

VNX[®] FAST Cache

A Detailed Review

Abstract

This white paper is an introduction to the EMC[®] FAST Cache technology in VNX[®] storage systems. It describes implementation of the FAST Cache feature and provides details of using it with Unisphere[®] and Secure CLI. Usage guidelines and major customer benefits are also included.

August 2012

Copyright © 2012 EMC Corporation. All Rights Reserved.

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

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

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

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

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

Part Number h8046.8

Table of Contents

Executive summary	4
Introduction	4
Audience.....	5
Terminology	5
Global FAST Cache and TCO	6
FAST Cache components	7
Theory of operation	8
FAST Cache Promotions	8
Flushing out of FAST Cache.....	10
Failure handling	11
Management	11
Best practices	17
Limitations.....	18
Conclusion	18
References	19
Appendix A: FAST Cache configuration options	20
Appendix B: FAST VP and FAST Cache	23
Appendix C: FAST Cache comparison with storage system cache	25

Executive summary

Since the original deployment of Flash technology in disk modules (commonly known as SSDs) in enterprise arrays, EMC® has expanded the use of this technology throughout the storage system. The combination of high performance and the rapidly falling cost-per-gigabyte of Flash technology led to the concept of a caching tier. A caching tier is a large-capacity secondary cache that uses enterprise Flash drives positioned between the storage processor's DRAM-based primary cache and hard-disk drives (HDD). On EMC VNX® storage systems, this feature is called EMC FAST Cache.

FAST Cache extends the storage system's existing caching capacity for better system-wide performance. It achieves this by extending the functionality of the DRAM cache by copying frequently accessed data to Flash drives, which are faster than HDDs, therefore boosting system performance. Flash drives also provide a much larger, scalable cache than the DRAM cache. FAST Cache capacities range from 100 GB to 2 TB, which is considerably larger than the available DRAM cache of existing storage systems.

At a system level, FAST Cache makes the most efficient use of Flash drive capacity. FAST Cache achieves this by using Flash drives for the most frequently accessed data in the storage system instead of dedicating the drives to a particular application. Configuring FAST Cache is a nondisruptive online process that uses the existing memory-allocation interface but 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 see the performance benefits of FAST Cache. You can use it in existing VNX series storage systems¹ running VNX Block OE Release 31 or higher. FAST Cache can be used for RAID-group-based LUNs and pool-based LUNs.

Introduction

This white paper provides an introduction to the FAST Cache feature. The FAST Cache helps realize performance benefits of Flash drives system wide, following is a high level overview of how it works. If a particular chunk of data is accessed frequently by the user application, VNX 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 of the FAST Cache. Details about the FAST Cache algorithm and basic FAST Cache operations are discussed in the "Theory of Operations" section of this white paper.

¹ The VNX5100 storage systems allow you to use either the FAST Cache or Thin Provisioning™ feature, but not both. All other models allow you to use both features at once.

Audience

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

Terminology

- **Cache page**—The smallest unit of allocation inside the FAST 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 and 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 that resides on an associated user LUN. A valid page may be either in the clean or dirty state.
- **Cache invalid page**—A page of FAST Cache that does not contain any usable data. It does not represent the data resident on any user LUN.
- **Cache warming**—The process of copying new pages into FAST Cache after FAST Cache has 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 (64 KB).
- **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 are served faster.
- **Extent**—A set of adjacent physical blocks.
- **FAST Cache promotion**—The process of copying data from the back-end user LUN, on which FAST Cache is 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, a Flash 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 and rotates at various speeds.
- **Hot spot**—A busy area on a LUN.
- **Locality of reference**—The concept that logical blocks located close to each other are accessed at approximately the same time and repeatedly.

- **Logical block address**—An addressing scheme that specifies the location of blocks of data on storage devices.
- **Memory map**—An array of addresses in which each bit represents a FAST Cache page. This map shows which pages are in FAST Cache and where they are located. 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 used in RAID groups.
- **Thin LUN**—A logical unit of storage created on a pool where physical space consumed by 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 consumed on the storage system is equal to the user capacity seen by the host server.

Global FAST Cache and TCO

FAST Cache allows you to leverage the lower response time and better IOPS of Flash drives without dedicating Flash drives to specific applications. This technology supplements the available storage-system cache (adding up to 2 TB read/write FAST Cache in VNX7500 storage systems; see Appendix A: FAST Cache configuration options). FAST Cache addresses a hot spot anywhere in the array, either on RAID Group LUNs or storage pool LUNs. Flash

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 bursts of read/write loads at Flash drive speeds.

Another important benefit is improved total cost of ownership (TCO) of the system. FAST Cache copies the hot or active subsets of data to Flash drives in chunks. By offloading many if not most of the IOPS to the cache, the customer can fill the remainder of their storage needs with low cost, high capacity disk drives. This ratio of a small amount of Flash paired with a lot of disk offers the best performance (\$/IOPS) at the lowest cost (\$/GB) with optimal power efficiency (IOPS/KWH).

FAST Cache is enabled by default on all RAID group LUNs and storage pools once the FAST Cache enabler is installed. RAID group LUNs and storage pools created before the FAST Cache enabler is installed have FAST Cache disabled. To use FAST Cache for these items, you must manually enable FAST Cache by using either Unisphere or the CLI.

FlashEMC conducted application-specific tests with FAST Cache to characterize the 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). Because 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.

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 EMC Online Support. See the “References” section of this paper for the document titles.

FAST Cache components

FAST Cache requires the FAST Cache enabler to take advantage of the feature. To create FAST Cache, you need at least 2 Flash drives in the system, which will be configured in a RAID 1 RAID group. Once the enabler is installed, the system uses the following main components to process and execute FAST Cache:

Policy engine—Manages the flow of I/O through FAST Cache. When a chunk of data on a LUN is accessed frequently, it is copied temporarily to FAST Cache (Flash drives). It is copied back to HDDs when other data is more heavily used. The policy engine decides when these operations should take place. This policy engine also maintains statistical information about the data access patterns. The policies defined by this policy engine are system-defined and cannot be changed by the user.

Memory map—Tracks extent usage and ownership in 64 KB chunks of granularity. This map 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 the FAST Cache enabler is installed, the existing SP read and write cache may need to be disabled temporarily to allocate space for the FAST Cache memory map. Once the memory map has been created, the resized SP read and write caches are re-enabled automatically 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

FAST Cache Promotions

Data on HDDs that becomes busy is promoted to FAST Cache, which is an asynchronous process. 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. If you have FAST VP enabled with Flash or “extreme performance” drives in the tier, I/Os from that tier are not promoted to FAST Cache, because the data already resides on the Flash drives. For more information on how FAST VP and FAST Cache work together, see *Appendix B: FAST VP and FAST Cache*.

For example, assume an I/O request is received from an application immediately after FAST Cache is created on the storage system. In this scenario, the FAST Cache memory is empty, because nothing has been promoted 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. Because the memory map is empty at this phase, the data is accessed from the HDD LUN. This is called a *FAST Cache miss*. EMC has found that there is minimal performance overhead when checking the memory map for every access to a FAST Cache enabled LUN.
- 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.
 - 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 is created and empty. This also happens when the working data set of the application changes 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 similar to 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 depends on the number and type of HDDs in the back end. For example, a setup of 80 SAS drives has a shorter warm-up time than a setup with 20 SAS drives due to the larger working data set. Similarly, FAST Cache with SAS HDDs in the back end warms up faster than when using NL-SAS HDDs in the back end, because NL-SAS drives typically have a higher response time than SAS drives. Therefore, when you design application layouts, remember that there is a warm-up time before stable FAST Cache performance is reached.
- When the application accesses this data again, the policy engine sees that it is in the FAST Cache. This is called a *FAST Cache hit*. Because the data is now accessed from


the Flash drives, the application gets very low response times and high IOPS. If a substantial part of the working set is promoted to FAST Cache over time, applications can see higher average performance, even with lower performing HDDs in the back end.

Reads

Incoming I/O from the host application is checked against the FAST Cache memory map to determine whether the I/O is for a chunk that is already in FAST Cache:

If the chunk is not in FAST Cache, the I/O request follows the same path it would follow if the storage system does not have FAST Cache.

However, if the data chunk is in FAST Cache, the policy engine redirects the I/O request to 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 FAST Cache and placed in the DRAM cache as it would with reads from HDD.




1. FAST Cache driver checks map to determine where page is located.
2. Page request satisfied from DRAM if available.
3. Page request satisfied from disk drive if not in FAST Cache.
4. Policy Engine promotes the page to FAST Cache if it is being used frequently.
5. FAST Cache promotions copied to map and DRAM to be served to host.

Figure 1: FAST Cache read operation

Writes

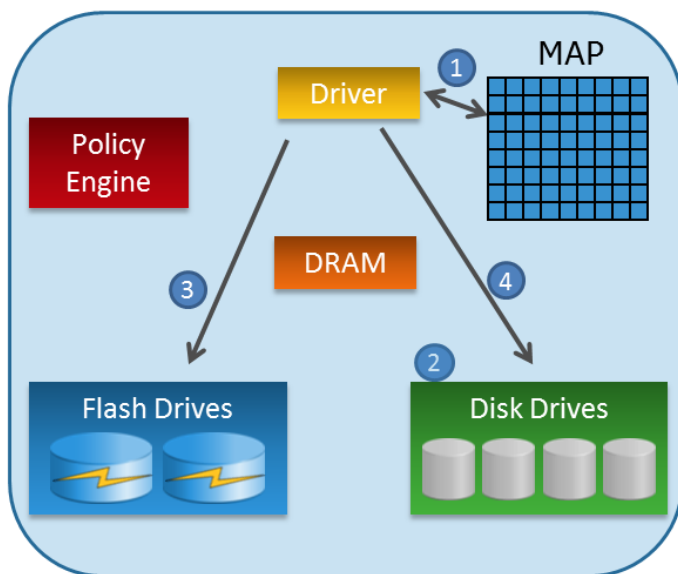
If the host I/O request is a write operation for a data chunk in FAST Cache, and the write cache is not disabled for the LUN, the DRAM cache is updated with the new “write,” and an acknowledgement is sent back to the host. The host data is not written directly to the FAST Cache. When data needs to be moved out of the DRAM cache, it is written to FAST Cache. Because the data is written to Flash drives instead of HDDs, this operation may help limit or decrease the number of dirty pages in the DRAM write cache.

Note: Even when FAST Cache is installed and enabled in the storage system, the I/O operations are serviced directly from DRAM cache whenever possible.



1. Driver checks the map to see if the I/O is located in FAST Cache.
2. Always write to the DRAM Cache If write Cache is not disabled for that LUN.
3. When data needs to be moved from DRAM it is written to FAST Cache if the data is already promoted to FAST Cache.
4. If the I/O is not in FAST Cache, the DRAM is flushed to the disk drives.

Figure 2: FAST Cache write operation



1. Driver checks the map to see if the I/O is located in FAST Cache.
2. Write Cache is disabled for the LUN.
3. Write operation goes directly to FAST Cache.
4. If I/O is not in FAST Cache, the write operations goes directly to HDD.

Figure 3: Write operation when write cache is disabled for the LUN

Flushing out of FAST Cache

In certain situations, data is copied from FAST Cache to the back-end HDDs. This is called a *write-back operation*. Write-back operations happen when a FAST Cache promotion is scheduled but there are no free or clean pages available in the FAST Cache. A dirty page is then copied from the FAST Cache and written to the HDD LUN to make room for the new data. The least recently used (LRU) algorithm determines which data blocks to flush to make room for the new promotions.

Small block sequential and high frequency access filter

In VNX OE Release 3.2, there have been improvements in dealing with small block sequential and short-lived bursts of activity with high spatial locality workloads. Previously, these workloads, with a low potential for re-hit, would trigger the promotion of pages into FAST Cache, resulting in very little benefit. With this enhancement, the FAST Cache will more intelligently identify longer-term access patterns of data that will be most beneficial for the system to use in FAST Cache, and avoid those with little to no benefit.

Failure handling

VNX 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 and 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.

If a hot spare is not available, a single pair of FAST Cache drives will remain in a degraded mode, and the cache-page cleaning algorithm increases the rate at which FAST Cache pages are copied from FAST Cache Flash drives to HDDs. Only read operations are allowed from this FAST Cache RAID group to reduce the potential exposure to data loss of a drive failure in a non-redundant RAID group. In this state, write operations that were destined for the degraded RAID group are now serviced by the HDDs. If there are other RAID groups in FAST Cache, those will continue to function as read / write. Once the degraded RAID group is repaired, the FAST Cache will then seamlessly start to re-warm data into it.

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 EMC Online Support. The following sections discuss the parts of Unisphere and Secure CLI that pertain to FAST Cache. For more information on FAST Cache configuration options, see *Appendix A: FAST Cache configuration options*.

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 4**).

To enable FAST Cache, click the **FAST Cache** tab in the 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 4: Storage System Properties dialog box

When 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 displays **RAID 1** when FAST Cache has been created. Clicking **Create** opens the Create FAST Cache dialog box (Figure 5).

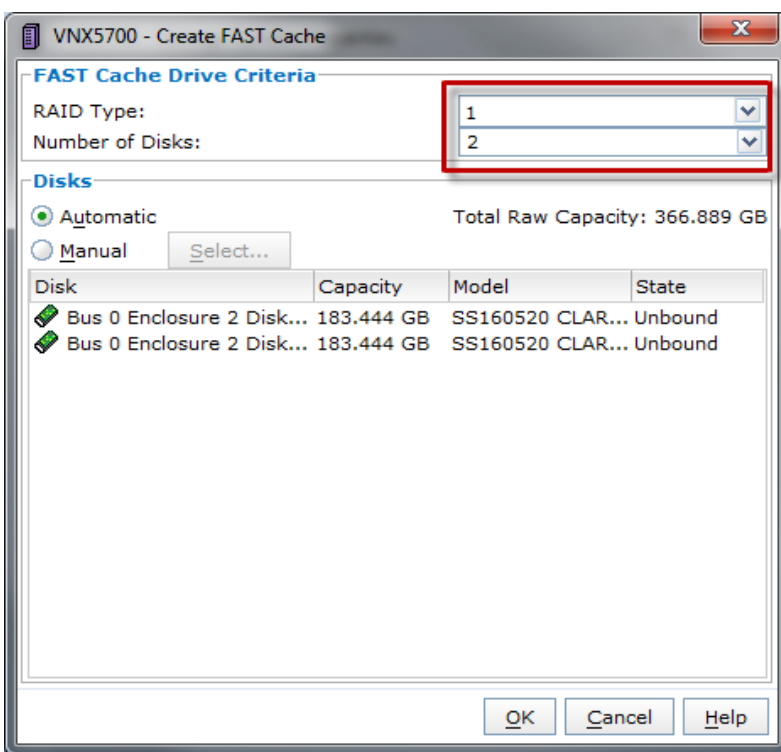


Figure 5: Create FAST Cache dialog box

If a sufficient number of Flash drives are not available to enable FAST Cache, Unisphere displays an error message, and FAST Cache cannot be created. 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. To change the size of FAST Cache after it is configured, you must destroy and recreate the FAST Cache. This requires FAST Cache to destage all of its data. When the FAST Cache is created again, it must repopulate its data (warm-up period).

Figure 3 shows how you can enable the FAST Cache for LUNs under the **Advanced** tab in the Create LUN dialog box.

If the LUN has already been created in a RAID group, click the **FAST Cache** tab in the LUN Properties dialog box to configure FAST Cache (shown in Figure 7).




Figure 6: Advanced tab in the Create LUN dialog box

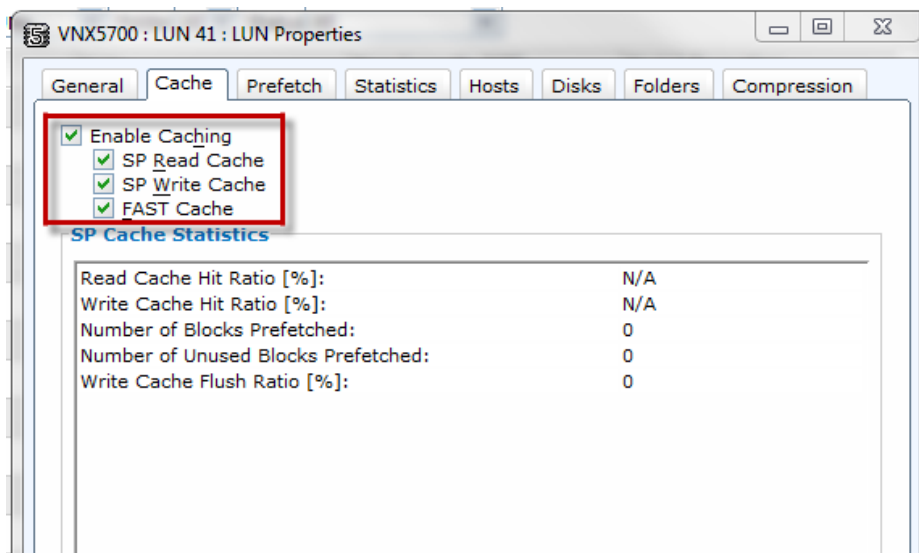


Figure 7: Cache tab in the LUN Properties dialog box

For FAST Cache with storage pools, you can enable FAST Cache per pool. All the LUNs created in the storage pool will have FAST Cache enabled or disabled collectively. You can configure FAST Cache on Pools by using the **Advanced** tab in the Create Storage Pool dialog box shown in Figure 8.




Figure 8: Advanced tab in the Create Storage Pool dialog box

If the storage pool has already been created, use the **Advanced** tab in the Storage Pool Properties dialog box to enable FAST Cache (see Figure 9).

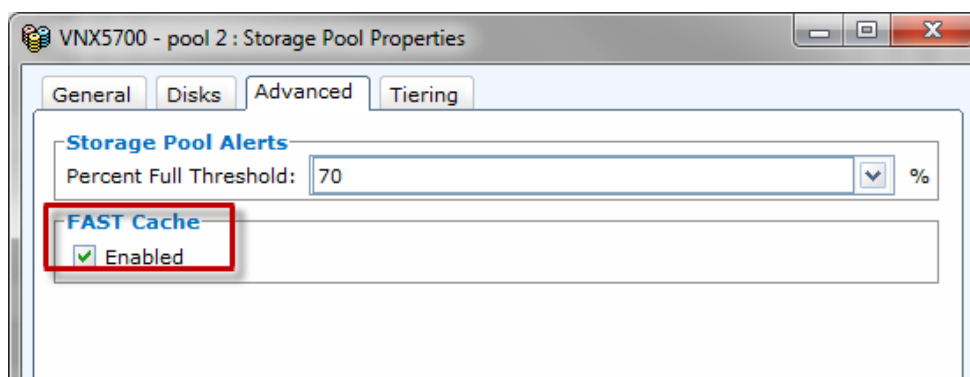


Figure 9: Advanced tab in the Storage Pool Properties dialog box

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, shown in **Figure 10**, where you can select **FAST Cache**. The FAST Cache property is displayed for every entry in the table.




Figure 10: Choosing to display FAST Cache information in Unisphere

Secure CLI

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

Task	CLI command
Create FAST Cache	cache -fast -create
Destroy FAST Cache	cache -fast -destroy
Get FAST Cache info	cache -fast -info
Configure FAST Cache when creating RAID Group LUN	bind ... -fastcache 0 1
Enable or disable FAST Cache on a RAID Group LUN	chglun -l <LUN #> -fastcache 0 1
Get FAST Cache config info from a RAID group LUN	getlun <LUN #> -fastcache
Configure FAST Cache when creating a storage pool	storagepool -create ... -fastcache on off
Configure FAST Cache on existing storage pool	storagepool -modify -id <#> -fastcache on off
Get FAST Cache state on storage pool	storagepool -list -id <#> -fastcache

... is indicative of more cli options required

Unisphere Analyzer

Unisphere Analyzer gathers FAST Cache statistics to monitor FAST Cache performance. To view these statistics, enable Analyzer's **Advanced** mode by using the following steps:

1. In Unisphere, click the **System** tab.
2. Click **Monitoring and Alerts**.
3. Click Statistics **for Block**.
4. Click **Customize Charts**.
5. Click the **General** tab.
6. Select the **Advanced** checkbox.
7. 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

EMC Online Support contains videos that will help you view these statistics. Log into EMC Online Support and search for “Analyzer Series” for the FAST Cache video.

Best practices

- Preferred application workloads for FAST Cache:
 - Small-block random I/O applications with high locality
 - Data rehits—High frequency of access to the same data
 - Systems where current performance is limited by HDD capability, not SP capability
- 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. Then 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. FAST VP only benefits the storage pool where the Flash drives reside. For more information on FAST Cache and FAST VP, see *Appendix A: FAST Cache configuration options*.
- The VNX5100 storage systems allow you to use either the FAST Cache or Thin Provisioning feature, but not both. If the Thin Provisioning™ feature is installed on the storage system, you cannot use FAST Cache. If you are using FAST Cache, you cannot use Thin Provisioning. On all other VNX models, you can use both features simultaneously.
- Unisphere allows you to choose the Flash drives to use to create FAST Cache. You can also choose these drives manually to ensure that you distribute the Flash drives across back-end buses. Avoid placing drives in the DPE or DAE-OS enclosure (0_0) that will be mirrored with drives in another enclosure. For example, DON'T mirror a disk in 0_0 with a disk in 1_0.
- FAST Cache can improve overall system performance if the current bottleneck is drive-related, but boosting the IOPS will result in greater CPU utilization on the SPs. Systems should be sized so that the maximum sustained utilization is 70 percent.

On an existing system, check the SP CPU utilization. If the utilization is over 80%, contact an EMC storage specialist to review the system health and determine next steps, before enabling FAST Cache.

- FAST Cache is most beneficial when the primary workload is random small block. To maximize the most out of FAST Cache avoid enabling it if the primary workload is sequential and/ or large block I/O.
- Random small block workloads provide the best use of FAST Cache, avoid enabling it for small block sequential workloads such as Database logs, Circular logs, and VNX File SavVol LUNs.

Note: For storage pools, FAST Cache is a pool wide feature so you will enable/disable at the pool level (for all LUNs in the pool).

More detailed best practice guidelines can be found in the *VNX Unified Best Practices for Performance* white paper available on EMC Online Support. Refer to application-specific white papers for guidelines on using FAST Cache with those applications.

Limitations

- Some optional applications, such as MirrorView™ and SnapView™, require private LUNs. These LUNs are already optimized for priority in the storage system's write cache, and therefore do not need to utilize FAST Cache. EMC recommends that you disable FAST Cache on MirrorView's write intent log and SnapView's clone private LUNs to prevent 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 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 the reserved LUNs are created in a pool and there are other LUNs in the pool that need FAST Cache, they can be left with FAST Cache enabled.
- 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. This does not happen if you have a VNX7500 with the extra memory upgrade.
- Flash drives installed in vault drive locations cannot be used to create FAST Cache. The VNX OE restricts this usage.

Conclusion

FAST Cache allows the storage system to provide Flash-drive class performance to data with a high locality of reference. This working data set increases IOPs without placing all of the data onto Flash drives. FAST Cache absorbs I/O bursts from applications, thereby reducing

the load on HDDs, which helps to improve the TCO of the storage solution. You can manage FAST Cache through Unisphere in an easy, intuitive manner.

FAST Cache might not be a natural fit for all types of workloads, ie sequential and large block workloads. Analyze the application I/O profile 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 EMC Support Zone:

- *EMC Unified Storage System Fundamentals for Performance and Availability*
- *EMC VNX FAST VP*
- *EMC Unisphere: Unified Storage Management Solution*
- *EMC VNX Virtual Provisioning*
- *Leveraging EMC FAST Cache with Oracle OLTP Database Applications*
- *EMC® Infrastructure for VMware® View™ 5.0*
- *Applied Best Practices Guide: EMC VNX Unified Best Practices for Performance*
- *Customer Technical Presentation: EMC Multisite Disaster Recovery for Microsoft SQL Server 2012 — EMC VNX5700, EMC FAST Cache, SQL Server AlwaysOn Availability Groups*

Appendix A: FAST Cache configuration options

Table 1: 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	

Table 2: 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); therefore, 200 GB Flash drives cannot be used to configure FAST Cache on this platform.

Appendix B: FAST VP and FAST Cache

FAST VP is a feature that was introduced in VNX OE release 30. 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 *EMCVNX FAST VP* white paper available on EMC Online Support.

Table 3: 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 is 64 KB.	Granularity 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.
Use when workload changes are unpredictable and very dynamic, and require a quick response time.	Use when workload pattern changes are predictable and relatively low.
Constantly promotes frequently accessed HDD data to FAST Cache. There are no 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.

You can use FAST Cache and FAST VP to yield high performance and TCO from the storage system. For example, you can use Flash drives to create FAST Cache, and use FAST VP 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 fewer Flash drives, while FAST VP optimizes disk utilization and efficiency with SAS and NL-SAS drives.

As a general rule, use FAST Cache in cases where storage system performance must be improved immediately for burst-prone data with a high locality of reference. 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, when 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.
- If a bursty workload starts accessing a particular chunk of a FAST Cache-enabled LUN, FAST VP does not immediately move that chunk to a different storage tier. Instead, FAST Cache promotes the chunk into the cache. After the chunk is promoted, a majority of I/O operations will be serviced from the FAST Cache. This can result in less activity on the back-end LUNs, and FAST VP might not need to move the chunk to a higher-storage tier. In this case, a FAST VP-initiated data move is avoided when there is a temporary burst in an application's workload.
- In contrast to the previous scenario, if the application workload has increased on a sustained basis, FAST Cache will need to write data back into the HDD LUNs to make space for new promotions. This will register as back-end activity and FAST VP will eventually schedule a move of the data chunks to higher storage tier—which may be Flash drives. When this move is completed, FAST Cache does not promote any data that is already in the Flash drive storage tier.
- You will see higher performance benefits and faster reaction time for changing I/O 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. Performance gains are due to the flushes from DRAM of an I/O that is in FAST Cache to the Flash drives instead of back to HDDs, which happens if FAST Cache is not enabled.

Appendix C: 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 HDDs.

Table 4: Comparison of DRAM memory and FAST Cache

Characteristic	DRAM Cache	FAST Cache
Position	Closest to the CPU, and has the lowest latency.	A step further from the CPU and is slower than DRAM cache.
Response time	Response time is from nanoseconds to microseconds.	Response time is from microseconds to milliseconds.
Upgradeability	Upgrade kit available for VNX 7500 only.	Upgradeable in all supported models, and the options depend on storage system model and type of Flash drive.
Operation	Has separate user-configurable areas for read and write operations.	Has a single area that serves read and write operations.
Capacity	Is limited in size, compared to FAST Cache.	Can scale to much larger capacity.
Granularity	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.	Operates in extents of 64 KB granularity.
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	Contents are volatile and cannot withstand a power loss.	Contents are non-volatile and can withstand a power loss.