

EMC FAST CACHE

A Detailed Review

Abstract

This white paper describes EMC® FAST Cache technology in CLARiiON®, Celerra® unified, and VNX™ storage systems. It describes the implementation of the FAST Cache feature and how to use it with Unisphere™ Manager and CLI.

October 2011

Copyright © 2011 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.

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

Part Number h8046.7

Table of Contents

Executive summary	3
Introduction	5
Audience.....	5
Terminology	5
Improved application performance and cost savings	6
Scope and characteristics	7
FAST Cache design concepts	8
FAST Cache comparison with storage system cache	9
FAST Cache components	10
Theory of operation.....	10
FAST Cache internal configuration.....	11
Failure handling	12
Management.....	12
Unisphere	12
Secure CLI.....	17
Unisphere Analyzer	17
FASTVP and FAST Cache.....	18
Best practices when you have a limited number of Flash drives.....	19
Configuration planning and usage guidelines	19
Application performance	23
Conclusion	23
References	24
Appendix A: FAST Cache configuration options	25

Executive summary

Since the original deployment of Flash technology in disk modules (commonly known as SSDs) by EMC in enterprise arrays, it has been our goal to expand the use of this technology throughout the storage system. The combination of high performance and the rapidly falling cost-per-gigabyte of Flash technology has led to the concept of a caching tier. A caching tier is a large-capacity secondary cache using enterprise Flash drives positioned between the storage processor's DRAM-based primary cache and

hard-disk drives. On EMC® CLARiiON®, Celerra® unified, and VNX™ storage systems, this feature is called EMC FAST Cache.

EMC FAST Cache works with FAST VP, which places data on the most appropriate storage tier based on how applications access various data chunks.

In CLARiiON and Celerra unified storage systems, starting with FLARE® release 30, you can use Flash drives as FAST Cache¹. VNX storage systems also support FAST Cache. FAST Cache extends the storage system's existing caching capacity for better system-wide performance. It does this by extending the functionality of DRAM cache by mapping frequently accessed data to Flash drives, which are an order of magnitude faster than mechanical drives (also known as hard disk drives), therefore boosting system performance. It also provides a much larger, scalable cache by using Flash drives, which provide very large capacities per drive as compared to DRAM capacities. FAST Cache capacities range from 73 GB to 2 TB, which is considerably larger than the available DRAM cache of existing storage systems.

At a system level FAST Cache helps make the most efficient use of Flash drive capacity. FAST Cache does this by using Flash drives for the most frequently accessed data in the storage system instead of dedicating Flash drives to a particular application. Configuring FAST Cache is a nondisruptive online process. It uses the existing memory-allocation interface and does not use host (server) cycles. FAST Cache is created in RAID-protected read/write mode, and the capacity options depend on the storage-system model and the number and type of installed Flash drives. You can create FAST Cache, enable it on storage volumes, and manage it through Unisphere™. No user intervention is required to enable applications to start seeing the performance benefits of FAST Cache. You can use it in existing CLARiiON CX4 and Celerra NS series storage systems (with a CLARiiON CX4 back end), and in VNX series storage systems. FAST Cache can also be used for RAID-group-based LUNs as well as pool-based LUNs.

¹ The VNX5100™ storage systems allow you to use either the FAST Cache or Virtual Provisioning™ feature. All other models allow you to use both features at once.

Introduction

This white paper provides an introduction to the FAST Cache feature on EMC CLARiiON, Celerra unified, and VNX series storage systems. If a particular chunk of data is accessed frequently by the user application, FLARE automatically promotes that chunk into FAST Cache by copying it from the hard disk drives into Flash drives. Subsequent I/O access to the same chunk is serviced at Flash-drive response times, thus boosting storage-system performance. If the access frequency of this chunk of data decreases and other chunks need to be promoted to FAST Cache, the older data is moved out from the FAST Cache. Details about the FAST Cache algorithm and basic FAST Cache operations are discussed in the “FAST Cache design concepts” section of this white paper.

Salient features of FAST Cache and its scope are discussed in the “Scope and characteristics” section. Basic considerations for configuration planning and best practices for deployment are discussed in the “Configuration planning and usage guidelines” section, and some application performance results with FAST Cache are discussed in the “Application performance” section.

Audience

This white paper is intended for EMC customers, partners, and employees who are considering the use of the FAST Cache feature in CLARiiON, Celerra unified, and VNX storage systems. It assumes familiarity with CLARiiON and VNX storage systems and its management software. A good introductory paper is EMC CLARiiON Storage System Fundamentals for Performance and Availability available on Powerlink®.

Terminology

- **Cache page** – The smallest unit of allocation inside the cache, typically a few kilobytes in size.
- **Cache clean page** – A page of FAST Cache that is valid, and contains a copy of data that has been synchronized with the user LUN.
- **Cache dirty page** – A page of FAST Cache that is valid, contains the most recent copy of data, but has not yet been synchronized with the user LUN.
- **Cache valid page** – A page of FAST Cache that contains a representation of data resident on an associated user LUN.
- **Cache invalid page** – A page of FAST Cache that does not contain any usable data; it does not represent data resident on any user LUN.
- **Cache warming** – The process of promoting new pages into FAST Cache after they have been created, or a change in the application access profile that begins to reference an entirely new set of data.
- **Chunk** – A portion of data in a particular address range.

- **DRAM cache** – A storage-system component that improves performance by transparently storing data in very fast storage media (DRAM), so that requests for that data can be served faster.
- **FAST Cache promotion** – The process of copying data from the back-end user LUN, on which FAST Cache has been enabled, to a FAST Cache page.
- **FAST Cache write-back** – The process of copying data from a FAST Cache page to a back-end hard-disk-based LUN.
- **Flash drive** – A data storage device that uses solid state media to store data. Because it does not have moving parts, this drive provides extremely low response times and high IOPS compared to rotating hard-disk drives (HDDs).
- **Hard disk drive (HDD)** – A data storage device that stores data on magnetic surfaces that rotate at various speeds.
- **Hot spot** – A busy area on a LUN.
- **Locality of reference** – The concept that logical blocks located *close* to each other will be accessed at around the same time.
- **Logical block address** – An addressing scheme that specifies the location of *blocks* of data on storage devices.
- **Memory map** – An array of bits in which each bit represents a FAST Cache *page*. This map shows which pages are in FAST Cache, and where they are located in FAST Cache. A copy of the memory map resides in DRAM cache; this ensures that the pages are accessed at memory speeds.
- **Pool** – A group of disk drives used by pool LUNs. There may be zero or more *pools* on a system. Disks may only be a member of one pool. Pool disks cannot be in RAID groups.
- **Thin LUN** – A logical unit of storage created on a pool where *physical space* allocated on the storage system may be less than the user capacity seen by the host server.
- **Thick LUN** – A logical unit of storage created on a pool where *physical space* reserved on the storage system is equal to the user capacity seen by the host server.

Improved application performance and cost savings

FAST Cache allows you to leverage the lower response time and better IOPS of Flash drives without dedicating Flash drives to specific applications. Flash drives need not be dedicated to specific applications anymore. This technology adds to the available storage-system cache (up to 2 TB read/write FAST Cache in CX4-960/NS-960 and VNX7500™ storage systems). It adjusts to a hot spot anywhere in the array, so you no longer need to analyze specific application requirements. It provides better performance to all applications in the storage system while using less Flash drives.

One of the major benefits of using FAST Cache is the improved application performance, especially for workloads with frequent and unpredictable large increases in I/O activity. The part of an application's working dataset that is frequently accessed is copied to the FAST Cache, so the application receives an immediate performance boost. FAST Cache enables applications to deliver consistent performance by absorbing heavy read/write loads at Flash drive speeds.

Another important benefit is improved total cost of ownership (TCO) of the system, which is realized because the load on back-end hard disk drives is reduced. FAST Cache copies the busy subsets of large storage capacities to Flash drives; as a result the busiest areas of many LUNs utilize a small set of Flash drives. This allows a large set of slower drives to deliver the performance typically provided by faster drives. Over a period of time faster Fibre Channel drives can be reduced or replaced with slower Fibre Channel or SATA drives, while maintaining the same application performance. This improves the financial (IOPS/dollar) and power (IOPS/KWH) efficiency of the storage system.

FAST Cache can be used by existing CLARiiON CX4 customers just by upgrading to FLARE release 30 and adding Flash drives, if necessary. This feature is very simple to set up and use. If you are using Celerra unified storage systems, the storage system needs to be upgraded to version 6.0. Celerra Startup Assistant (CSA) detects Flash drives. Before CSA uses the Flash drives for RAID group LUNs, it asks if you want to use the Flash drives for configuring FAST Cache on the storage system. You can then use Unisphere to configure FAST Cache on the block side before continuing with CSA. All models in the VNX series support FAST Cache.

In Unisphere, creating FAST Cache is as easy as and similar to creating a storage pool. A set of Flash drives is selected either automatically by the storage system or manually by the user, then Unisphere uses those drives to create FAST Cache. By default, after you have installed the FAST Cache enabler, FAST Cache is enabled on new RAID group LUNs and storage pools. When FAST Cache has been created, the storage system automatically promotes frequently accessed chunks of data to FAST Cache to improve application performance.

As with any cache, the results that are obtained through this type of architecture may vary, depending on the I/O profile of the application.

Scope and characteristics

FAST Cache acts as an *application accelerator* by providing Flash-drive-level performance to the data that is accessed most frequently. It also reduces the load on other storage system resources. Some of the salient features of FAST Cache are:

- It can be used on all CLARiiON CX4 storage systems running FLARE release 30 or later. If you have a Celerra unified storage system, Data Movers need to be running Celerra version 6.0 or later (along with a back-end CX4 running FLARE 30) to see FAST Cache benefits. All models in the VNX series support FAST Cache.

- It leverages the strong history of Flash drives and uses existing Flash drives, Disk Array Enclosures, and back-end buses. For example, the Flash drives do not need to be inserted in a particular drive slot or on dedicated back-end buses.
- A single FAST Cache instance per storage system serves as a system-wide resource. You can create and destroy the FAST Cache while the storage system and applications are online.
- When you create a FAST Cache, the Flash drives are configured as *RAID 1 mirrors* (redundant).
- Once a Flash drive has been used to configure FAST Cache, the entire Flash drive is dedicated to FAST Cache.
- Unisphere allows you to create FAST Cache in different configurations for different storage models; these are listed in “Appendix A.”
- FAST Cache can be applied to any LUN in a storage system. If a LUN (or component LUNs of a metaLUN) is created in a RAID group, FAST Cache is enabled or disabled at the individual LUN (or component LUN) level. If a LUN is created in a storage pool, FAST Cache must be configured at the pool level.
- FAST Cache is disabled for all LUNs and storage pools that were created before the FAST Cache enabler was installed (through an NDU process). After installing the FAST Cache enabler, the *existing* LUNs and storage pools will have FAST Cache disabled, but FAST Cache will be enabled (by default) on *new* LUNs and storage pools. Unisphere can then be used to enable or disable FAST Cache on any RAID group LUN or storage pool.
- If a LUN has been created on Flash drives, FAST Cache is disabled by default. You can use Unisphere to enable or disable FAST Cache on any Flash drive LUN.
- Once the FAST Cache enabler has been installed on the storage system, there is a separate process to create FAST Cache using dedicated Flash drives. Creating or destroying FAST Cache has no effect on the FAST Cache setting on LUNs and storage pools; it just determines whether caching actually occurs when FAST Cache is set up.
- Once FAST Cache has been created, expansion or shrink operations require disabling the cache and re-creating the FAST Cache.
- FAST Cache consumes a portion of the storage system memory that was formerly available for read or write cache. The amount of memory consumed is dependent on the storage system model and FAST Cache size.
- Flash drives installed in vault drive locations cannot be used to create FAST Cache. This is restricted through FLARE.

FAST Cache design concepts

FAST Cache was designed to follow these basic principles:

- It should be large enough to be able to contain a substantial part of the highly active portion of the application working set. FAST Cache can scale up to 2 TB in size in the CX4-960/NS-960.

- It should be fast enough to provide an order of magnitude improvement in performance to the application under certain conditions. Once the data has been promoted to the FAST Cache, any I/O to those chunks of data is serviced directly from the Flash drives at much lower response times and higher IOPS.
- It should keep the most heavily accessed data on the Flash drives for as long as possible. When new data needs to be copied into FAST Cache, a *Least Recently Used (LRU)* algorithm decides which data should be moved out of FAST Cache. This ensures that the frequently accessed data stays in FAST Cache for as long as possible.

FAST Cache comparison with storage system cache

FAST Cache is semi-conductor-based storage technology. It provides a large-capacity secondary tier of Flash-memory-based caching between the storage system’s fast, limited-capacity DRAM cache, and slower, higher-capacity mechanical hard disk drives. Table 1 compares DRAM memory with FAST Cache.

Table 1. Comparison of DRAM memory and FAST Cache

Characteristic	DRAM Cache	FAST Cache
Position	It is closest to the CPU, and has the lowest latency.	It is a step further from the CPU and is slower than DRAM cache.
Granularity	It has very high granularity, which is effectively the I/O size. The cache page size is user-configurable and can vary from 2 KB to 16 KB.	It operates in extents of 64 KB granularity.
Upgradeability	It is not upgradeable.	It is upgradeable in models and the options depend on storage system model and type of Flash drive.
Operation	It has separate areas for read and write operations.	It has a single area that serves read and write operations.
Capacity	It is limited in size compared to FAST Cache.	It can scale up to much larger capacity.
Scope	It services FAST Cache LUNs as well as other LUNs.	It allows you to enable FAST Cache on selected RAID group LUNs or storage pools.
Response time	Response time is in the order of nanoseconds to microseconds.	Response time is in the order of microseconds to milliseconds.
Availability	In case of failure, replacement requires service by qualified personnel.	In case of failure, another Flash-drive hot spare automatically replaces the failing drive, and the faulted component is customer-replaceable.
Power failure	Its contents are volatile; they cannot withstand a power loss.	Its contents are non-volatile and can withstand a power loss.

FAST Cache components

The two main blocks of FAST Cache are:

- **Policy engine** – This module manages the flow of I/O through the FAST Cache. When a chunk of data on a back-end LUN is accessed frequently it is temporarily copied to Flash drives; it is copied back to HDDs if other data becomes more heavily used. The policy engine module decides when these operations should take place. This module also maintains statistical information about the data access patterns. The policies defined by this module are system-defined and may not be changed by the user.
- **Memory map** – This map tracks extent usage and ownership in 64 KB chunks of granularity. It maintains information on the state of 64 KB chunks of storage and their contents in FAST Cache. A copy of the memory map is stored in DRAM memory, so when FAST Cache is installed in the storage system the existing SP read and write cache may need to be temporarily disabled to allocate space for the FAST Cache memory map. Once the memory map has been created, the resized SP read and write caches are automatically enabled by the storage system. The size of the memory map increases linearly with the size of FAST Cache being created. A copy of the memory map is also mirrored to the Flash disks to maintain data integrity and high availability of data.

Theory of operation

If FAST Cache is installed in the storage system and enabled on the LUN and/or pool, incoming I/O from the host application is checked against the FAST Cache memory map to determine if the I/O is for a data block that has already been promoted into FAST Cache or not. If the chunk *is not* in FAST Cache, the I/O request follows the same path it would follow if the storage system did not have FAST Cache.

However, if the data chunk *is* in FAST Cache, the policy engine redirects the I/O request to the FAST Cache. If the host I/O request is for a read operation, and the target data is in the DRAM cache, the data is read from the DRAM cache. If the data is not in DRAM cache, the data is read from the FAST Cache Flash drives; it is also put in the DRAM cache (as it would be with reads from HDDs).

If the host I/O request is a write operation for a data chunk that is in FAST Cache, the DRAM cache is updated with the new “write,” and an acknowledgement is sent back to the host. If DRAM write cache is not disabled for the HDD LUN, write operations usually go through the DRAM cache; the host data is not written directly to the FAST Cache Flash drives. When data needs to be moved out of DRAM cache, it is written to FAST Cache Flash drives. Since the data is written to Flash drives instead of HDDs, this operation is very fast and may help limit the number of dirty pages in the DRAM write cache. It should be noted that even when FAST Cache is installed in the storage system, the I/O operations are serviced directly from DRAM cache whenever possible.

Data that is in FAST Cache and is used less frequently is copied back to the HDD from FAST Cache. This chunk of data must be re-promoted into FAST Cache to be serviced by the FAST Cache. Data on HDD that becomes busy is promoted to FAST Cache,

which is an asynchronous process (does not block host I/O). Data promotion into FAST Cache depends on the number of accesses (read and/or write) within a 64 KB chunk of storage, and is not dependent on whether the data already exists in the DRAM cache or not.

For example, assume a storage system receives an I/O request from an application right after FAST Cache is created on the storage system. In this scenario, the FAST Cache memory is empty since nothing has been copied into the FAST Cache.

- When the first I/O is sent by the application, the FAST Cache policy engine looks for an entry in the FAST Cache memory map for the I/O's data chunk. Since the memory map is empty at this stage, the data is accessed from the HDD LUN. This is called a *FAST Cache Miss*. (We have found that the performance overhead of checking the memory map for every access to a FAST Cache enabled LUN is very minimal.)
- If the application frequently accesses data in a 64 KB chunk of storage, the policy engine copies that chunk from the hard-disk LUN to FAST Cache. The memory map is updated to indicate that the data chunk is now resident in FAST Cache. This operation is called *Promotion* and this period is called the *warm-up* period for FAST Cache. The storage system controls this activity to make sure that the overhead for promotion does not exceed a specified percentage of the storage system capability.
- When the application accesses this data again, the policy engine sees that it is in the FAST Cache. This is called the *FAST Cache Hit*. Since the data is accessed directly from the Flash drives, the application gets very low response times and high IOPS. If a substantial part of the working set gets promoted to FAST Cache over a period of time, applications can see higher average performance even with lower performing HDDs in the back end.

In certain situations, data is copied from FAST Cache back to the back-end HDDs. This is called a *Write Back* operation. The rate at which data blocks are written back to the HDD LUNs is controlled so that the process does not significantly impact storage system performance. Write Back operations happen in the event that a FAST Cache promotion is scheduled but there are no free or clean pages available in the FAST Cache. A "dirty" page will then be copied from the FAST Cache and written to the HDD LUN to make room for the new data. The LRU algorithm determines which data block is least likely to be accessed again and this block is processed to make room for the new promote.

FAST Cache internal configuration

When FAST Cache is created, the Flash drives are configured as RAID 1 mirror pairs. The application gets the acknowledgement for an I/O operation once it has been serviced by the FAST Cache. In write operations, the FAST Cache page is marked as dirty and an acknowledgement is sent back to the host. Read operations are also directly serviced from the FAST Cache if the corresponding data chunks have already been promoted into FAST Cache. The FAST Cache algorithm is designed so that the workload is spread evenly across all Flash drives in the FAST Cache.

Failure handling

Traditional CLARiiON global hot sparing algorithms are used for the Flash drives configured as FAST Cache. Global hot sparing provides automatic, online rebuilds of redundant RAID groups when any of the group's drives fail. EMC has further advanced this functionality with *proactive* hot sparing. Proactive hot sparing recognizes when a drive is nearing failure and preemptively copies the drive content before it fails. The combination of these features minimizes each RAID group's vulnerability to additional drive failures; this prevents data loss. For performance reasons, only Flash drives configured as global hot spares replace failing Flash drives in the FAST Cache. Also, a hot spare can be shared between the Flash drives used for FAST Cache and Flash drives used elsewhere in the storage system. For more details on this feature, refer to the *EMC CLARiiON Global Hot Spares and Proactive Hot Sparing* white paper available on Powerlink.

In the event of loss of a Flash drive in FAST Cache, the normal FLARE RAID recovery routine attempts to initiate a repair by establishing a hot-spare Flash drive and rebuilding the data. If a hot spare is not available, then the FAST Cache remains in Degraded mode, and the cache-page cleaning algorithm increases the rate at which FAST Cache pages are copied back from FAST Cache Flash drives to HDDs. In this case, only read operations are allowed from FAST Cache; this reduces the potential exposure to data loss of a drive failure in a non-redundant RAID group. In this state, write operations from the application are serviced by the HDD LUNs in the back end.

Management

You can use Unisphere or Secure CLI to create, manage, and monitor FAST Cache. Unisphere details can be found in the *EMC Unisphere: Unified Storage Management Solution* white paper available on Powerlink. The following sections discuss the parts of Unisphere and Secure CLI that pertain to FAST Cache.

Unisphere

The **System** tab in Unisphere has links on the right-hand side for System Properties and Manage Cache. Both these links open the System Properties window (Figure 1).

To create FAST Cache, click the **FAST Cache** tab in the Storage System Properties window to view FAST Cache information. If FAST Cache has not been created on the storage system, the **Create** button in the bottom of the dialog box is enabled. The **Destroy** button is enabled when FAST Cache has been created.




Figure 1. Storage System Properties dialog box

When the FAST Cache has been created, the State, Size and RAID Type fields are updated to reflect the FAST Cache configuration details. The RAID Type field is displayed as RAID 1 when the FAST Cache has been created. Clicking **Create** opens the Create FAST Cache dialog box (Figure 2).




Figure 2. Create FAST Cache dialog box

FAST Cache can be created in certain configurations, depending on the storage system model, and number and size of Flash drives installed in the storage system. These criteria are used to present you with the available options for your configuration. For example, if a sufficient number of Flash drives is not available, Unisphere displays an error message and FAST Cache cannot be created. The number of Flash drives can also be chosen in this screen. The bottom portion of the screen shows the Flash drives that will be used for creating FAST Cache. You can choose the drives manually by selecting the **Manual** option.

If the LUN is created in a RAID group, you can enable or disable FAST Cache at the LUN level. It is enabled by default if the FAST Cache enabler is installed on the storage system. Figure 3 shows how you can configure the FAST Cache under the **Advanced** tab in the Create LUN dialog box. If the LUN has already been created in a RAID group, click the **Cache** tab in the LUN Properties dialog box to configure FAST Cache (shown in Figure 4).

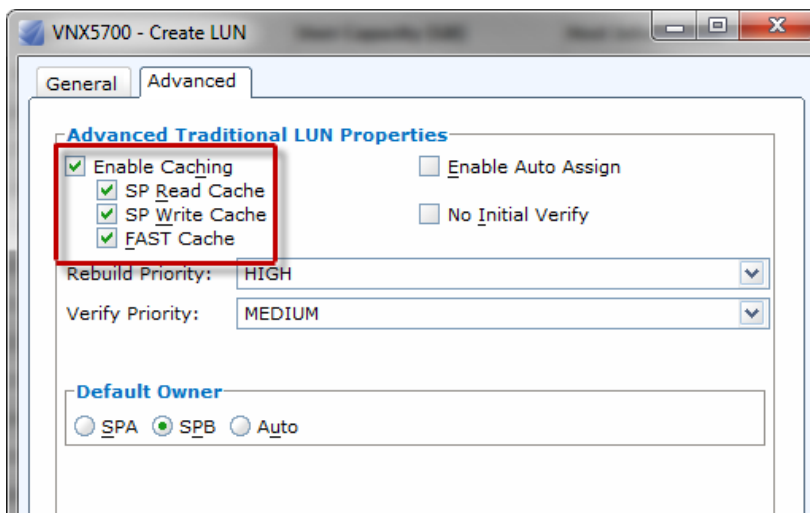


Figure 3. Advanced tab in the Create LUN dialog box

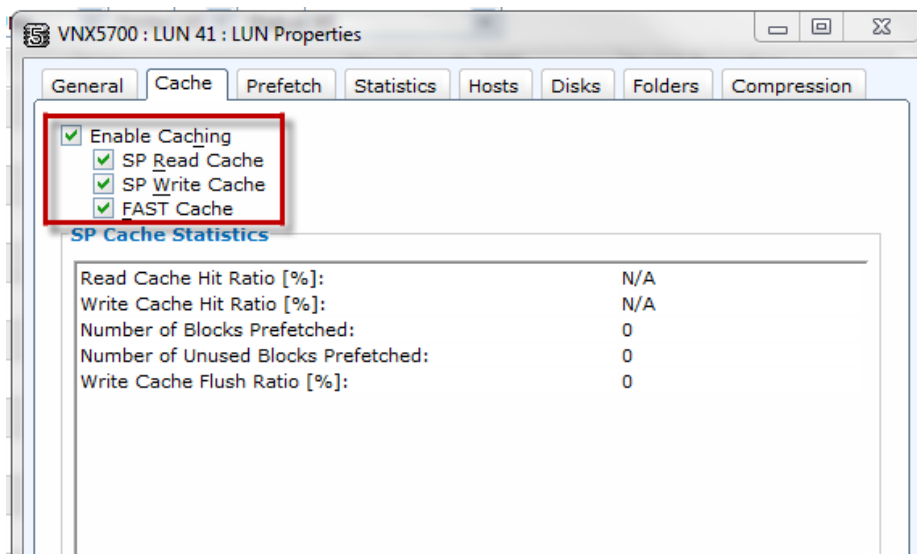


Figure 4. Cache tab in the LUN Properties dialog box

If a LUN is created in a storage pool, you can only configure FAST Cache for that LUN at the storage pool level. In other words, all the LUNs created in the storage pool will have FAST Cache enabled or disabled. You can configure them under the **Advanced** tab in the Create Storage Pool dialog box shown in Figure 5. If the storage pool has already been created, then you can use the **Advanced** tab in the Storage Pool Properties dialog box to configure FAST Cache (see Figure 6).




Figure 5. Advanced tab in the Create Storage Pool dialog box

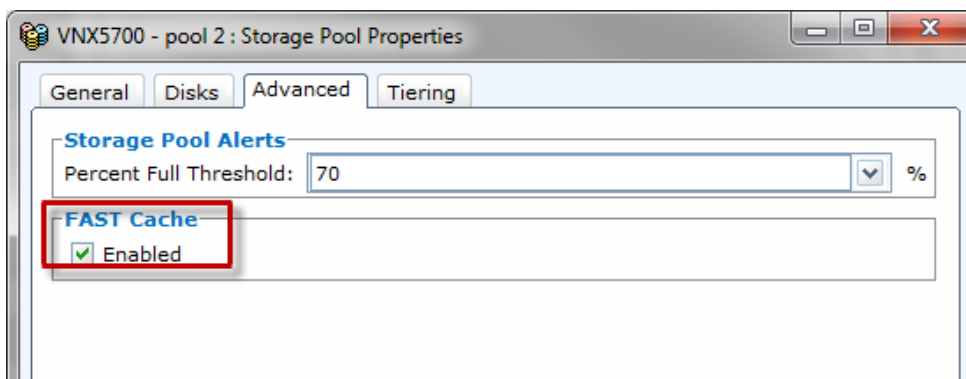


Figure 6. Advanced tab in the Storage Pool Properties dialog box

You can display FAST Cache properties in any Unisphere table (for example, the LUNs table) by right-clicking the table header, and selecting **Choose Columns**. You can also click the Tools icon at the top-right corner of the table and select **Choose Columns**. This opens a dialog box in which you can select **FAST Cache**. This is shown in Figure 7. The FAST Cache property will be displayed for every entry in the table.

The screenshot shows the 'LUNs' table in Unisphere. The table has columns for Name, FAST Cache, State, ID, User Capacity (GB), and Host Information. The 'FAST Cache' column is highlighted with a red box. A context menu is open over the table, with 'Choose Columns' selected. The table contains the following data:

Name	FAST Cache	State	ID	User Capacity (GB)	Host Information
LUN 28	On	Ready		1.000	
LUN 46	On	Ready		10.000	
LUN 48	On	Ready		20.000	
LUN 52	On	Ready	52	3.000	

Figure 7. Choosing to display FAST Cache information in Unisphere

Secure CLI

The management functions described in the previous section are also available with Unisphere CLI. CLI commands for FAST Cache include:

- **Create FAST Cache:** `cache -fast -create`
- **Destroy FAST Cache:** `cache -fast -destroy`
- **Get FAST Cache information from the storage system:** `cache -fast -info`
- **Configure FAST Cache when creating a RAID group LUN:**
`bind ... -fastcache 0|1`
- **Enable or disable FAST Cache on a RAID group LUN:**
`chglun ... -fastcache 0|1`
- **Get FAST Cache configuration information from a RAID group LUN:**
`getlun ... <-fastcache> <-all>`
- **Configure FAST Cache when creating a storage pool:**
`storagepool -create ... -fastcache on|off`
- **Configure FAST Cache on an already created storage pool:**
`storagepool -modify ... -fastcache on|off`
- **Get FAST Cache configuration information from a storage pool:**
`storagepool -list ... <-fastcache> <-all>`

Unisphere Analyzer

After you install FAST Cache, Unisphere Analyzer gathers FAST Cache statistics to help you monitor FAST Cache performance. To view these statistics, you must enable Analyzer's Advanced mode. To enable Analyzer's Advanced mode, follow these steps:

In Unisphere, click the System tab.

1. Click Monitoring and Alerts.
2. Click Statistics for Block.
3. Click Customize Charts.
4. Click the **General** tab.
5. Select the **Advanced** checkbox.
6. Click **OK** to apply the settings.

The following FAST Cache statistics are available at the storage processor level:

- FAST Cache Dirty Pages (%)
- FAST Cache MBs Flushed (MB/s)

The following FAST Cache statistics are available for RAID group LUNs and storage pools:

- FAST Cache Read Hits/s
- FAST Cache Read Misses/s
- FAST Cache Read Hit Ratio
- FAST Cache Write Hits/s
- FAST Cache Write Misses/s
- FAST Cache Write Hit Ratio

FAST VP and FAST Cache

FAST VP is a new feature introduced in FLARE release 30 and is also supported in the VNX series. FAST VP performs storage tiering for 1 GB chunks of data at a sub-LUN level in pools that contain multiple drive types. FAST VP automatically moves more active chunks (data that is more frequently accessed) to the best performing storage tier, and it moves less active chunks to a lower performing (and less expensive) tier for a better TCO. For more details on this feature, refer to the EMC FAST VP for Unified Storage Systems white paper available on Powerlink. Table 2 compares FAST VP and FAST Cache feature.

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

FAST Cache	FAST VP
Allows Flash drives to be used to extend the existing caching capacity of the storage system.	Allows a single LUN to leverage the advantages of multiple drive types through the use of storage pools.
Granularity level is 64 KB.	Granularity level is 1 GB.
Data that is accessed frequently is <i>copied</i> from HDDs to Flash drives.	Data is <i>moved</i> between different storage tiers based on weighted-average-of-access statistics collected over a period of time.
This should be used when workload changes are unpredictable and very dynamic, and require a quick response time.	This should be used when workload pattern changes are predictable and relatively low.
Constantly promotes HDD data to FAST Cache or vice-versa. There are no polling or relocation cycles.	Data movement occurs in scheduled or manually invoked relocation windows.
Calculation to decide which data needs to be promoted to FAST Cache is performed continuously.	Calculation to decide which portion of data needs to be moved is performed once every hour.

FAST Cache and the FAST VP feature can be used together to yield high performance and TCO from the storage system. For example, you can use Flash drives to create FAST Cache, and use the FAST VP feature for storage pools consisting of SAS and NL-SAS disk drives. From a performance point of view, FAST Cache provides an immediate performance benefit to bursty data, while FAST VP moves more active data to SAS drives and less active data to NL-SAS drives. From a TCO perspective, FAST Cache can service active data with less Flash drives, while FAST VP optimizes disk utilization and efficiency with SAS and NL-SAS drives.

As a general rule, you should use FAST Cache in cases where storage system performance needs to be improved immediately for burst-prone data. On the other hand, FAST VP optimizes TCO by 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. These features complement each other and, used together, can improve performance *and* TCO.

FAST Cache works with FAST VP to ensure that resources are not wasted on unnecessary tasks. For example, if FAST VP moves a chunk of data to Flash drives, FAST Cache will not promote that chunk of data into FAST Cache, even if the FAST Cache criteria is met for promotion. This ensures that resources are not wasted by copying data from one Flash drive to another.

Best practices when you have a limited number of Flash drives

If you have a limited number of Flash drives and an option to use them either for FAST VP or FAST Cache, EMC recommends that you use the Flash drives to create FAST Cache, and use the remaining Flash drives in a FAST VP-enabled storage pool. FAST Cache is global in nature and benefits all the LUNs and pools in the storage system, instead of only benefiting the storage pool where the Flash drives are used for FAST VP.

FAST Cache has a higher granularity than FAST VP; FAST Cache uses 64 KB chunks while FAST VP uses 1 GB chunks. As a result, you will see higher performance benefits and faster reaction time for changing usage patterns when using Flash drives for FAST Cache. The downside of higher parity overhead in FAST Cache because of the RAID 1 architecture is offset by improved performance of the DRAM cache. This is because of the faster write operations from DRAM cache to FAST Cache Flash drives as compared to HDDs.

Configuration planning and usage guidelines

You should consider the following guidelines when planning the configuration and deployment of FAST Cache feature in a storage system:

- The VNX5100 storage systems allow you to use either the FAST Cache or Virtual Provisioning feature. If the Virtual Provisioning feature enabler is installed on the storage system, you cannot use FAST Cache. If you are using FAST Cache, you

cannot use Virtual Provisioning On all other models you can use both features at the same time.

- FAST Cache and FAST VP can be used in the same pool, but we highly recommend that you use FLARE 04.30.000.5.517 or later for storage systems with CLARiiON CX4 backend and VNX Operating Environment for Block 05.31.000.5.502 or later for VNX storage systems.
- For FAST Cache, different storage system models require different numbers and sizes of Flash drives. When you create FAST Cache, Unisphere presents the configurations for FAST Cache that are available in your storage system. All acceptable configurations are listed in “Appendix A.”
- The performance for a LUN with FAST Cache will fall somewhere between HDD performance and Flash drive performance, depending on the access pattern of the workload.
- You should analyze application workload characteristics before deciding to use FAST Cache in a storage system. FAST Cache is not equally helpful for all kinds of I/O profiles. For example, sequential streams with large I/O size may not cause any FAST Cache promotions to occur because these I/O streams may not access the same 64 KB block of data multiple times, which is what is needed for FAST Cache promotions to occur. Sequential applications such as backups typically have I/Os that are larger than 64 KB, which will result in only one access in one 64 KB chunk. This will not result in any FAST Cache promotions. Sequential workloads with small I/O sizes (for example, 8 KB sequential write) may result in excessive promotion activity for FAST Cache; this can cause drop in achievable performance of the primary workload on the storage system. Detailed analysis of application workloads to determine potential FAST Cache benefits can be done through EMC.
- The key to maximizing FAST Cache performance is locality of reference in the workload. Applications that access a small area of storage with very high frequency will benefit the most from using FAST Cache. Examples of this are database indices and reference tables. If the locality of reference is low, the storage system may get less benefit after promoting a data chunk into FAST Cache; very low locality will result in few or no promotions and thus no benefit.
- The FAST Cache policy engine has to track every I/O to calculate whether a block needs to be promoted to FAST Cache for LUNs that have this feature enabled. This requires some SP CPU processing cycles. Therefore, enabling the feature only for LUNs that will benefit from it results in optimal resource utilization. For example, enabling FAST Cache for LUNs that serve as ‘secondary’ mirror or as clone target LUN will gain no significant performance benefit. In the case of Thin or Thick LUNs, the FAST Cache would have to be disabled at a pool level, so it is best to have separate pools for LUNs that do not need FAST Cache. This gives a higher performance gain to the LUNs that really do need FAST cache.

- If the application “working set” is not very large and can fit in the FAST Cache, then after all the promotions have finished, all application I/Os will be serviced at Flash drive speeds. Working-set size for all the applications using the FAST Cache should be considered.

FAST Cache takes some time to *warm up* before it shows performance improvement. Warm-up time consists mostly of promotion operations in FAST Cache. This happens when the FAST Cache has just been created and is empty. This also happens when the working data set of the application has changed dramatically, and the current FAST Cache data is no longer being referenced. During this phase, the FAST Cache Hit rate is low, so the response time is more like that of an HDD LUN. As the FAST Cache Hit rate increases, the response times gradually shift to those of the Flash drives.

Among other things, the warm-up time will depend on the number and type of HDDs in the back end. For example, a setup of 80 SAS drives will have a shorter warm-up time than a setup with 20 SAS drives. Similarly, FAST Cache with SAS HDDs in the back end will warm up faster than with NL-SAS HDDs in the back end, because NL-SAS drives typically have a higher response time than SAS drives. When you design application layouts, it is important to remember that there is a warm-up time before stable FAST Cache performance is reached.

- Layered applications such as MirrorView™ and SnapView™ require private LUNs. These LUN are optimized for priority in the storage system’s write cache. For this reason, EMC recommends that you disable FAST Cache on MirrorView’s write intent log and SnapView’s clone private LUNs. This prevents unnecessary promotions into the FAST Cache.
- SnapView snapshots and related replication software such as MirrorView/A and SAN Copy™ (incremental sessions) require reserved LUNs. FAST Cache does not improve reserved LUN performance, but it is not a detriment to performance as it is with write intent logs and clone private LUNs. Disabling FAST Cache for reserved LUNs can help to minimize the overall FAST Cache workload. FAST Cache can be disabled at the LUN level if the reserved LUN is created in RAID groups. If they are created in a pool and there are other LUNs in the pool that need FAST Cache, reserved LUNs can be left with FAST Cache enabled. Please refer to the *EMC CLARiiON Reserved LUN Pool Configuration Considerations* white paper on Powerlink for more reserved LUN configuration guidance.
- When using SnapSure, EMC recommends that you not enable FAST Cache on the SavVol volumes. SavVol volumes are typically used for small-block sequential-write operations, and read operations from the checkpoint are usually satisfied by the production file system.
- If possible, the size of the FAST Cache should be chosen carefully, keeping future needs in consideration. If FAST Cache needs to be expanded or reduced in size after it has been created, it needs to be destroyed first and then re-created in the new configuration. Destroying a FAST Cache can be a long process depending on the amount of dirty pages currently in the FAST Cache. It also depends on the

storage system workload because the speed of FAST Cache destroy process is throttled internally by FLARE to make sure that application workload is not adversely affected.

- For FAST Cache to show performance improvement, the processing power of the storage system needs to be sufficient so that it does not create a bottleneck. If the bottleneck is the limited number of HDDs in the back end, then FAST Cache can help improve overall system performance by promoting the most frequently accessed data and thereby providing Flash drive class service to the application. In other words, FAST Cache attempts to offload commonly accessed I/O from spinning media, but does not improve the overall performance potential of the back-end SPs.
- You should consider utilization percentage of the Storage Processors (SP) when enabling FAST Cache on a large number of LUNs. The following guidelines can be used for this:
 1. If the SP utilization is less than 60%, you can enable FAST Cache on all LUNs (or pools) but it should be done in a phased manner. Divide the total number of LUNs into multiple sets, logically grouped by application workload, with each set representing a reasonable part of the total I/O workload on the array. Enable FAST Cache on each of these sets one at a time and monitor the SP utilization during this process. The SP utilization should stabilize before you enable it on the next set of LUNs or pools.
 2. If the SP utilization is between 60% and 80%, enable FAST Cache on one or two LUNs (or a single pool) at a time and monitor the SP utilization during this process. The SP utilization should stabilize before you enable it on any other LUNs or pools.
 3. If the SP utilization is more than 80%, do not enable FAST Cache on any LUNs or pools; contact your service representative before proceeding.
- Unisphere allows you to choose the Flash drives that are used to create FAST Cache. You may need to choose these drives manually to make sure that you spread the Flash drives across all back-end buses. This helps avoid the bandwidth limitation of a single bus. If all the FAST Cache Flash drives are on a single bus, then HDD LUNs with a high bandwidth requirement can be put on a different back-end bus.
- In unified Celerra environments, most file systems reside on multiple dvols (LUNs). FAST Cache should be enabled for every LUN underlying a given file system. Also, FAST Cache should be enabled for Celerra UxFS log LUNs. Although this data should always be cached in the SP given the high temporal locality of access to it, the UxFS log is very small and ensuring that the writes to it are cached can only benefit performance.

More detailed best practice guidelines can be found in the *EMC CLARiiON Best Practices for Performance and Availability* white paper available on Powerlink. Refer

to application-specific white papers for guidelines on using FAST Cache with those applications.

Application performance

EMC conducted application-specific tests with FAST Cache to determine potential performance benefits when this feature is used. Here is a summary of FAST Cache benefits with various applications:

- **VMware View™** – With linked-clone desktops, FAST Cache improves overall performance by reducing I/O accesses to hard-disk drives and directing them to Flash drives. Specific use cases include boot storms (up to 99 percent reduction in hard-disk I/O access), recompose operations (up to 70 percent reduction in hard-disk I/O access), and virus-scan operations (77 percent reduction in time required for a full scan of the desktop). Since less I/O accesses are done from the HDDs, the same number of users can be supported with 44 percent less HDDs.
- **Oracle** – With an OLTP workload in an Oracle 11g R2 environment, FAST Cache improved performance from 12,500 transactions per minute to 30,000 transactions per minute. At the same time, latency was reduced from 14 ms to 5 ms.
- **SQL Server** – In a SQL Server OLTP environment, FAST Cache doubled the number of users from 25,000 to 50,000 while keeping the same number of Fibre Channel drives in the back end. At the same time, response time improved from 40 seconds to 5 seconds and transactions per second improved from 1,300 to 2,400 when FAST Cache was used.
- **Celerra unified file services** – In a test configuration with a random NFS-based file serving workload on an NS-960, a 100-drive FC-based system with 2 TB of FAST Cache matched the performance of a 300-drive FC-based system with no FAST Cache. This solution provided a 46 percent TCO benefit on a cost-per-I/O basis and also delivered *a 65 percent reduction in power consumption*.

Details about these performance numbers, the lab setup, and best practice recommendations for using FAST Cache with specific applications can be found in separate white papers available on Powerlink.

Conclusion

FAST Cache enables the system cache to be expanded by using Flash drives as a midtier of cache. This allows the storage system to provide “Flash drive” class performance to the most heavily accessed chunks of data. FAST Cache absorbs I/O bursts from applications, thereby reducing the load on back-end hard disks. This helps improve the TCO of the storage solution. FAST Cache can be managed through Unisphere in an easy and intuitive manner. However, FAST Cache might not be a natural fit for all types of workloads, and the application I/O profile should be analyzed to determine the potential performance benefits.

FAST Cache works in a complementary way with FAST VP technology. Both technologies help place data segments on the most appropriate storage tier based on their usage pattern.

References

The following white papers are available on Powerlink:

- *EMC CLARiiON Storage System Fundamentals for Performance and Availability*
- *EMC CLARiiON Best Practices for Performance and Availability*
- *EMC FAST VP for Unified Storage Systems – A Detailed Review*
- *EMC Unisphere: Unified Storage Management Solution*
- *EMC CLARiiON Virtual Provisioning – Applied Technology*
- *An Introduction to EMC CLARiiON and Celerra Unified Platform Storage Device Technology*
- *EMC CLARiiON Reserved LUN Pool Configuration Considerations*
- *Leveraging EMC FAST Cache with Oracle OLTP Database Applications*
- *Performance Review with Microsoft Exchange 2010 and EMC CLARiiON CX4 Utilizing FAST, FAST Cache, and Compression*

Appendix A: FAST Cache configuration options

Table 3. FAST Cache configuration options in CLARiiON CX4 and Celerra unified storage systems

Model	FAST Cache capacity (GB)	Flash drive capacity (GB)	Number of drives
CX4-120/NS-120	100	100	2
	73***	73	2
CX4-240	200	100	4
	146***	73	4
	100	100	2
	73***	73	2
CX4-480/NS-480	800	200	8
	400*	200	4
	400*	100	8
	200	100	4
	146***	73	4
	100**	100	2
	73***	73	2
CX4-960/NS-960	2,000	200	20
	1,000	200	10
	800**	200	8
	400**	200	4
	400	100	8
	292***	73	8
	200**	100	4
	146***	73	4
	100**	100	2
	73***	73	2

* This configuration requires FLARE version 04.30.000.5.507 or later.

** This configuration requires FLARE version 04.30.000.5.512 or later.

*** This configuration is supported if it already exists, but configurations with 73 GB Flash drives are no longer sold.

Table 4. FAST Cache configuration options in VNX storage systems using 100 GB Flash drives

Model	FAST Cache capacity (GB)	Number of 100 GB Flash drives
VNX5100	100	2
VNX5300™	100	2
	200	4
	300	6
	400	8
	500	10
VNX5500™	100	2
	200	4
	300	6
	400	8
	500	10
	600	12
	700	14
	800	16
	900	18
	1,000	20
VNX5700™	100	2
	200	4
	300	6
	400	8
	500	10
	600	12
	700	14
	800	16
	900	18
	1,000	20
	1,100	22
	1,200	24
	1,300	26
	1,400	28
1,500	30	

VNX7500	100	2
	200	4
	300	6
	400	8
	500	10
	600	12
	700	14
	800	16
	900	18
	1,000	20
	1,100	22
	1,200	24
	1,300	26
	1,400	28
	1,500	30
	1,600	32
	1,700	34
	1,800	36
	1,900	38
	2,000	40
2,100	42	

The VNX5100 storage systems allow you to use either the FAST Cache or Virtual Provisioning feature. If the Virtual Provisioning feature is installed on the storage system, you cannot use FAST Cache. If you are using FAST Cache, you cannot use Virtual Provisioning. On all other models, you can use both features at the same time.

Table 5. FAST Cache configuration options in VNX storage systems using 200 GB Flash drives

Models	FAST Cache capacity (GB)	Number of 200 GB Flash drives
VNX5100	N/A*	N/A*
VNX5300	200	2
	400	4
VNX5500	200	2
	400	4
	600	6
	800	8
	1,000	10
VNX5700	200	2
	400	4
	600	6
	800	8
	1,000	10
	1,200	12
	1,400	14
VNX7500	200	2
	400	4
	600	6
	800	8
	1,000	10
	1,200	12
	1,400	14
	1,600	16
	1,800	18
	2,000	20

* For VNX5100, the only FAST Cache option is 100 GB (usable) and therefore 200 GB Flash drives cannot be used to configure FAST Cache on this platform.