS Flash 4 drives are only supported within an All Flash Unity system.

THE PERFORMANCE TIER

The Performance Tier provides high, all-around performance with consistent response times, high I/O, and good bandwidth at a mid-level price point. The Performance Tier contains SAS drives, which provide high levels of performance, reliability, and capacity. SAS drive technology is based on rotating mechanical hard-drive technology which uses magnetic platters to store data. The Performance Tier can be created using 15K RPM or 10K RPM SAS drives. Though it is possible, mixing SAS drives with different speed within a Pool is not suggested.

THE CAPACITY TIER

The Capacity Tier is a tier of storage within a Pool comprised of large capacity, low cost 7.2K RPM NL-SAS drives. By leveraging NL-SAS drives, the Capacity Tier helps to lower the cost per GB of a Pool. The Capacity Tier is utilized by FAST VP to store data which is primarily static, and not accessed often. It is best used for data that does not have a strict performance requirement. As data ages, and the activity level decreases, more and more data considered “cold” by FAST VP will be tiered to the Capacity Tier.

TIER CONSIDERATIONS

While FAST VP handles the placement of data within a Pool based on data access patterns and storage efficiency, it is good to understand the differences between the multiple tier categories. The following table, Table 2, outlines the main differences between the Tiers and the drives that comprise them.

Table 2. Comparison of the Extreme Performance Tier, Performance Tier, and Capacity Tier.

	Extreme Performance (Flash)	Performance (SAS)	Capacity (NL-SAS)
Performance	User Response Time		
	<1 ms – 5 ms	≈5 ms	7 ms – 10 ms
	Multi-Access Response Time		
	< 10 ms	10 ms – 50 ms	≤ 100 ms
	High IOPs / GB Low Latency	High Bandwidth with contending workloads	Low IOPs / GB Leverages System Cache for sequential and large block I/O
Strengths	Extremely fast access for reads	Sequential reads leverage prefetching of data	
	Handles multiple sequential workloads better than SAS or NL-SAS	Sequential writes leverage system optimizations favoring drive	
		Read/write mixes provide predictable performance	Large I/O is serviced efficiently
Other Observations	Writes perform slower than reads	Long response times typically seen for heavy write workloads	
	Heavy concurrent writes affect read performance		
	Single-threaded, large, sequential I/O performance equivalent to SAS		

FAST VP TIERING POLICIES

FAST VP is a software feature which monitors each Pool within the system, and makes tiering choices for data within each of the Pools. Each storage resource created on a Pool also has settings which can impact how FAST VP tiers the data. This setting is called a tiering policy. At time of creation in Unisphere, CLI, or REST API, you can specify a tiering policy for storage resources being created on a Pool. This tiering policy determines which tier the resource’s data will initially be placed in, and how the 256 MB slices of data for the resource will be relocated within the Pool. FAST VP runs on a customizable schedule, or can be manually started at any time. The

tiering of data is based on statistics gathered hourly on the Pool. This will be discussed further in *The FAST VP Algorithm* section of this paper.

TIERING POLICY OPTIONS

FAST VP is an automated feature which utilizes user specified tiering policies to meet performance goals based on the storage resource. Relocations occur on 256 MB blocks of data, called slices, which provide storage to each resource created on the Pool. FAST VP bases relocation decisions on an algorithm which takes into account the activity level of each slice, and the resource's tiering policy. The algorithm then orders the slices based on rank, which will later be used during the relocation window. This process is repeated hourly, and a data movement candidate list is created for all Pools within the system.

The following are the available Tiering Policies for storage resources created on a Pool in EMC Unity:

- Highest Available Tier
- Auto-Tier
- Start High then Auto-Tier (Default/Recommended)
- Lowest Available Tier

All Tiering Policies can be set at time of creation, or changed later.

HIGHEST AVAILABLE TIER

The Highest Available Tier tiering policy is designed to prioritize the placement of slices for the assigned storage resource over others within the Pool. It is suggested to be used when performance and minimal response times are a priority. Using this tiering policy does not require the data to be highly accessed, but helps to provide the best performance when the data is accessed. Slices used for the resources with Highest Available Tier assigned to it will be taken from the highest tier with available space when slices are needed.

If the highest tier in the Pool does not have any space, space from the next tier with available space will be taken. Slices are prioritized in the following manner:

- Existing slices on the highest tier of a Pool have priority over new slices being consumed for storage resources. New slice allocations do not immediately force slices out of the highest tier, regardless of the tiering policy set on the resource. This priority is revisited during the next FAST VP relocation window.
- When multiple resources have the Highest Available Tier policy assigned, and there is not enough space within the top tier to store all data slices, they compete for top tier placement based on each slice's temperature. The temperature is based on the activity level of the slice. Slices for resources with Highest Available Tier assigned always take top tier priority over other tiering policies.

AUTO-TIER

A storage resource tends to contain regions which have higher activity levels than others. To efficiently utilize the tiers within a Pool, FAST VP will move the "hot" slices to the higher tiers, while placing less active data on lower tiers of the Pool. The Auto-Tier policy automatically places slices of data for the storage resource on the various tiers of a Pool based on the data's usage level. Although a slice for a storage resource with Auto-Tier assigned may be more active than a resource with Highest Available Tier assigned, the resource with Highest Available Tier takes precedence. When allocating new slices to a storage resource, slices are taken from all tiers based on the usage of each tier. If a large portion of a Pool's free space resides in the Capacity Tier, slices will be allocated from the Capacity Tier.

START HIGH THEN AUTO-TIER (DEFAULT/RECOMMENDED)

Start High then Auto-Tier is not only default tiering policy when creating a storage resource, but it is also the EMC recommended policy. When allocating slices to a storage resource, this policy assigns slices from the highest tier with available space in the Pool. For existing slices for resources using the Start High then Auto-Tier policy, it uses the Auto-Tier tiering policy. FAST VP relocates the slices based on performance statistics of the slice, and the auto-tiering algorithm. This tiering policy maximizes initial performance by taking slices from the highest tier with available space, while also efficiently utilizing the storage and tiering existing slices with lower activity levels to the lower tiers. This allows other slices to utilize the higher tiers, which helps to reduce the Total Cost of Ownership of the configuration. This tiering policy is recommended when activity on a storage resource has a skew.

LOWEST AVAILABLE TIER

When using this policy, new slices for a resource will be assigned from the lowest tier with free capacity. This policy is recommended for resources that are not performance sensitive or response time sensitive. Regardless of the activity level for the slices assigned to resources with this policy, all slices will remain in the lowest available tiers. If the lowest tier within a Pool does not have free capacity, all slices for resources with this tiering policy assigned to it will be compared. The slices with the lowest activity will have priority to reside on the lowest tier of the Pool.

COMPARING THE TIERING POLICIES

The following table, Table 3, summarizes the various Tiering Policies in EMC Unity.

Table 3. EMC Unity Tiering Policies.

Tiering Policy	Corresponding Initial Tier Placement	Description
Highest Available Tier	Highest Available Tier	Initial data placement and subsequent data relocations set to the highest performing tier of drives with available space
Auto-Tier	Optimized for Pool Performance	Initial data placement optimizes Pool capacity, then relocates slices to different tiers based on the activity levels of the slices
Start High then Auto-Tier (Default)	Highest Available Tier	Initial data placed on slices from the highest tier with available space, then relocates data based on performance statistics and slice activity
Lowest Available Tier	Lowest Available Tier	Initial data placement and subsequent relocations preferred on the lowest tier with available space.

THE FAST VP ALGORITHM

The FAST VP algorithm uses multiple strategies to improve the overall performance of the system, efficiently utilize the capacity of each Pool, and reduce the Total Cost of Ownership. These processes identify and relocate slices to appropriate tiers based on statistics collected on each slice, analysis of the data, and the relocation operation based on the activity levels of each slice. Each slice contains 256 MBs of data, which is the granularity of FAST VP.

FAST VP also attempts to free space within each tier to allow for new slice allocations or slice promotions. Leaving 10% free capacity in each tier allows FAST VP to be more efficient when tiering slices to higher tiers. If needed, during a relocation window least recently used slices within the tier will be tiered to lower tiers to reach the 10% free capacity target.

STATISTICS COLLECTION

The activity level of a slice within a Pool is often referred to as the temperature of the slice, and terms “Hot” and “Cold” are often used to describe how active the slices are. When comparing multiple slices, a slice which has more activity than another is considered “Hotter”, and the one with less activity is considered “Colder”. The activity of a slice is determined by tracking the amount of I/O that is sent to each slice, which includes both reads and writes. FAST VP keeps these statistics and “weighs” the I/O based on the time of arrival. Recent activity on a slice receives a higher weight, and the weight deteriorates over time. The slice statistics are collected continuously on the system for all storage resources.

ANALYSIS

Once an hour, FAST VP analyzes the data collected and ranks each slice, based on each slice’s temperature. This list is ordered from “Hottest” to “Coldest”, and lists are created for each Pool within the system. Based on this list, a relocation candidate list is compiled with information regarding which slices should be moved up, moved down, or moved within a tier in a Pool. This candidate list also takes into account each storage resource’s tiering policy, to ensure each policy is being followed. The next time a relocation is started, either via the schedule or manually, the latest candidate list is used. You can influence the candidate list by changing the tiering policies on storage resources, as the tiering policy takes precedence over the activity levels of slices.

RELOCATION

When a relocation window starts, either by the schedule or manually, FAST VP begins promoting or demoting slices according to the candidate list created in the analysis phase. The hottest slices are moved to the higher tiers, and colder slices are moved down to the lower tiers. During the relocation window, priority is given to slices moving to the higher tiers, as they will benefit most from the relocation. For slices relocating to a lower tier, relocations only occur when slices being promoted need the space it occupies. By leveraging the space available, FAST VP ensures that top tier drives are utilized.

Storage resources with a tiering policy of Lowest Available Tier may also relocate during a FAST VP relocation window. If slices for these resources do not already reside on the lowest tier, and space becomes available on the lowest tier, relocations may occur. Only the “Coldest” slices for storage resources with the Lowest Available Tier policy will be stored on the lowest tier if not enough capacity exists for all slices.

Another factor for relocations is tier capacity. FAST VP will also review the capacity of a tier to make relocation decisions. If a tier has less than 10% of free space, “Cold” slices will be tiered down to free enough space to reach the 10% mark. Leaving free capacity in each tier allows storage resources to allocate slices efficiently based on their tiering policies. The free capacity is also utilized when slices are relocated into higher tiers when the relocation window starts. Using free space within a tier is more efficient on the system than relocating slices from a tier before relocating slices into a tier.

Figure 1 shows an illustration of how FAST VP can improve the performance of the Pool by relocating slices. On the left is a Storage Pool before FAST VP relocations have occurred. Notice that slices across each of the tiers have different levels of activity. After analyzing the activity on these slices, FAST VP will determine the best placement for the data within the Pool. The right side of Figure 1 shows the Pool after relocations have occurred. Notice that activity levels have been corrected and slices have been placed on the appropriate tiers.

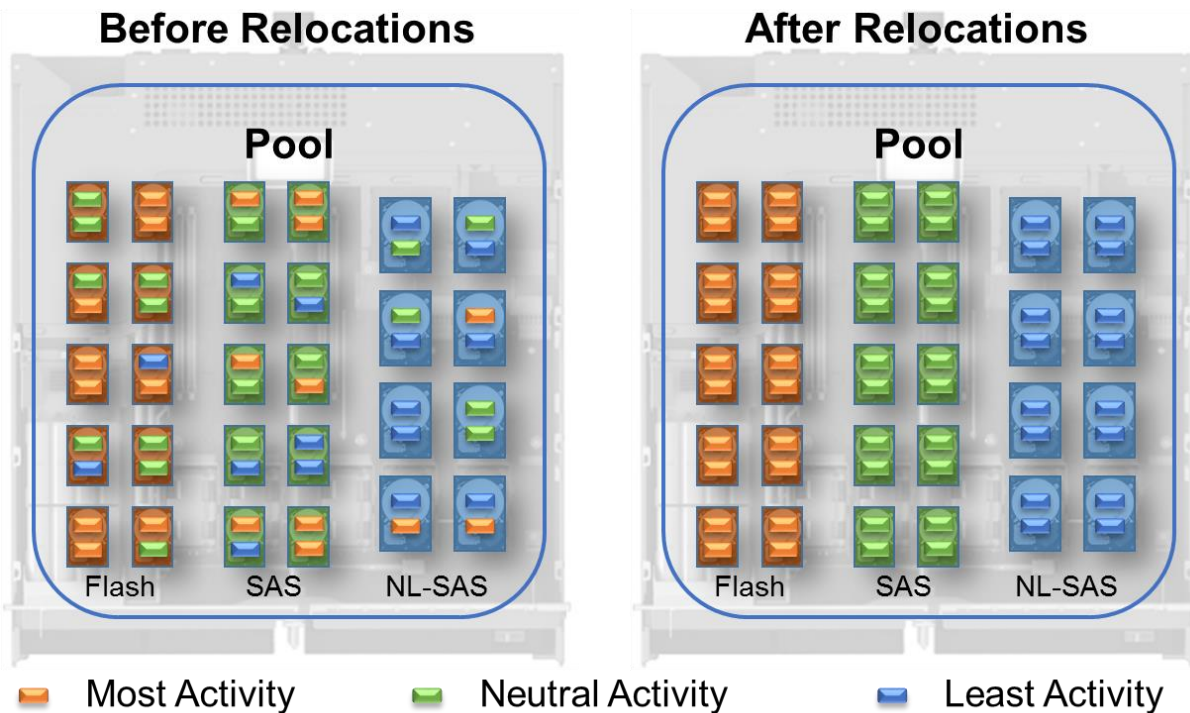


Figure 1. FAST VP Slice Relocations.

MANAGING FAST VP

The following sections outline how to manage FAST VP.

THE UNITYVSA

As UnityVSA is a virtual EMC Unity system, no physical drives exist. Pools on a UnityVSA system are created using Virtual Disks, which have been provisioned to the system from VMware. As there is a layer of abstraction between the UnityVSA system and the storage providing capacity for the Virtual Disks, FAST VP cannot automatically differentiate and assign the proper Storage Tier for each Virtual Disk. You must manually assign the Storage Tier to each Virtual Drive before they can be utilized within a Pool. Correctly matching the Storage Tier to the type of technology the Virtual Drive is created on is a crucial step, as FAST VP will use this information when tiering slices within the Pool. Typical EMC tier classification denotes Flash drives to be “Extreme Performance”, SAS drives to be “Performance”, and NL-SAS drives to be “Capacity”. EMC recommends adhering to this schema to ensure FAST VP relocation is to the appropriate Virtual Disk. Figure 2 below shows the **Tier Assignment** step, which only exists in the **Create Pool Wizard** within the UnityVSA. In this step, you must specify the **Storage Tier** for each Virtual Drive that will be used in the Pool. To do so, click the pencil icon located in the **Storage Tier** column, and select the appropriate Storage Tier label. Once the Storage Tier is specified for a Virtual Disk, and the Virtual Drive is added to a Pool, the Storage Tier cannot be changed. If many Virtual Drives exist in the UnityVSA, but only a subset of them will be utilized at this time, you will not need to assign a Storage Tier to Virtual Drives that will be left unused.

To differentiate between the Virtual Drives on the system, match the **SCSI ID** of the Virtual Drive in Unisphere to the SCSI ID of the Hard Drives in VMware. For more information on EMC UnityVSA, please see the paper *EMC Unity: UnityVSA* white paper on EMC Online Support.

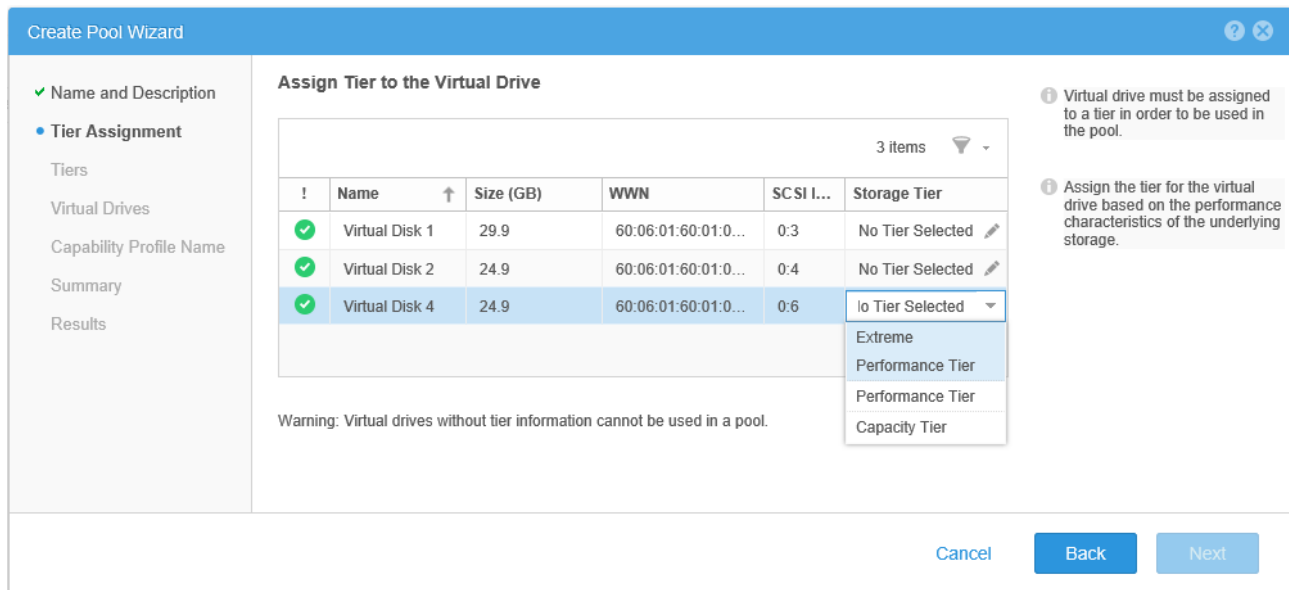


Figure 2. Create Pool Wizard. Tier Assignment Step.

SYSTEM LEVEL FAST VP MANAGEMENT

To configure FAST VP in Unisphere, select the **Update system settings** icon (⚙️), which is displayed as a gear icon on the top bar of Unisphere. Clicking this icon will open the **Settings** window. In this window many of the system settings can be configured. For FAST VP, select **Storage Configuration** in the left pane, then select **FAST VP**. Figure 3 shows an example of the **FAST VP Settings** page.

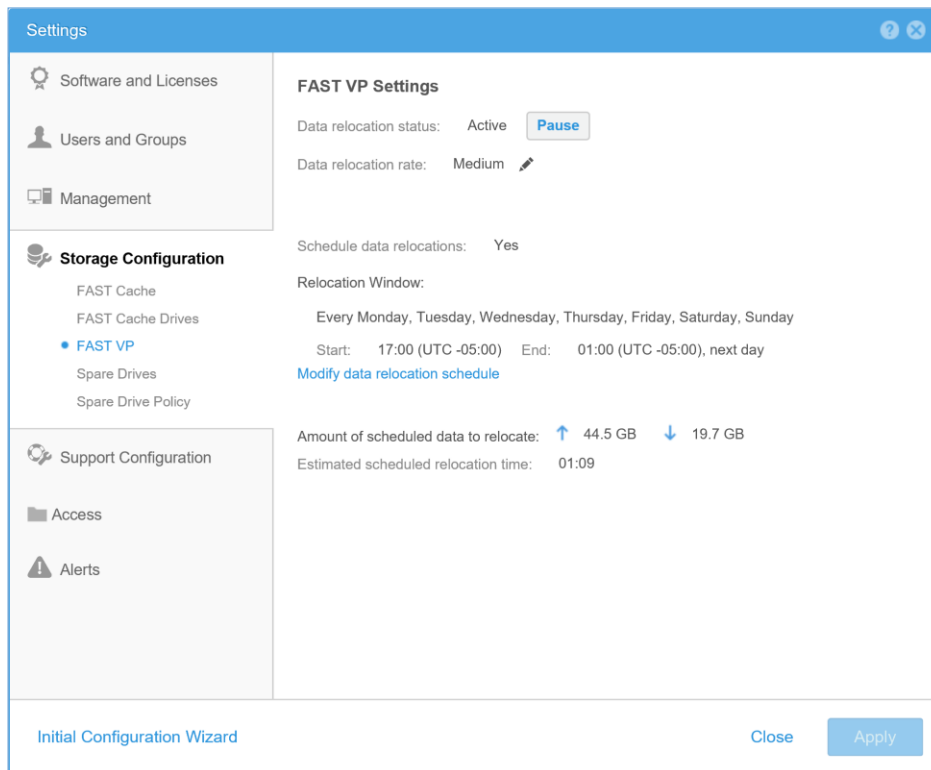


Figure 3. Settings Window. Storage Configuration – FAST VP. FAST VP Settings.

On this page you can customize FAST VP to the settings you wish to use. Near the top of the page is the **Data relocation status**. In Figure 3 above, the status is **Active**, which means FAST VP is currently active on the system. You can pause all data relocations and manual data relocations on the system by clicking the **Pause** button at any time. While the status is Paused, the button will say Resume. Resuming will resume all paused relocations on the system.

Below the Data relocation status is the **Data relocation rate**. By clicking the pencil icon, you can change the relocation rate to either High, Medium, or Low. High utilizes the most system resources to relocate data, while Low utilizes the least. The default relocation rate is Medium.

In the middle of the FAST VP Settings page, shown in Figure 3, **Schedule data relocations** displays whether or not the system is scheduled for relocations to occur. If this shows “No”, all data relocations must be manually started by the user. The **Relocation Window** is shown next. This displays which days of the week relocations are schedule for, and the **Start** and **End** times for the relocation window. By default, relocations are scheduled daily, between 22:00 (UTC +00:00), and 06:00 (UTC +00:00) of the next day. Clicking the **Modify data relocation schedule** link allows you to customize the relocations schedule further.

Lastly on the FAST VP Setting page is the **Amount of scheduled data to relocate** and the **Estimated scheduled relocation time**. These give you an idea how much data needs to move on the system based on the FAST VP algorithm, and the amount of time it will take based on the relocation rate. In this example, a large amount of data needs to move up, move down, and move within a Pool’s tier.

Shown in Figure 4 is the window that appears after selecting the **Modify data relocation schedule** link shown in Figure 3. In this window you can check or uncheck the **Schedule data relocation** checkbox to enable or disable FAST VP from running on a schedule. Below this are checkboxes for each day of the week. To have FAST VP relocations run on a particular day, ensure that day’s corresponding checkbox is checked. By default, FAST VP is scheduled to run on every day of the week. You have to ability to customize which days relocation runs on, and which days to avoid.

Also shown are the **Start time** and **End time** for FAST VP relocations. By default relocations will occur between 22:00 (UTC +00:00) and 06:00 (UTC +00:00) the next day. To customize the start or end time, simply change the values within the boxes and click **OK**. FAST VP will attempt to complete all relocations within this time period, and if this is not possible, all relocations are stopped from running. The next time FAST VP runs will be based on a new relocation priority list.

Figure 4. Modify data relocation schedule window.

POOL LEVEL FAST VP MANAGEMENT

To configure FAST VP settings on a Pool in Unisphere, simply review the **FAST VP** tab within a Pool's properties window. An Example of the tab can be found in Figure 5. On this tab, you can enable or disable FAST VP from running on this Pool by editing the **Data relocations scheduled** setting. To do so, simply click the pencil icon and change the setting. To the right of this is a link labeled **Manage FAST VP system settings**. Selecting this link redirects you to the FAST VP system settings page in Unisphere.

Also shown in this tab is **Relocation information** specific to this Pool. In Figure 5, you can see that the relocation **Status** is currently **Active**. This means relocations are currently running on this Pool.

A full list of Status information is shown below:

- **Paused** – Data relocation is paused on the system.
- **Active** – Pool is actively relocating data.
- **Not Started** – Data relocations have not been started on the Pool.
- **Completed** – Data relocation have completed for the Pool.
- **Stopped by user** – Relocations have been manually stopped by the user.
- **Failed** – Data relocations have failed for the Pool.
- **Not Applicable** – No applicable information for FAST VP exists for the Pool.

The estimated **Time to relocate** the data on the Pool, displayed in hours and minutes, is also displayed. From this page you can also see the **Last start time** and **Last end time** for relocations within this Pool. You will also notice that the button below this information says **Stop Relocations**. As relocation are currently occurring on the Pool, you can stop them at any time by selecting this button. If relocations were not currently running, the button would say **Start Relocation**, and you could select it to start a relocation on the Pool.

In the bottom of the FAST VP tab is the Pool's Tier information. As shown in Figure 5, each Tier configured within this Pool is displayed, along with how many drives and which RAID type is used for the Pool. Also in the chart is information regarding how much data per tier needs to **Move Up**, **Move Down**, and **Rebalance** within the tier. Each of the totals are displayed in **GBs**. In this example, a large amount of data is being relocated across tiers. Lastly, the chart displays the **Total Size** and **Free Size**, both in **TBs** for each tier within the Pool. From this chart, you are able to see the configuration, the scheduled relocation, and the total and free capacity for the Pool.

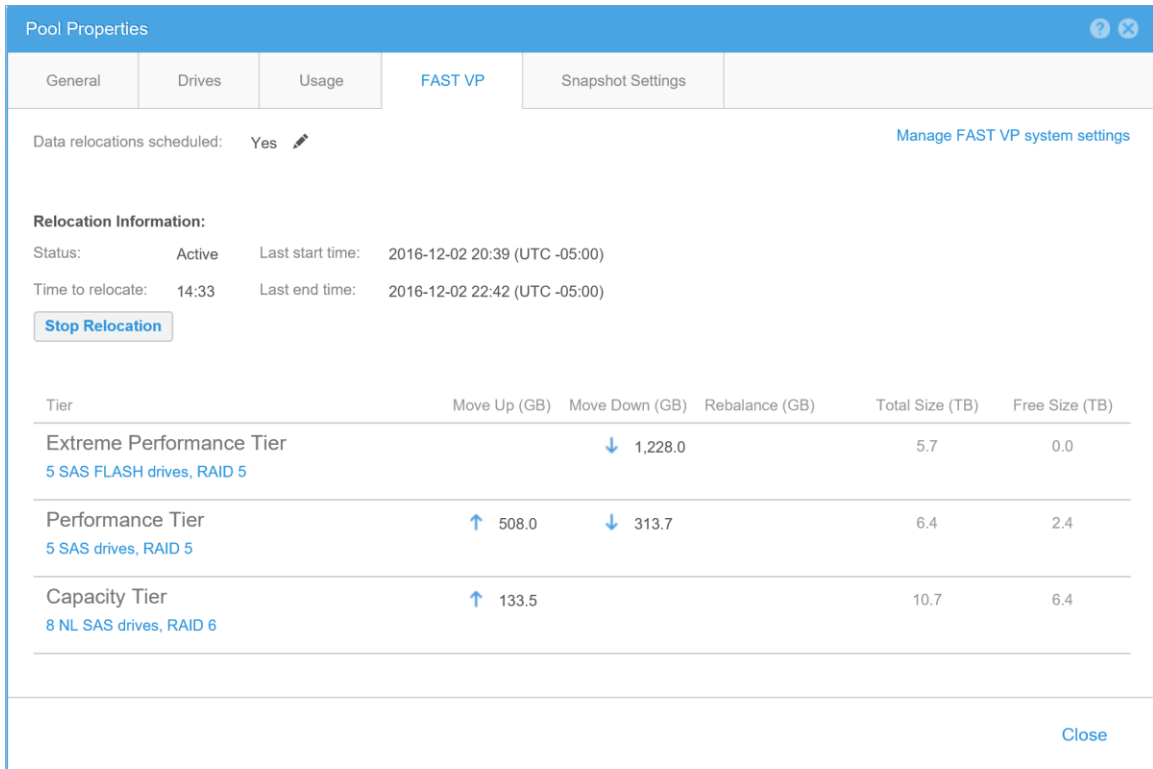


Figure 5. Pool Properties. FAST VP tab.

When available, the **Start Relocation** button brings you to the **Start Data Relocation** window, shown in Figure 6. From this window you are able to choose the **Data relocation rate**, either High, Medium, or Low, and choose an **End** time for relocations. By default, the end time is 2 hours from the local time within the system. After changing any settings within the window, click **OK** to start relocations manually on the Pool.

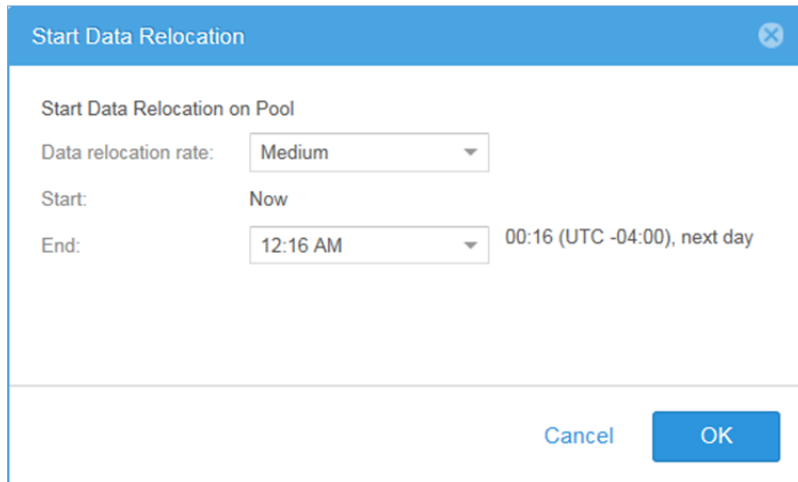


Figure 6. Start Data Relocation window.

STORAGE RESOURCE LEVEL FAST VP MANAGEMENT

When creating a storage resource, you are able to customize the Tiering Policy that the resource will use. Figure 7 shows an example of the **Configure** step within the **Create LUNs** wizard. In this step, you are able to choose which tiering policy the resource will use. By default, the **Tiering Policy** will be **Start High then Auto-Tier**. Other choices include **Auto-Tier**, **Highest Available Tier**, and **Lowest Available Tier**.

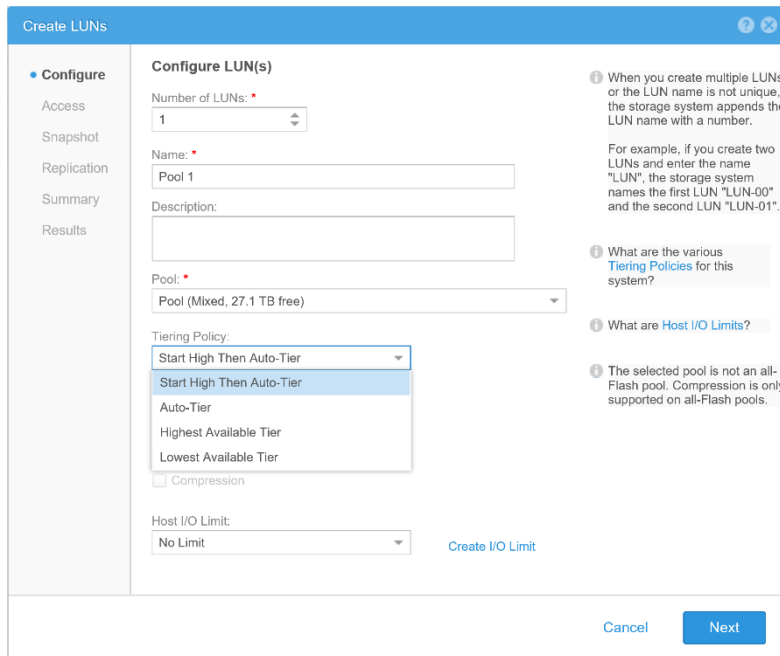


Figure 7. Create LUNs Wizard. Configure step.

You can also change FAST VP settings on a storage resource after it has been created. Figure 8 shows an example of the FAST VP tab within the LUN Properties window. From here, you can change the FAST VP **Tiering policy** for the resource at any time by selecting another one from the drop down list. Also shown are the Tiers within the Pool, and the resources **Data Distribution** across those tiers. In this example, 32% of the LUN's data resides on the **Extreme Performance Tier**, 67% resides on the **Performance Tier**, and 1% resides on the **Capacity Tier**.

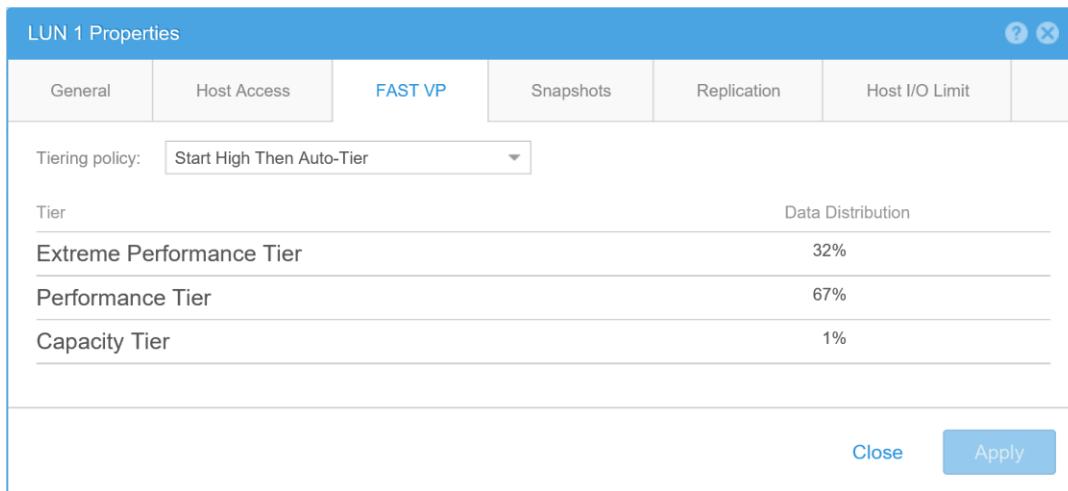


Figure 8. LUN Properties Window. FAST VP Tab.

EXPANDING A POOL

At any time you can expand a Pool in Unisphere by selecting the **Expand Pool** option in the **Pools** page in Unisphere. The location of the **Expand Pool** option can be seen in Figure 9. This will launch the Expand Pool wizard, which allows you to select which drives to add to the selected Pool. When expanding a Pool, you can either add a new tier to the Pool, or add drives to an existing tier.

Pools							
				Size (TB)	Free (TB)	Used (%)	Subscription (%)
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Pool	↑	22.9	8.8	<div style="width: 40%;"></div>	62.2

Figure 9. Unisphere Pools Page.

Figure 10 below shows the Pool Properties Window after expanding a Pool. The Pool was expanded by adding drives to the Performance Tier. Once a tier is expanded, a Rebalance is started to spread the existing data within the tier across all drives within the tier. In the **Rebalance** column, notice a double arrow is displayed with the amount of data to rebalance.

Pool Properties ? ×

General

Drives

FAST VP

Snapshot Settings

Data relocations scheduled: Yes ✎ Manage FAST VP system settings

Relocation Information:

Status: Active Last start time: 2016-12-02 20:39 (UTC -05:00)

Time to relocate: 21:30 Last end time:

Stop Relocation

Tier	Move Up (GB)	Move Down (GB)	Rebalance (GB)	Total Size (TB)	Free Size (TB)
Extreme Performance Tier 5 SAS FLASH drives, RAID 5		↓ 1,227.0		5.7	0.0
Performance Tier 10 SAS drives, RAID 5	↑ 504.7	↓ 312.7	↔ 1,046.7	12.8	8.8
Capacity Tier 8 NL SAS drives, RAID 6	↑ 135.7			10.7	6.4

Close

Figure 10. Pool Properties. FAST VP Tab After Expanding a Pool.

FAST CACHE

INTRODUCTION

The EMC Unity FAST software includes FAST Cache and FAST VP. FAST Cache utilizes flash drives as an additional cache layer within the system to temporarily store highly accessed data. For data not already located on Flash, the system copies the highly accessed 64 KB chunks of data from their current locations on spinning drive to FAST Cache. Repeated access to this data will benefit by taking advantage of the high IOPs and low response times Flash drives provide. As FAST Cache is a global resource on the system, all data can benefit from this caching layer and the overall performance of the system can increase. When a piece of data located on spinning drive is marked for promotion into FAST Cache and there is currently no free FAST Cache pages, FAST Cache will free a page by removing the Least Recently Used (LRU) chunk of data. If the data being removed from FAST Cache is dirty, meaning the data has not been synchronized with the location on the Pool, the data is first copied back to its location on drive before being removed from FAST Cache.

FAST CACHE LICENSING

In EMC Unity, FAST Cache is only supported on a Unity Hybrid system, which includes models Unity 300, Unity 400, Unity 500, and Unity 600. The FAST Cache license is bundled in the EMC Unity Essentials Software package, which is included with all Unity Hybrid systems. To create FAST Cache on a Unity Hybrid system, the system must first have the FAST Cache license installed. Once the license is installed, you may configure FAST Cache up to the limits of the system, which are discussed later in this document.

FAST CACHE COMPONENTS

FAST Cache consists of multiple hardware and software components, including the drives, the Policy Engine, and the Memory Map. To create FAST Cache, you must first have at least two drives supported for use in FAST Cache available in the system. FAST Cache is configured using RAID 1 mirrored pairs, and only configuring FAST Cache drives in multiples of two is supported. Supported drives and system limits are outlined in the *Supported Drives and Configurations* section of this paper. The Policy Engine and Memory Map are defined below:

POLICY ENGINE

The FAST Cache Policy Engine is the software which monitors and manages the I/O flow through FAST Cache. The Policy Engine keeps statistical information on blocks on the system and determines what data is a candidate for promotion. A chunk is marked for promotion when an eligible block is accessed from spinning drives three times within a short period of time. The block is then copied to FAST Cache and the Memory Map is updated. More information on which data is eligible for FAST Cache promotion is discussed later in this document.

MEMORY MAP

The FAST Cache Memory Map contains information of all 64 KB blocks of data currently residing in FAST Cache. Each time a promotion occurs, or a block is replaced in FAST Cache, the Memory Map is updated. The Memory Map resides in DRAM memory and on the system drives to maintain high availability. When an I/O reaches FAST Cache to be completed, the Memory Map is checked and the I/O is either redirected to a location in FAST Cache or to the Pool to be serviced.

THEORY OF OPERATION

FAST CACHE PROMOTIONS

A FAST Cache Promotion is the operation of copying data from spinning media Hard Disk Drives (HDDs) to a location within FAST Cache. This operation is labelled as a promotion due to the block of data being copied into higher performing drives due to its access pattern. While the block of data is located in FAST Cache, the data still resides in its original location within the Pool. The data residing within the Pool is periodically updated during a FAST Cache Cleaning operation or a FAST Cache flush, both described later in this document. A FAST Cache promotion is only considered on Pools with FAST Cache enabled.

A FAST Cache promotion occurs when the Policy Engine determines the performance for a chunk of data would benefit by residing in FAST Cache. While the Policy Engine is monitoring the I/O to FAST Cache enabled Pools, data access patterns are reviewed. When a chunk of data is accessed three times within a certain period of time, the eligibility of the block is checked and the block is marked for promotion into FAST Cache. If there are free FAST Cache blocks available, the data is copied into FAST Cache. If FAST Cache is full, the access pattern of the data being considered for promotion is compared to the access pattern of data in FAST Cache. If the access

pattern of the data considered for promotion exceeds that of a chunk of data in FAST Cache, the least accessed data in FAST Cache is flushed out of FAST Cache and the new promotion replaces it. The Memory Map is then updated to include all changes to the contents of FAST Cache. The next time the promoted block is accessed, assuming System Cache could not complete the I/O, the FAST Cache Memory Map will be checked and the I/O will be serviced from FAST Cache. While promoted into FAST Cache, the chunk of data has the potential for higher overall throughput and lower Response Time. When a large portion of a data set resides in FAST Cache, applications can also benefit with the increased performance FAST Cache can provide.

There are multiple circumstances in which the access pattern of an application is assumed to cause a FAST Cache promotion, but it doesn't occur. In some cases, the efficiencies in system cache are more suited to handle the I/O, while other times the configuration or location of the data stops the promotion. Some of these circumstances are:

- FAST Cache is not enabled on the system or Pool
In this instance, FAST Cache is not running on the resource, so no FAST Cache promotions will occur.
- The data currently resides on Flash Drives with a Pool
FAST Cache will not promote data which already is serviced from Flash
- Small block sequential I/O
System Cache within the EMC Unity system is optimized to handle sequential I/O. When a sequential workload is identified, prefetching larger blocks of data into System Cache is done to increase the performance of the workload. The larger I/Os to the drives helps avoid FAST Cache promotions from occurring.
- High-frequency access patterns
Data which is accessed frequently enough to continually reside in System Cache will not be promoted to FAST Cache. As the I/O is continuously serviced from System Cache, I/O is not issued to the Pool to cause a FAST Cache promotion.
- Zero fill requests
Zero fill requests are typically a one-time occurrence, and will be handled by System Cache. System Cache will coalesce the I/O into large I/Os which tend to avoid promotion into FAST Cache.
- I/O sizes larger than the stripe length
FAST Cache promotions are avoided when I/O sizes exceed the RAID configuration's stripe length. For example, in an 8+1 RAID 5 configuration, I/Os above 512KB (8x64KB) in size will not cause promotions.

READ OPERATIONS

When a read is sent to the system, the contents of System Cache are reviewed to determine if the I/O can be completed. If the data being requested resides in System Cache, the I/O is completed from the data currently residing in System Cache. The Pool is not sent an I/O request, and the Fast Cache Memory Map is not checked.

If System Cache cannot complete the I/O, a System Cache miss occurs. If FAST Cache is enabled, the Memory Map is reviewed to see if the contents of FAST Cache can complete the I/O. If the data resides in FAST Cache, the I/O is redirected to the location within FAST Cache the data resides, the data is copied into System Cache, and the read request is completed.

When the data being requested is not currently located in FAST Cache, the data must be requested from the drive the data resides on. The data is then copied from drive to System Cache, and the read operation completes with the requestor of the information. If the data has been accessed frequently, the Policy Engine will cause a promotion to occur and the data will be copied into FAST Cache. Subsequent requests for this information will either come from System Cache or FAST Cache.

Figure 11 shows a read operation.

1. The Unity system receives a Read operation.
2. If the data being requested is located in System Cache, the I/O is complete using the contents of System Cache. This is also known as a System Cache hit.
3. If System Cache does not contain the data requested, this is known as a System Cache miss. If FAST Cache is enabled, the FAST Cache Memory Map is reviewed to determine if the data being requested resides in FAST Cache.
4. If the data resides in FAST Cache, a FAST Cache hit is seen. The information is then copied into System Cache, and System Cache completes the Read request.
5. If a FAST Cache miss is seen, the data is requested from the Pool HDDs, and is copied into System Cache. System Cache then completes the Read request.
6. The FAST Cache Policy Engine promotes the data into FAST Cache and updates the Memory Map if the data has been accessed frequently.

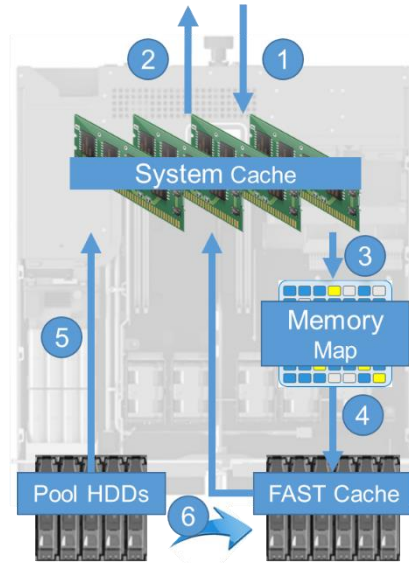


Figure 11. Read Operation.

WRITE OPERATIONS

In EMC Unity, all write operations to the system are completed using System Cache. The system accepts the write request into System Cache, and an acknowledgment is sent to the host. No interaction with FAST Cache happens during this time. Write Cache on a EMC Unity system cannot be disabled by the user, and is enabled in almost all circumstances.

If Write Cache on the system becomes disabled, writes to the system will need to be saved on drive before the write operation can be acknowledged. During this operation, the data is temporarily held in System Cache, while the data is saved to drive. If FAST Cache is enabled, the Memory Map is reviewed to see if a copy of the data resides in FAST Cache. If so, the data in FAST Cache is updated and the write operation is acknowledged. If the data does not reside in FAST Cache, a write to the Pool HDDs occurs and the write is acknowledged. This write to the Pool HDDs may cause a FAST Cache promotion to occur.

1. A write operation is received by the system while Write Cache is disabled. The data is temporarily held in System Cache while the data is saved to disk.
2. When FAST Cache is enabled, the contents of the Memory Map are reviewed to determine if a copy of the data being updated resides in FAST Cache.
3. If the data resides in FAST Cache, the data is copied from System Cache to FAST Cache.
4. If a copy of the data does not reside in FAST Cache, the data is written directly to the Pool HDDs.
5. The write is acknowledged once the data resides on disk, either in FAST Cache or within the Pool.
6. A FAST Cache promotion may occur if the block is accessed frequently.

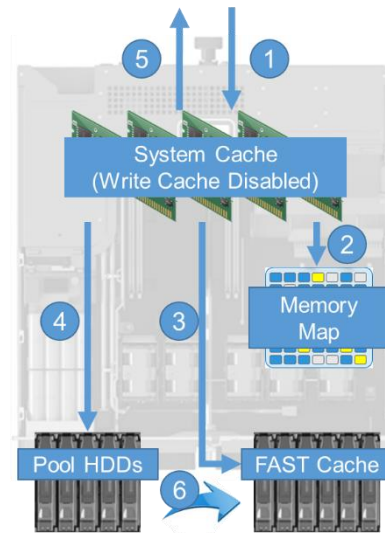


Figure 12. Write Operation with System Write Cache disabled.

In the instance when System Cache is proactively cleaning cache pages or flushing cache pages, outlined in Figure 13 below, updates to FAST Cache may be seen. During this operation, the FAST Cache Memory Map is reviewed to determine if the data being overwritten resides in FAST Cache. If the data is located within FAST Cache, the data being cleaned from System Cache is synchronized with the contents of FAST Cache. As FAST Cache now contains data which is newer than what resides on the Pool, the data is considered dirty. This is also known as a FAST Cache dirty page. This data will be synchronized with the data on the Pool when a FAST Cache page cleaning operation occurs. If FAST Cache does not contain a copy of the data being updated, the data is written directly to the Pool drives. This operation may cause a FAST Cache Promotion to occur.

1. During a System Cache page cleaning or page flushing operation, the FAST Cache Memory Map is reviewed if FAST Cache is enabled.
2. If a copy of the data resides in FAST Cache, System Cache will update the data block in FAST Cache with the latest information.
3. If the data does not reside in FAST Cache, System Cache will write the data to the Pool HDDs directly.
4. If the data is accessed frequently, the write operation may cause the data to be promoted to FAST Cache. If a promotion occurs, the data is copied to FAST Cache and the Memory Map is updated.

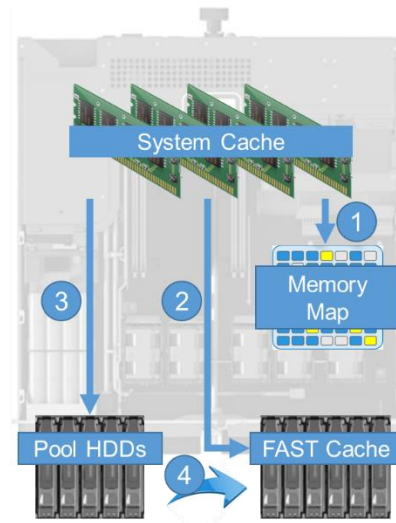


Figure 13. System Cache cleaning operation.

FAST CACHE PAGE CLEANING

FAST Cache page cleaning occurs when FAST Cache needs to synchronize the contents of a FAST Cache page with the data's location in the Pool. This happens when FAST Cache has more recently updated data than the Pool HDDs. Cleaning occurs to reduce the amount of FAST Cache dirty pages. When page cleaning starts, a 64 KB FAST Cache page is copied to the Pool HDD in the corresponding location for the data. Once the data is written and the contents of FAST Cache and the Pool are identical, the page in FAST Cache is marked as clean. The data remains in FAST Cache for future use.

For FAST Cache promotions to occur efficiently, FAST Cache free or clean pages need to exist. If no free or clean pages exist, a page cleaning operation needs to happen before the page can be freed for the next promotion to occur. When a promotion is scheduled, pages are used in the following order:

1. Consume a free, unused FAST Cache page.
2. If no free page exists, remove the data from the Least Recently Used clean page, and use the page for the scheduled promotion.
3. If no free or clean pages exist, copy the contents of the Least Recently Used page to HDD, free the page, and use the page for the promotion. Cleaning a FAST Cache Dirty page and freeing it for a promotion is also called a FAST Cache flush operation.

FAST CACHE ONLINE EXPAND

EMC Unity allows users to increase the configured size of FAST Cache online, without impacting FAST Cache operations on the system. Online expansion gives users the option of first configuring FAST Cache with a minimal amount of drives, and growing the configuration as demands on the system are increased. To expand FAST Cache, free drives of the same size and type currently used in FAST Cache must exist within the system. FAST Cache is expanded in pairs of drives, and can be expanded up to the system maximum.

When a FAST Cache expansion occurs, a background operation is started to add the new drives into FAST Cache. This operation first configures a pair of drives into a RAID 1 mirrored set. The capacity from this set is then added to FAST Cache, and is available for future promotions. These operations are repeated for all remaining drives being added to FAST Cache. During these operations, all FAST Cache reads, writes, and promotions occur without being impacted by the expansion. The amount of time the expand operation takes to complete depends on the size of drives used in FAST Cache and the number of drives being added to the configuration.

FAST CACHE ONLINE SHRINK

In EMC Unity, removing drives from FAST Cache is possible while FAST Cache is configured and servicing I/O. If at any time a number of drives need to be removed from FAST Cache, a shrink operation can be started. FAST Cache shrink is issued in pairs of drives, and allows the removal of all but 2 drives from FAST Cache. To remove all drives from FAST Cache, the Delete operation is used. FAST Cache shrink is often utilized when drives need to be repurposed to a Pool for expanded capacity. Removing drives from FAST Cache can be a lengthy operation, and can impact system performance.

When a FAST Cache shrink occurs, a background operation is started to remove drives from the current FAST Cache configuration. Removing drives from FAST Cache reduces the size of FAST Cache by the number of drives selected. After starting a shrink operation, new promotions are blocked to each pair of drives selected by the system to be removed from FAST Cache. Next, each FAST Cache dirty page within the drives to be removed is cleaned to ensure that data is synchronized with the locations on the Pool. After all dirty pages are cleaned within a set of drives, the capacity of the set is removed from the FAST Cache configuration. Data which existed on FAST Cache drives that were removed may be promoted to FAST Cache again through the normal promotion mechanism.

FAST CACHE MANAGEMENT

The following sections outline the steps to create FAST Cache, view the current FAST Cache configuration on the system, expand FAST Cache, shrink FAST Cache, and delete FAST Cache. Examples for each of the operations will be shown in Unisphere, though they can also be completed using UEMCLI or REST API.

CONFIGURING FAST CACHE

To configure FAST Cache in Unisphere, select the **Update system settings** icon (⚙️), which is displayed as a gear icon on the top bar of Unisphere. Clicking this icon will open the Settings window. In this window many of the system settings can be configured. For FAST Cache, select **Storage Configuration** in the left pane, then select FAST Cache. When FAST Cache is not configured on the system, you will see what is shown in Figure 14. To configure FAST Cache, click **Create**.

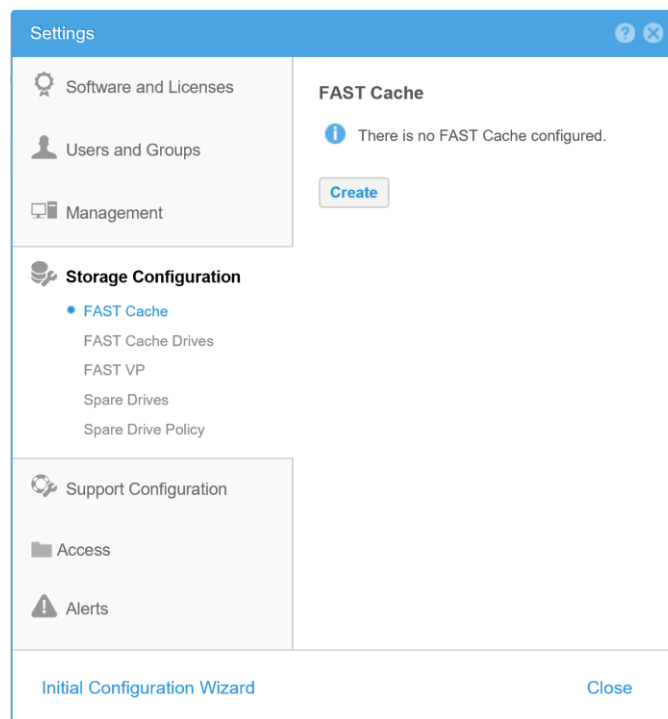


Figure 14. Settings Window. Storage Configuration – FAST Cache. FAST Cache not configured.

The **Drives** step of the Create FAST Cache wizard is now shown. Figure 15 shows an example of what is seen when multiple sizes of drives supported for use in FAST Cache are present in the system. In this example, a Unity 600 system has multiple 800 GB and 400 GB free drives. FAST Cache can only be created using drives of the same size, so a radio icon is displayed before each drive size in the list. To select a certain drive size to use for FAST Cache, select the radio icon in front of the desired drive size. Next, click the drop-down box and select the number of drives you will use for FAST Cache.

At this time you can also choose to have FAST Cache enabled on all existing Pools on the system. FAST Cache is a global resource, which can be utilized by all Pools within the system. The **Enable FAST Cache for existing pools** checkbox is checked by default, and can be deselected before proceeding. To change the FAST Cache setting for a particular Pool, view the General tab within the Pool Properties window. After selecting the desired drive size and the number of drives to use in FAST Cache, click the **Next** button.

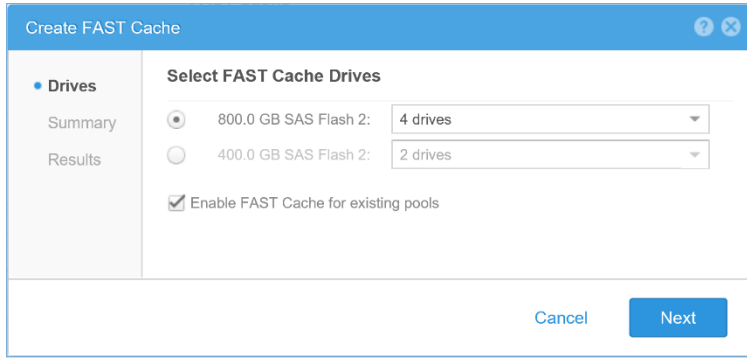


Figure 15. Create FAST Cache Wizard. Drives step.

The **Summary** step for the Create FAST Cache Wizard is now displayed, and an example of this screen can be seen in Figure 16. This screen displays the choices made in the previous step, and allows you to confirm the proper selections were made. If the incorrect drive size was selected, or you wish to change the number of drives selected, you can select the Back button to correct the information. After reviewing this screen and confirming the correct information is displayed, select **Finish** to create FAST Cache with these settings.

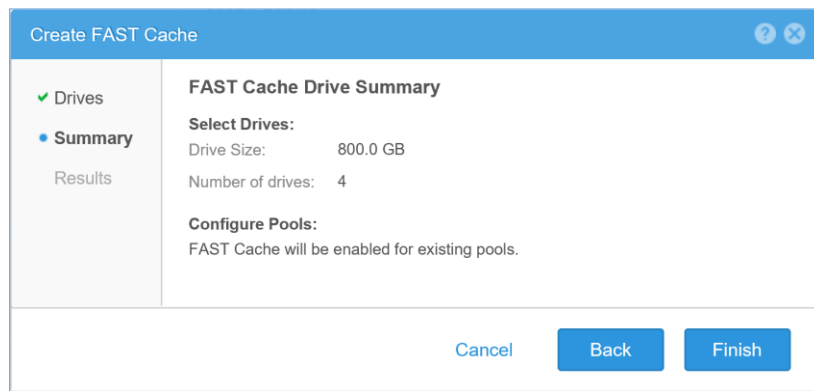


Figure 16. Create FAST Cache Wizard. Summary step.

The **Results** step for the Create FAST Cache Wizard is now shown. The overall status of the FAST Cache creation is shown, along with each job and its status. Figure 17 shows an example of this window when all processes are complete. For each pair of drives, a RAID 1 RAID Group is created, and the capacity of the group is added to FAST Cache. Once capacity is added to FAST Cache, FAST Cache is enabled and available for data promotions. A process is also started to enable FAST Cache on all Pools on the system if that option was selected. While FAST Cache is being created, you can click **Close** to close out of this window. The process is a Unisphere job and will continue to run in the background.

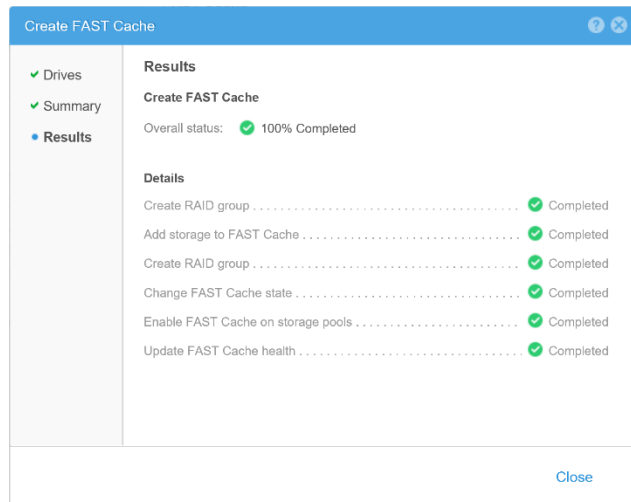


Figure 17. Create FAST Cache Wizard. Results step.

EDITING THE FAST CACHE SETTING ON A POOL

To change the FAST Cache settings of a Pool, you must view the Properties window for the Pool. On the General tab you can change if the Pool is allowed to use FAST Cache by either clearing or checking the **Use FAST Cache** checkbox and clicking Apply. An example of the General tab is shown in Figure 18. By default, FAST Cache will be enabled at creation as shown in Figure 15 unless you uncheck the **Enable FAST Cache for existing Pools** checkbox. The **Use FAST Cache** checkbox cannot be changed while FAST Cache is disabled, or FAST Cache is being deleted.

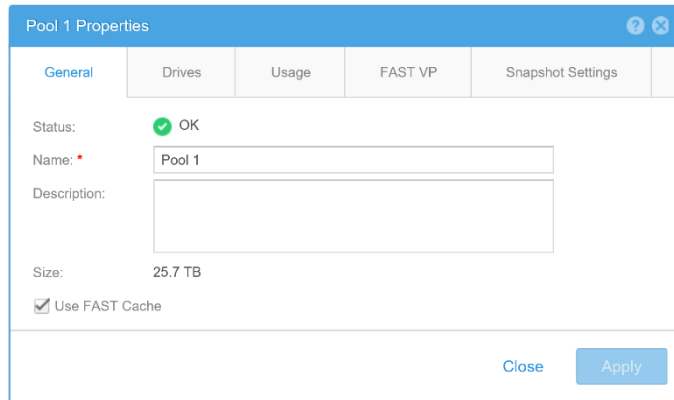


Figure 18. Pool Properties Window. General Tab.

VIEWING THE FAST CACHE CONFIGURATION

When FAST Cache is enabled on the system, you can view the FAST Cache configuration in the system Settings window. To open the Settings window in Unisphere, select the **Update system settings** icon. While in the Settings window, select **Storage Configuration** in the left pane, then select **FAST Cache**. Shown is the current FAST Cache Status and configuration. As shown in Figure 19, the Status of FAST Cache is OK, and FAST Cache currently is configured with 4 Drives. From this window, you have the option to Expand, Shrink, or Delete FAST Cache.

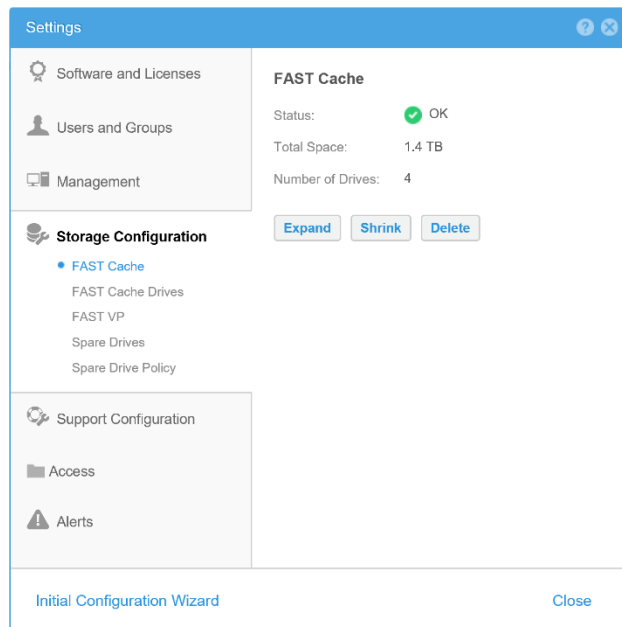


Figure 19. Settings Window. Storage Configuration – FAST Cache. FAST Cache configured.

To view which drives were used to configure FAST Cache, select **FAST Cache Drives**, which is found under **Storage Configuration** in the Settings window. As shown in Figure 20, the FAST Cache drives are displayed, shown with their locations and size. In this example, drives from the Drive Processor Enclosure were selected to configure FAST Cache. This window makes it easy to locate the physical location of each of the FAST Cache drives.

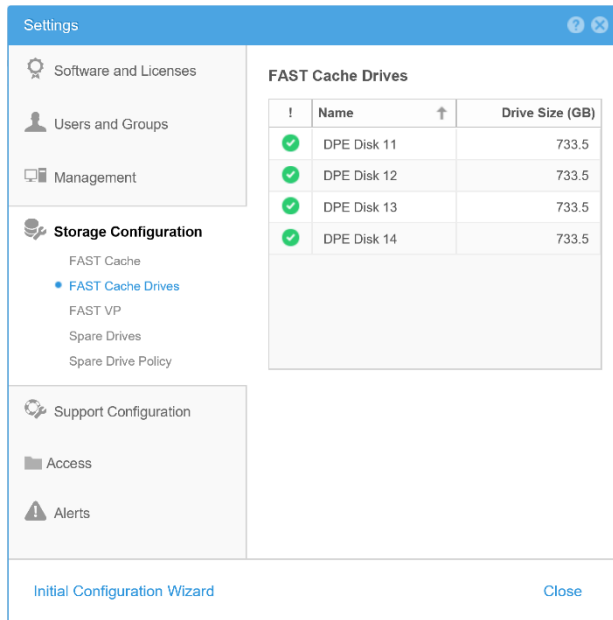


Figure 20. Settings Window. Storage Configuration – FAST Cache Drives.

EXPANDING FAST CACHE

When FAST Cache is enabled on the EMC Unity system, you have the option to expand the capacity of FAST Cache up to the system maximum. To expand FAST Cache on a EMC Unity system from Unisphere, you must navigate to the **FAST Cache** page found under **Storage Configuration** in the **Settings** window. From this window, select **Expand** to start the **Expand FAST Cache** wizard.

Figure 21 shows an example of the Expand FAST Cache wizard. When expanding FAST Cache, you may only select free drives of the same size and type as what is currently in FAST Cache. In this example, only 800 GB SAS Flash 2 drives are available to be selected, as FAST Cache is currently created with those drives. From the drop-down list, you are able to select pairs of drives to expand the capacity of FAST Cache up to the system maximum. In this example, only 2 free drives were found. Click **OK** to start the expansion process.

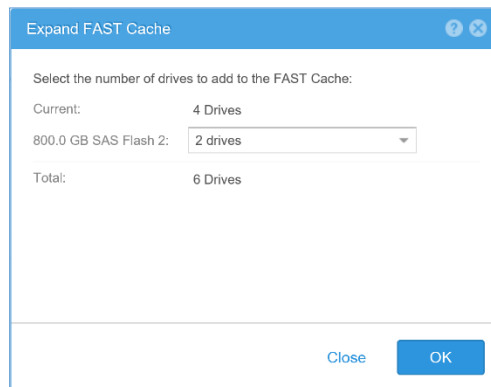


Figure 21. Expand FAST Cache Wizard.

After clicking OK, an Expand FAST Cache job is created to add the drives to the FAST Cache configuration. This process occurs in the background, and does not impact I/O or promotions to FAST Cache. Figure 22 shows an example of the Job Properties window for the Expand FAST Cache job. Shown is the overall status of the operation, and the individual steps taken for the process. In this example, only two drives were added to the configuration.

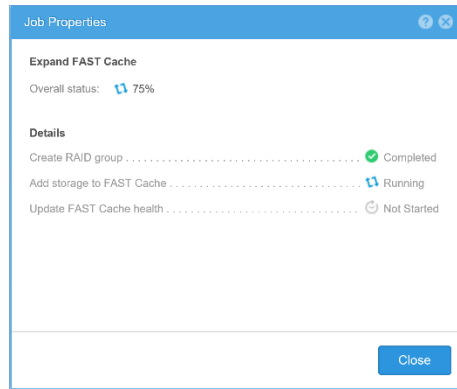


Figure 22. Job Properties. Expand FAST Cache Job.

SHRINKING FAST CACHE

As with expanding FAST Cache, you also have the option to shrink the capacity of FAST Cache when FAST Cache is enabled on the system. To shrink FAST Cache on a EMC Unity system from Unisphere, you must navigate to the **FAST Cache** page which is found under **Storage Configuration** in the **Settings** window. From this window, select **Shrink** to start the **Shrink FAST Cache** wizard.

Figure 23 shows an example of the Shrink FAST Cache wizard, with and without the drop-down box selected. In this window you can see that FAST Cache currently contains six Drives. During a shrink operation you can remove all but two of the drives currently configured in FAST Cache. In this example, two drives will be removed from FAST Cache.

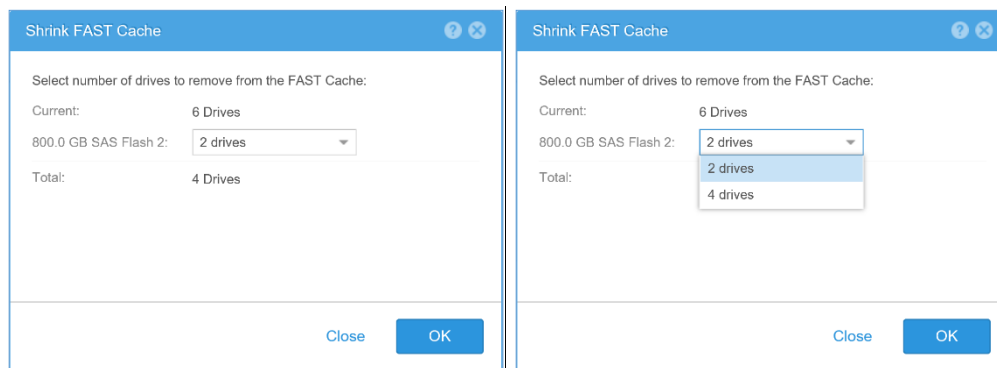


Figure 23. Shrink FAST Cache Wizard.

Figure 24 shows the warning message received when shrinking drives out of the FAST Cache configuration. The warning message outlines that all FAST Cache data must be flushed from the drives being removed from FAST Cache. This operation takes time to complete, and can vary based on the I/O workload being seen by FAST Cache and the Pool drives. Performance of the system may also be impacted due to data no longer residing in FAST Cache. Hot data will need to promote again once flushed out of FAST Cache.

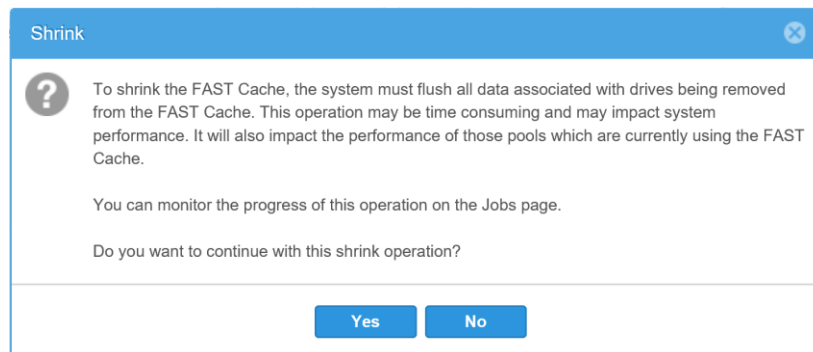


Figure 24. Shrink Fast Cache Warning.

While the FAST Cache Shrink operation is completed, no changes to FAST Cache can be made. Figure 25 shows the FAST Cache page within system settings during the time a shrink operation was occurring. Notice that no options within the FAST Cache are available while the operation is running.

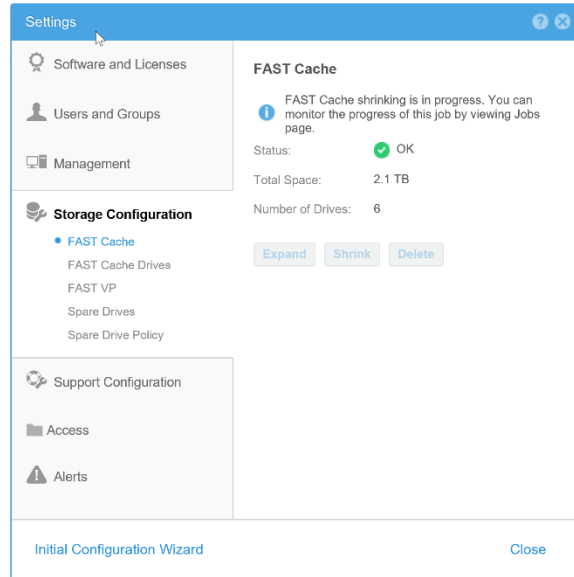


Figure 25. Settings Window. Storage Configuration – FAST Cache.

At any time you can view the status of the FAST Cache Shrink job by navigating to the **Jobs** page found under the **Events** heading in the left pane of Unisphere. Figure 26 shows an example of a completed FAST Cache Shrink job.

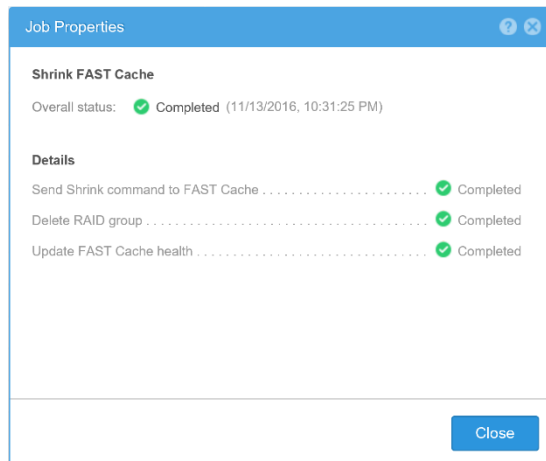


Figure 26. Job Properties. Shrink FAST Cache Job.

DELETING FAST CACHE

The last option found on the FAST Cache page in the Settings window is **Delete**. Delete is utilized when you want to disable FAST Cache and remove all drives from the FAST Cache configuration. This may be done to repurpose the drives to a Pool, or to recreate FAST Cache with different capacity drives. From the FAST Cache page in the Settings window, select **Delete** to delete the current FAST Cache configuration.

Figure 27 shows the message received after selecting Delete. This message states that all data must be flushed from FAST Cache which can be a time consuming operation. Performance of the system may also be impacted during the delete operation as the contents of FAST Cache will need to be copied to the Pool drives on the system. Because the data flushed from FAST Cache may no longer reside on Flash, this data may see increased Response Times.

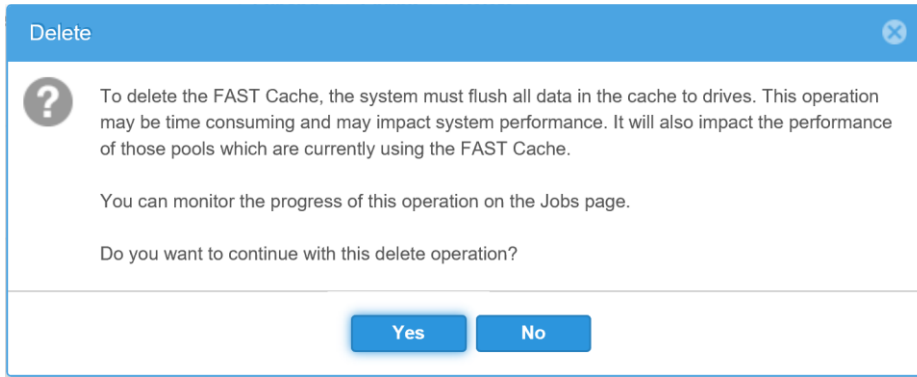


Figure 27. Delete FAST Cache Warning.

Figure 28 shows the FAST Cache page within the Settings window while the FAST Cache Delete operation is running. No changes can be made to FAST Cache while a Delete is running. More information on the progress of the operation can be seen on the Jobs page in Unisphere. Once the operation completes you will be able to create FAST Cache again.

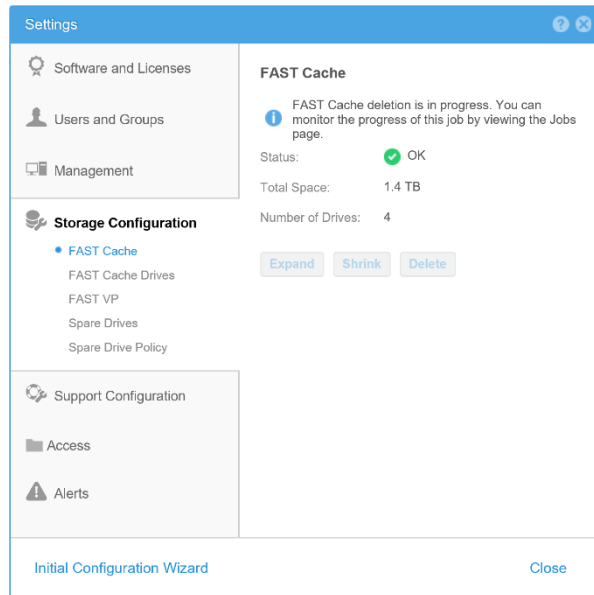


Figure 28. Settings Window. Storage Configuration – FAST Cache. FAST Cache Disabling.

You can review the current status during a Delete FAST Cache operation by reviewing the Jobs page. Figure 29 shows an example of the processes seen during a delete job. In this example the Delete operation was just started, and the first operation was still running.

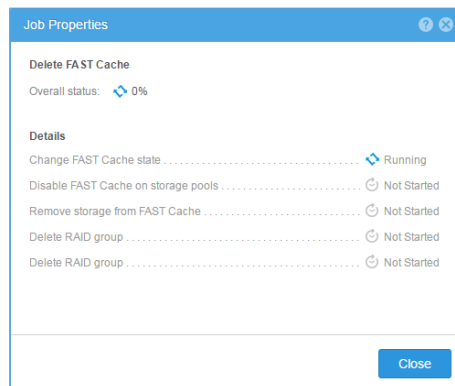


Figure 29. Job Properties. Delete FAST Cache Job.

SUPPORTED DRIVES AND CONFIGURATIONS

In EMC Unity, FAST Cache is only supported on Unity Hybrid systems, which includes models Unity 300, Unity 400, Unity 500, and Unity 600. The Unity 300, Unity 400, and Unity 500 all support either 200 GB SAS Flash 2 drives, or 400 GB SAS Flash 2 drives in FAST Cache. The Unity 600 system supports 200 GB SAS Flash 2 drives, 400 GB SAS Flash 2 drives, or 800 GB SAS Flash 2 drives in FAST Cache. The following table, Table 4, shows each Unity model, the SAS Flash 2 drives supported for that model, and the maximum FAST Cache capacities. For each of the maximum FAST Cache capacity, the total number of drives needed to reach that capacity are listed.

Table 4. Maximum FAST Cache Configurations

UNITY SYSTEM	SAS FLASH 2 SUPPORTED DRIVES	MAXIMUM FAST CACHE CAPACITY (TOTAL # OF DRIVES NEEDED)
UNITY 300	200 GB	800 GB (8 Drives)
	400 GB	800 GB (4 Drives)
UNITY 400	200 GB	1200 GB (12 Drives)
	400 GB	1200 GB (6 Drives)
UNITY 500	200 GB	3200 GB (32 Drives)
	400 GB	3200 GB (16 Drives)
UNITY 600	200 GB	4000 GB (40 Drives)
	400 GB	6000 GB (30 Drives)
	800 GB	5600 GB (14 Drives)

FAILURE HANDLING

In EMC Unity, if a single drive failure occurs within FAST Cache, the private RAID 1 mirrored pair that is affected will enter a degraded state. FAST Cache page cleaning begins copying dirty pages from the degraded group to the Pool HDDs to ensure the data is protected on the system. This cleaning operation runs at a faster rate than the normal cache cleaning operation. While in degraded mode, only read operations are allowed to the degraded group as to reduce the chance of data loss. Write operations targeted for the degraded group are redirected to the remaining FAST Cache drives within the configuration. Only after the failed drive is replaced or a spare is chosen and fully rebuilt will the degraded pair enter a normal state and allow for promotions to occur.

When a failed drive is encountered, as with Pool drives, FAST Cache drives can permanently spare to free drives within the system. The system's Hot Spare Policy however does not apply to drives within FAST Cache. This means that if a system only contains 2 drives of a supported size and type, FAST Cache could still be created. It is always recommended to have spare drives within the system however. All drives not used in FAST Cache are subject to the Unity Hot Spare Policy of 1 spare drive for every 31 drives.

COMPARISON OF FAST CACHE AND FAST VP

The EMC Unity FAST software contains both FAST Cache and FAST VP, and both have been designed to take advantage of Flash technology. These software features complement each other, and both help improve the overall performance of the system.

FAST Cache is a global resource on the system which leverages Flash to provide high throughput and low response times. The most heavily accessed chunks of data not already residing on Flash are promoted into FAST Cache, which boosts performance for active workloads. FAST Cache absorbs IO bursts to the system, which helps reduce I/O workloads on the Pool.

FAST VP helps to optimize TCO by relocating data across storage Pools to meet the demands of workloads over time. As data ages, and the activity level for data reduces, data is relocated to provide capacity for active workloads.

The table below compares FAST Cache and FAST VP.

Table 5. FAST Cache and FAST VP Comparison.

CHARACTERISTIC	FAST CACHE	FAST VP
FUNCTION	Utilizes SAS Flash 2 Flash drives to create a large capacity secondary cache on the system.	Allows a storage resource to take advantage of multiple drive types within a Pool.
GRANULARITY	64 KB chunks of data.	256 MB slices of data.
OPERATION	Frequently accessed data on spinning drive is copied to FAST Cache for increased performance.	Data is physically moved between tiers of a Pool based on the tiering policy and a statistical analysis of the activity level of the slices.
TIMING	Always active. Promotes frequently accessed data residing on spinning drive to FAST Cache.	Relocations occur during a user specified FAST VP relocation window, or started manually.
ANALYSIS	Real-time monitoring of activity to data within a Pool.	Analysis of statistics information occurs every hour. This data is used to determine which data needs to be relocated.

When FAST Cache and FAST VP are used together they deliver improved TCO for the system and high performance and efficiency. As FAST Cache is a global resource, highly active data from all Pools can utilize FAST Cache. Bursts of data sent to Pool HDDs can be handled by FAST Cache, while FAST VP optimizes drive utilization and efficiency within the Pool. With highly active data on FAST Cache, FAST VP can prioritize the placement of slices for data accessed from the Pool. FAST Cache is a cost efficient way to add Flash to a configuration.

FAST VP works on a schedule to optimize the storage within a Pool. Even if a burst of I/O is seen, no slice movements occur until the relocation window. If the activity is frequent enough to cause a promotion, the data will be promoted to FAST Cache. FAST VP only monitors I/O which reaches the drives within the Pool. I/O handled by FAST Cache does not affect the analysis done by FAST VP. However, I/O activity due to FAST Cache page cleaning or flushing is monitored and weighed like normal I/O. If the activity of the slice is hot enough, the slice may relocate to a higher tier.

INTEROPERABILITY

FAST Cache and FAST VP are designed to work with other features of the system, as well as each other. The following features are compatible for use with FAST Cache and FAST VP.

UNIFIED SNAPSHOTS

The Unified Snapshots technology is fully compatible with FAST Cache and FAST VP. Unified Snapshots is a software feature which uses redirect-on-write technology to take Snapshots for LUNs, Consistency Groups, and File Systems on the system. Writes to a storage resources are redirected to new locations within the Pool. These blocks inherit the Tiering Policy and initial allocation policy of the resource the Snapshot was taken on. All data associated with a storage resource and its Snapshot can tier and leverage the different drive types within the Pool.

SYSTEM CACHE

The FAST Cache and FAST VP algorithms only track and base decisions on I/O which reach the underlying storage for a given resource. I/O serviced by System Cache are not factored into the promotion algorithms of FAST VP, or the statistical tracking by FAST VP. As System Cache is servicing the I/O, unnecessary promotions or slice relocations are avoided entirely. System Cache page cleaning or flushing operations are counted as these are I/O to the drives in the Pool.

FAST Cache and System Cache both have specific functions which complement each other. System Cache handles many workloads, such as high-frequency workloads, which are best handled by System Cache. System Cache also consolidates I/O requests where possible, which can be seen in prefetching of sequential data, or coalescing of sequential writes. This helps to reduce the amount of I/O sent to the Pools. Each feature helps to improve the overall performance of the system. The table below shows the differences between System Cache and FAST Cache.

Table 6. System Cache and FAST Cache Comparison.

CHARACTERISTIC	SYSTEM CACHE	FAST CACHE
POSITION	Closest to the CPU, Lowest Latency	Between System Cache and Pool drives
BEST SUITED FOR	Sequential I/O I/O larger than 64 KB Zero fill requests High-Frequency Access Patterns	Random I/O I/O smaller than 64 KB Data with a high locality
RESPONSE TIME	Nanosecond to microsecond response time	Microsecond to millisecond response time
OPERATION	Single memory region which services read and write requests	Single region which services read and write requests
CAPACITY	Limited in size, capacity based on model	Capacities scale to higher levels than System Cache. Maximums based on model.
GRANULARITY	Fixed 8 KB page size	Fixed 64 KB page size
AVAILABILITY	Memory modules are customer replaceable	Failing drives can proactively hot spare to free drives within the system. Faulted component is customer replaceable.
POWER FAILURE	Contents of System Cache are volatile. Battery Backup Units power the system until all dirty pages are copied to the mSATA drive.	Contents of FAST Cache are non-volatile and can withstand power loss.

CONCLUSION

The FAST software optimizes EMC Unity systems which reduces the Total Cost of Ownership and increases the overall performance of the system by efficiently utilizing the resources within the system. The FAST software optimizes the use of Flash drives, by leveraging them to service highly active data within the system. Utilizing Flash with FAST Cache avoids dedicating these drives to Pools, and allows them to be a global resource on the system. FAST Cache and FAST VP serve different purposes, and help to achieve storage efficiency in different ways.

While FAST Cache handles bursts of activity and unpredictable workloads, FAST VP tiers data over time, and continuously matches resource needs to tiers within the Pool. These features complement each other, as they work on different granularities, and have different frequencies in which they function. Implementing both FAST Cache and FAST VP on your storage system can improve performance, and reduce storage costs.

FAST Cache and FAST VP can be customized once, and left to run on the storage, with no other manual intervention needed. FAST Cache and FAST VP will efficiently manage the placement of data and leverage the drives in the system. Each of these features can be managed easily through Unisphere, Unisphere CLI, and REST API.

REFERENCES

The following white papers can be found on EMC Online Support:

- EMC Unity: Introduction to Unity Platform
- EMC Unity: Unisphere Overview
- EMC Unity: UnityVSA
- EMC Unity: Snapshots