

USING EMC FAST SUITE WITH SYBASE ASE ON EMC VNX STORAGE SYSTEMS

Applied Technology

Abstract

This white paper introduces EMC's latest innovative technology, FAST Suite, and emphasizes how users can leverage that technology with Sybase ASE 15.5 database applications. This white paper covers several use cases that were tested in EMC labs, along with their test results. This paper also discusses best practices for implementing FAST Suite technologies (FAST VP and FAST Cache) with Sybase OLTP database applications.

June 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.

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

Part Number h8256.1

Table of Contents

Executive summary	4
The need for FAST Suite	4
Audience	5
Introduction	5
Technology overview	6
VNX platform	6
EMC FAST Suite	7
FAST VP technology	7
FAST Cache technology.....	7
Sybase Adaptive Server Enterprise	8
Application I/O characteristics	9
Benchmark setup	9
Workload description	10
Use case descriptions	11
Baseline scenario.....	12
Use Case 1	13
Use Case 2	15
Use Case 3	17
Results summary.....	19
Storage considerations for using FAST Suite	19
Impact of data placement and locality of data on FAST Suite performance	20
Conclusion	21
References	21

Executive summary

EMC® unified storage platforms continue to provide improved storage use, enhanced capabilities, and greater interoperability and security, as well as multiple ease-of-use improvements. The business challenge that many Sybase Adaptive Server Enterprise (ASE) customers face today is reducing total cost of ownership (TCO) while improving service-level delivery. Frequently, responsiveness to sensitive Online Transaction Processing (OLTP) applications has deteriorated over time due to increased data volumes, unbalanced data stores, and changed business requirements. Currently, the Serial Attached SCSI (SAS) drive technology does not keep up with the demand for more storage capacity and performance. By using EMC's newest storage platform, EMC VNX™ storage with EMC's FAST Suite, which includes Fully Automated Tiering for Virtual Pools (FAST VP) and FAST Cache, Sybase OLTP applications can often gain a significant responsiveness boost without the need to redesign the applications, tune the ASE database, adjust the data layouts, and/or reload significant amounts of data. As a result, business applications can take full advantage of the investments made in Flash drives, delivering optimal application service time and improving the TCO of the solution. This white paper addresses the considerations and guidelines for deploying FAST Suite with Sybase ASE OLTP applications.

The need for FAST Suite

Sybase ASE OLTP database applications tend to be mission-critical and have stringent transaction responsiveness requirements. Sybase ASE server performance depends heavily on the I/O subsystem, particularly I/O service latencies. The ultimate limit of throughput and input/output operations per second (IOPS) for data reads and writes by the ASE server is typically limited by disk I/Os. Physical I/O performance can make or break a mission-critical Sybase ASE database application due to its huge impact on the overall user response time.

FAST VP provides performance improvements by using automated sub-LUN-level tiering and allows a single LUN to leverage the advantages of Flash, SAS, and Near Line Serial Attached SCSI (NL-SAS) drives by using storage pools. FAST VP automatically balances data distribution over the tiers of storage to optimize both space utilization and performance delivered from the storage investment.

FAST Cache technology uses Flash drives to add an extra layer of cache between Dynamic Random Access Memory (DRAM) cache and rotating spindles, which increases the I/O service responsiveness. The Sybase OLTP application's hot and cold data is automatically identified, and hot data is either placed in the FAST VP's Flash tier or cached in the FAST Cache layer to speed up data access.

FAST VP and FAST Cache technology provide very low latencies to frequently accessed data, thus improving overall application response times. By using Flash drives for the most frequently accessed data, the investment made in Flash tier/FAST Cache is leveraged to the most optimal extent in order to deliver a magnitude of application-level service improvement. As the focus of business changes, the data that becomes

the most important, and most frequently used, will either be automatically migrated based on relocation schedule to the Flash tier or cached in Flash drives. Therefore, business applications can always take full advantage of the performance benefits of Flash drives. The performance improvements gained by using the FAST Suite in Sybase OLTP applications provide quantifiable operational cost savings over time. The following figure shows the FAST Suite consisting of FAST VP with three tiers of storage layers and FAST Cache.

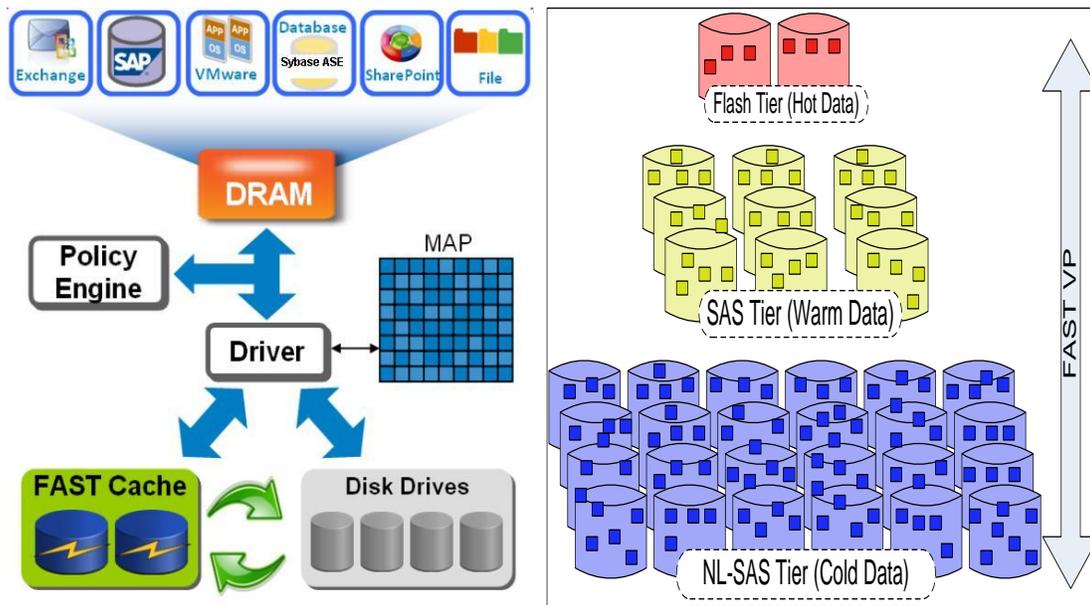


Figure 1. FAST Cache and FAST VP

Audience

This white paper is intended for Sybase ASE database administrators, storage architects, customers, and EMC field personnel who want to improve the performance of business applications by implementing FAST VP and FAST Cache technology in their Sybase OLTP environments using EMC VNX storage.

Introduction

The new EMC VNX family of unified storage platforms continues the EMC tradition of providing some of the highest data reliability and availability in the industry, but they also include in their design a boost in performance and bandwidth to address the sustained data access bandwidth rates. The new system design has also placed heavy emphasis on storage efficiencies and density, as well as crucial green storage factors, such as a smaller data center footprint, lower power consumption, and improvements in power reporting.

This white paper details the positive effects of deploying the EMC FAST Suite. The goal is to verify that the use of FAST VP and FAST Cache together can boost performance and reduce TCO for certain kinds of Sybase ASE applications.

Technology overview

VNX platform



Figure 2. EMC VNX models

A key distinction of this new generation of EMC unified storage products is support for both block- and file-based external storage access over a variety of access protocols, including Fibre Channel (FC), iSCSI, FCoE, NFS, and CIFS network shared file access. Furthermore, data stored in one of these systems, whether accessed as block- or file-based storage objects, is managed uniformly under a single EMC Unisphere™ web-based interface. For more details about the new features that Unisphere supports, refer to the [References](#) section.

The rapid increase in Sybase deployments has added a tremendous amount of new data to storage systems. In addition to being able to readily accommodate storing new data at an unprecedented rate, our new generation of systems must also push the bandwidth envelopes to a totally new level to keep up with that data growth.

Another consequence of the rapid data growth is the need to store the data more compactly. Data centers now continually contend with physical space and power and cooling challenges in order to store an ever-increasing volume of data that the business demands. The VNX storage family now supports the 2.5" SAS drives in a 2U disk-array enclosure (DAE) that can hold up to 25 drives, one of the densest offerings in the industry. For example, compared to the older technology of storing 15 x 600 GB worth of data using 3.5" FC drives in a 3U DAE, the new DAE uses 25 x 600 GB drives in a 2U footprint, which results in a 1U footprint reduction and a capacity 2.5 times greater. The power efficiency of the new DAEs also makes it more cost-effective to store the increased data in a more compact footprint without the need to increase power consumption and cooling.

Most of the data points discussed in this paper apply to all VNX storage platforms except the VNX5100™ model.

EMC FAST Suite

The FAST Suite is a software bundle offered on the VNX series systems and is composed of two storage efficiency technologies and two management products—FAST VP and FAST Cache, and Unisphere Analyzer and Unisphere Quality of Service Manager, respectively. For the purposes of this paper we will focus solely on the efficiency technologies, FAST VP and FAST Cache.

Both technologies are aimed at reducing the number of Flash drives required to meet a given performance requirement, significantly reducing the TCO of a storage system while achieving the desired level of performance.

FAST VP technology

EMC Fully Automated Storage Tiering for Virtual Pools (FAST VP) technology is the storage-array-based software that transparently provides a policy-based, auto-tiering solution for enterprise applications. FAST VP operates at a granularity of 1 GB, referred to as a “slice.” The goal of FAST VP technology is to leverage storage tiers to lower customers’ TCO and to increase performance by keeping hotter slices of data on performance tiers and colder slices of data on high-capacity tiers. High locality of data is important to realize the benefits of FAST VP. FAST VP can relocate sub-LUN-level pieces of data from one storage tier to another by moving more active data to faster drives like Flash drives for better performance and less active data to slower drives like NL-SAS drives for reduced storage costs. This occurs automatically and transparently to the host environment. FAST VP operates on the storage data by relocating the most active data to the highest available tier (typically Flash) in the scheduled relocation window. To ensure that sufficient space is available in the higher tiers, FAST VP relocates less active data to lower tiers (SAS or NL-SAS). When FAST VP relocates data, it moves the entire slice to the new storage tier. To successfully identify and move the correct slices, FAST VP performs statistics collection, analysis, and relocation. You can initiate the relocation of slices manually or automatically with a user-configurable, automated scheduler that is accessible from the Unisphere management tool. The multi-tiered storage pool allows FAST VP to fully use all three storage tiers: Flash, SAS, and NL-SAS. The creation of a storage pool allows for the aggregation of multiple RAID groups, using different storage tiers, into one object. You can provision the LUNs created from the storage pool as either thick or thin. These “pool LUNs” are no longer bound to a single storage tier. Instead, you can spread them across different storage tiers within the same storage pool. If you create a storage pool with one tier (Flash or SAS or NL-SAS), then the FAST VP feature has no impact on the performance of the system. To operate FAST VP, you need a minimum of two tiers in the pool.

FAST Cache technology

In traditional storage arrays, the DRAM caches are relatively small, making it a challenge to maintain the hot data for long periods of time. Very few storage arrays

enable users to nondisruptively expand DRAM cache, even if they support DRAM cache expansion. FAST Cache technology builds a faster medium (sort of cache) for data on Flash drives, so that the hotter data stays longer on a faster medium. FAST Cache tracks the number of hits at a 64 KB chunk size and copies the chunks to Flash drives once the chunks get three or more hits. After a data chunk is copied to FAST Cache, the subsequent accesses to that chunk of data will be served at Flash latencies. Eventually when the hotness of data chunk decreases, they are evicted from FAST Cache and are replaced by new, hot data. FAST Cache uses a simple least reused (LRU) mechanism to evict the data chunks when it needs room for additional hot data that has met the access threshold.

A complete introduction to FAST Cache technology is beyond the scope of this paper. This paragraph will provide a brief overview of FAST Cache. FAST Cache is built on the premise that overall application latencies can improve when most frequently accessed data is maintained on a smaller sized, but faster, storage medium, such as Flash drives. FAST Cache technology identifies the frequently accessed data used by currently running business applications and moves it to Flash drives automatically and nondisruptively. The data move is completely application-transparent and it is as simple as turning a switch on or off. The LUN/pool-level granularity enables FAST Cache to be selectively enabled on a few LUNs or pools within a storage array depending on the application performance requirements and SLAs.

FAST Suite technology enhances the user choices for deploying Flash drives in their environment with the following options:

- Flash drives as a separate tier of storage by themselves
- Flash drives as a tier within a pool of storage by using FAST VP
- Flash drives as a logical extension to DRAM cache by using FAST Cache

The FAST Suite and FAST VP algorithms also work hand-in-hand to avoid having data on Flash medium twice. FAST Cache does not cache any data that already resides on a Flash tier of storage pools or RAID groups.

Sybase Adaptive Server Enterprise

Sybase ASE is a database management system (DBMS) designed to support transaction-intensive, mission-critical OLTP, decision support, and mixed-load applications. ASE also manages multiple databases and users, keeps track of the data's location on disks, maintains the mapping of the logical data description to the physical data storage, and maintains data and procedure caches in the memory. ASE 15 meets the demands of large databases and high transaction volumes while providing a DBMS. The most recent release, ASE 15.5, adds several important features for enterprise-class customers. Its key features include on-disk encryption, smart partitions, In-Memory databases and incremental data transfer utility, multiple Tempdb groups, efficient backup compression, deferred name resolution for user-defined stored procedures, new query processing technology that demonstrates a significant increase in performance, and enhanced support for unstructured data management.

Application I/O characteristics

The relative improvement from any faster technology depends on if the application is bottlenecked by the storage subsystem. For Sybase DBAs, it is frequently easier to work with storage administrators to avoid I/O bottlenecks than to redesign the application to minimize the amount of inefficient or extra data access by the database system. Nevertheless, some amount of analysis at the database level is required to identify I/O bottlenecks, and to determine that the proper application of FAST Suite leveraging Flash drives is worthwhile.

Benchmark setup

To characterize the EMC FAST Suite with Sybase ASE OLTP workloads, the following benchmark setup was used: The benchmark testing was done on a single Sybase ASE 15.5 database deployed on a VNX5500™ storage array and Dell R610 server. The configuration in Figure 3 shows two back-end buses on VNX5500 used to connect all the disk trays in a daisy chain fashion. You should configure your back-end buses depending on your I/O requirement and by following EMC best practices for bus and disk layout. Figure 3 also shows the hardware system setup with front-end connections on the left side and the VNX5500 back-end connections on the right side.

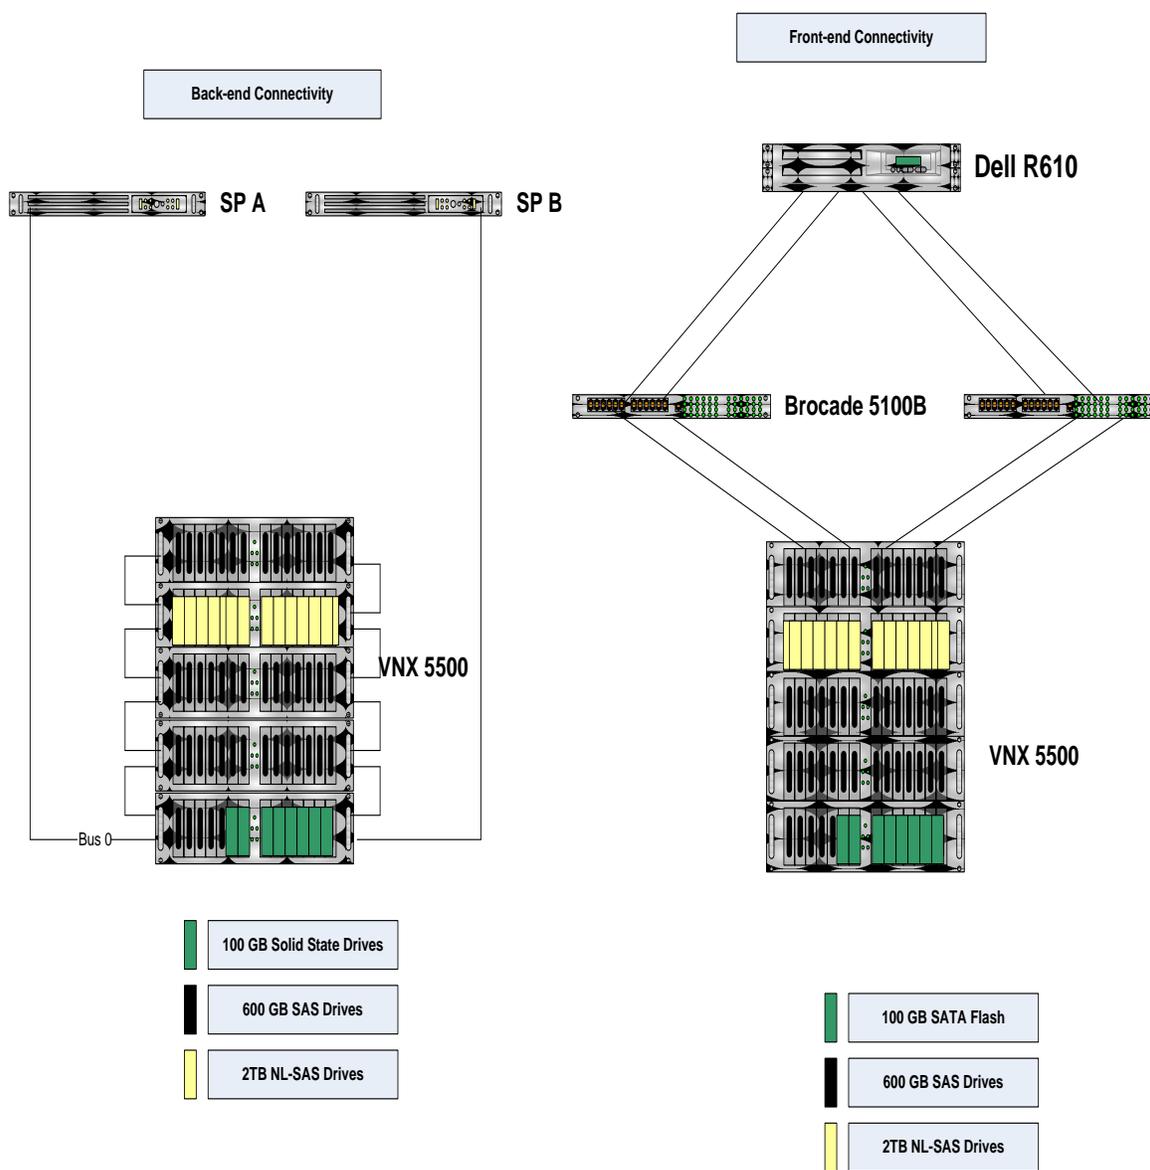


Figure 3. Hardware setup

Workload description

The database workload simulates an OLTP application (Order Entry) environment fashioned after the business transaction profile described in the popular industry TPC-C benchmark. The test workload is TPC-C-like (but not strictly TPC-C specification conforming), including a mixture of read-only and update-intensive transactions. The Sybase ASE database devices were distributed against LUNs in such a way that both storage processors handle an equal number of I/O requests.

To ensure that FAST Suite can handle workloads involving a high mixture of writes, EMC created a suitable test scenario. Generally, most of the OLTP databases tend to

have a read/write mix of 70/30. The Sybase OLTP database benchmark used in the test has the characteristics shown in the following table.

Table 1. Workload characteristics

Property	Value / Description
Database size	1.4 TB
Hot Working Dataset	200GB
Database version	Sybase ASE 15.5
Read / write ratio	70 / 30
Database metric	TPM
Number of concurrent users	200
Storage pools	SybaseDATA , SybaseLOG
Total LUNs	12 LUNs – 5 x 100 GB LUNs for database devices and segments; 1 x 300 GB LUN for a Log device; 2 x 200 GB LUNs for a Stock table; 2 x 200 GB LUNs for a Customer table; and 2 x 200 GB LUNs for an OrderLine table

EMC kept consistent workload characteristics for all the use cases described next so that users can analyze the application metrics.

Use case descriptions

EMC conducted the following tests to understand the impact of the FAST Suite on Sybase ASE OLTP databases and to compare the impact of the FAST Suite on the workload using the application-level and database-level metrics. The goal of this study was to show how the FAST Suite can significantly reduce the number of IOPS received from rotating spindles and improve performance by relocating the hot data to faster Flash drives. The following use cases are crafted to demonstrate this very basic point.

- **Baseline:** Establishes the baseline metric on a storage pool of 45 x 600 GB SAS drives.
- **Use Case 1:** Enables FAST Cache with 100 GB Flash drives to the baseline environment. Number of Flash drives used: {2, 4}.
- **Use Case 2:** Repeats Use Case 1 where SybaseDATA pool was expanded with 5 x 100 GB Flash drives and 15 NL-SAS drives.
- **Use Case 3:** Enable FAST Cache with 2 x 100 GB Flash drives to the Use Case 2 environment.

Baseline scenario

The baseline was established on all SAS drives in this configuration. A total of 45 x 600 GB FC disks was used to create the first storage pool (40 SAS) for Data and the second storage pool (5 SAS) for Log in this setup. Then, 12 LUNs (see Table 1) were carved out to store a Sybase ASE database master device, log device, and segments. Since the Stock table, Customer table, and OrderLine table are bigger tables and will be frequently accessed during TPC-C like testing, the Stock table was stored on a striped logical volume created on 2 x 200 GB LUNs and similarly the Customer and OrderLine tables were stored on a striped logical volume created on 2 x 200 GB LUNs, respectively. This configuration helps in distributing the I/O load equally among storage processors. The next figure shows the exact disk and database layout for this use case.



Figure 4. Baseline disk and database layout

Use Case 1

The database layout for this use case is identical to the baseline scenario except that FAST Cache (read/write) was created initially on 2 x 100 GB Flash drives and was enabled on the storage pool containing Sybase devices and segments. A second iteration was done using 4 x 100 GB Flash drives as a comparison (refer to Figure 5). *FAST Cache was not enabled on the storage pool that has a LUN for ASE's log device because you will get better results by enabling FAST Cache on a storage pool with LUNs for database devices and segments.* After FAST Cache was enabled for Sybase ASE, the frequently accessed data in this workload was cached by FAST Cache. This warming up of FAST Cache resulted in improved latencies to the most accessed data, which enhanced overall application performance.

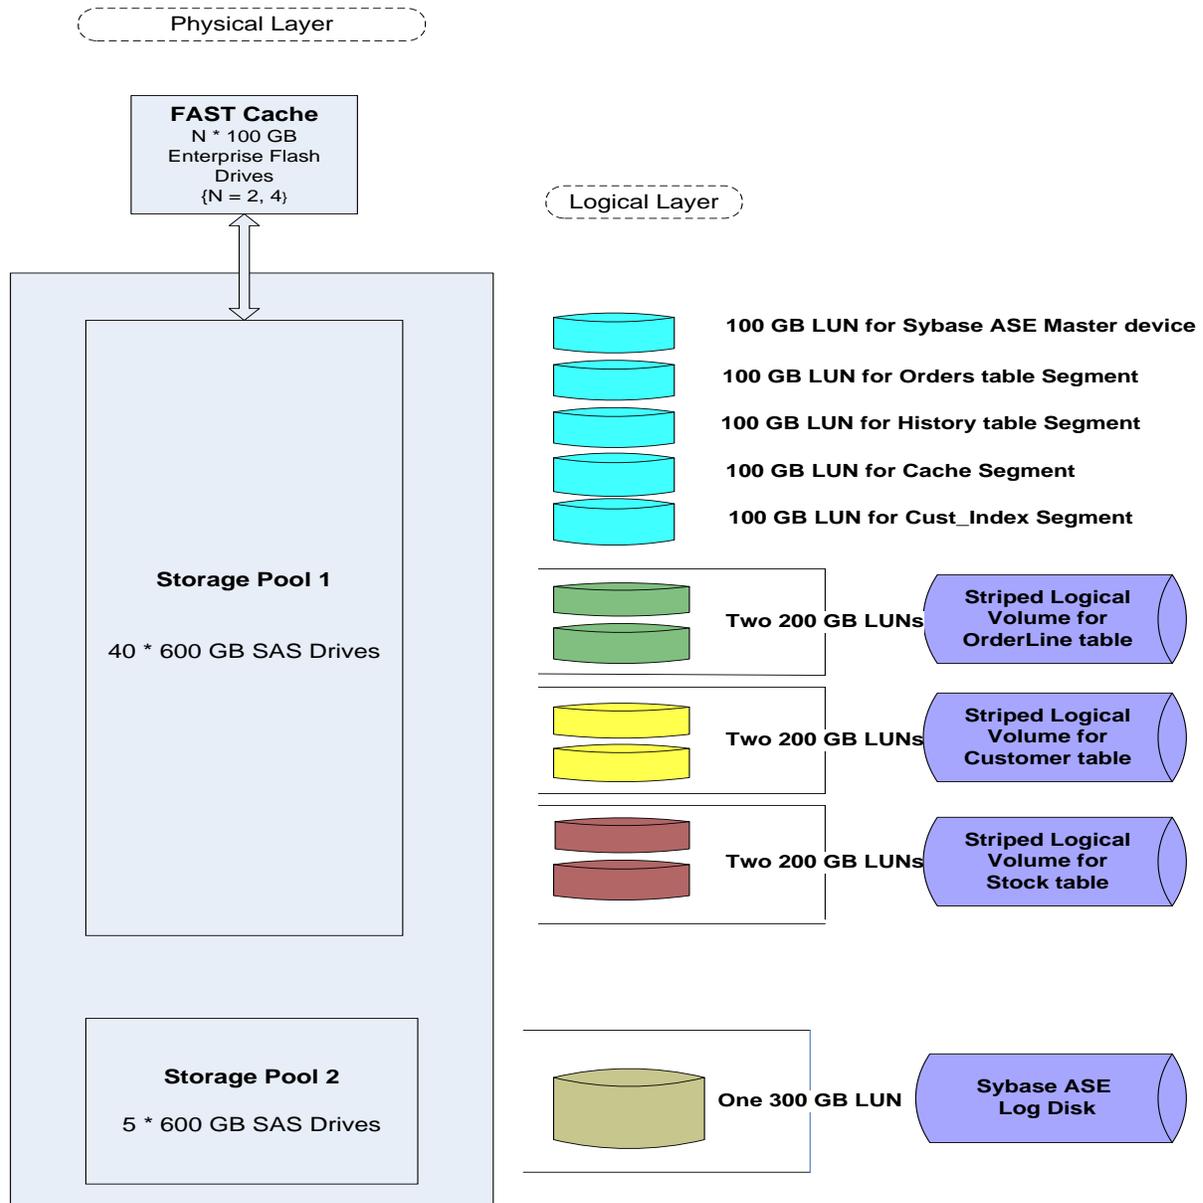


Figure 5. Use Case 1 (FAST Cache) disk and database layout

After FAST Cache was enabled for Sybase ASE, FAST Cache cached the data that the OLTP workload frequently accessed. This warming up of FAST Cache resulted in improved latencies to the most accessed data, which enhanced overall application performance. Once the active working set of the database is cached by FAST Cache, adding more drives will simply spread the same IOPS over newly added drives and effectively drive fewer IOPS out of each Flash drive. The results shown in Figure 6 reveal that Sybase OLTP application performance can be improved significantly by enabling FAST Cache, without performing any fine-tuning at the OLTP application level.

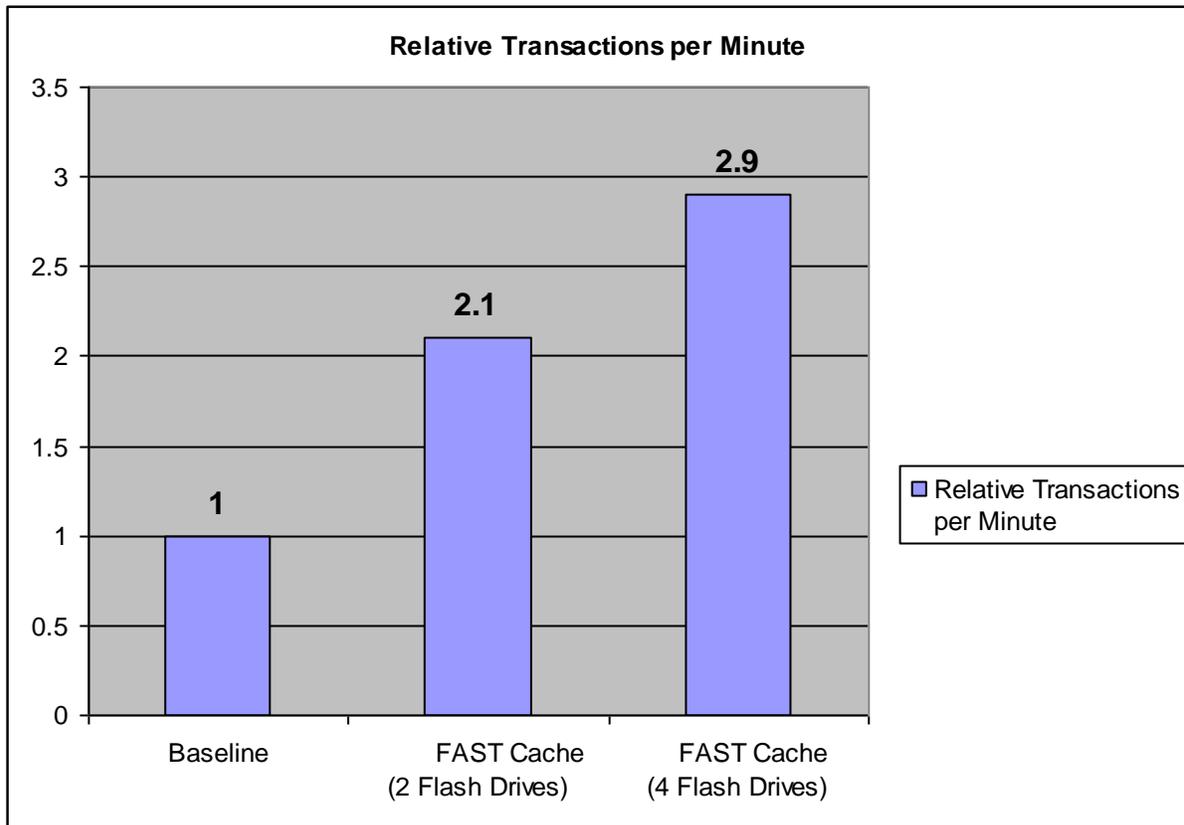


Figure 6. FAST Cache advantage

Use Case 2

Figure 7 shows the Sybase TPC-C-like workload running on the storage pool configuration established on 45 SAS drives, 15 NL-SAS drives, and 5 Flash drives. One storage pool (40 SAS + 15 NL-SAS + 5 EFDs¹) for data and another storage pool (5 SAS) for log were carved out from this setup.

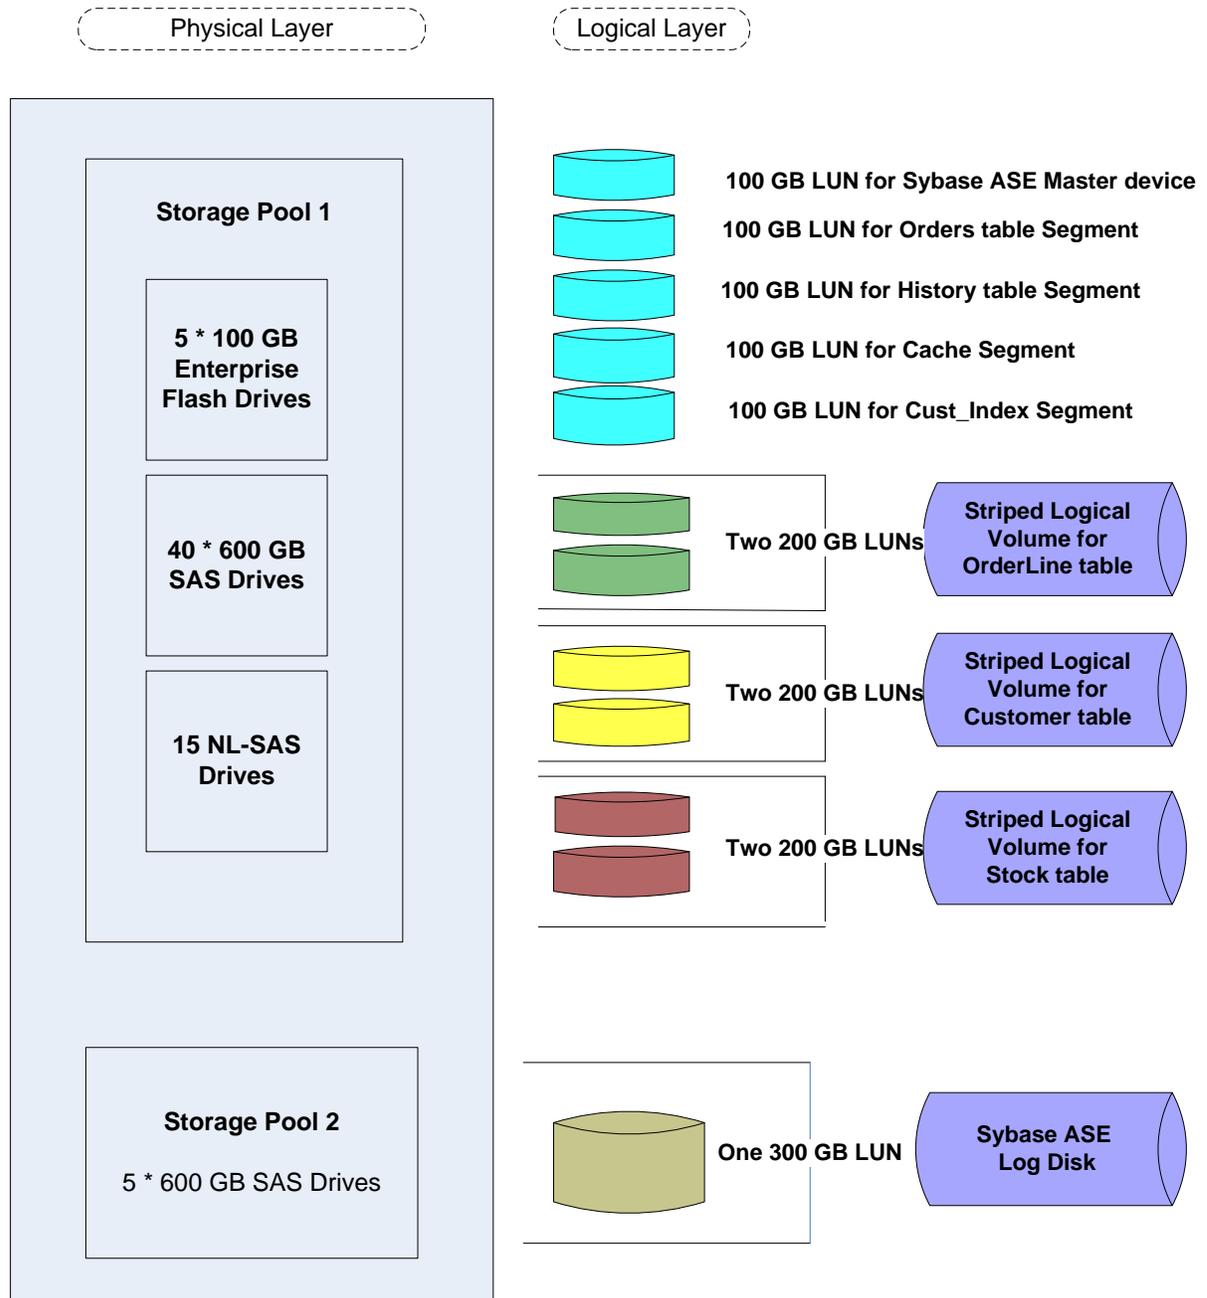


Figure 7. Use Case 2 (FAST VP) disk and database layout

¹ Enterprise Flash Drives

As discussed earlier the storage pools can be grown nondisruptively and dynamically as requirements for capacity and performance increase. When new drives are added to a pool, the FAST VP policy engine makes data migration recommendations based on historic data temperature information it already has on that storage pool. Users can choose to wait until the next scheduled data migration window for data movement or implement the recommendations immediately by starting the data relocation operation manually on that pool. In this use case, a manual migration is started to implement the recommendations immediately. After FAST VP relocated the hot/cold data, the OLTP workload was run against the tiered dataset. This solution achieves approximately 3.7 times the number of transactions compared to the baseline configuration after the relocation hot/cold data among storage tiers.

Use Case 3

The database layout for this use case is identical to Use Case 2 except that FAST Cache (read/write) was created on 2 x 100 GB Flash drives and was enabled on the storage pool containing Sybase devices and segments.

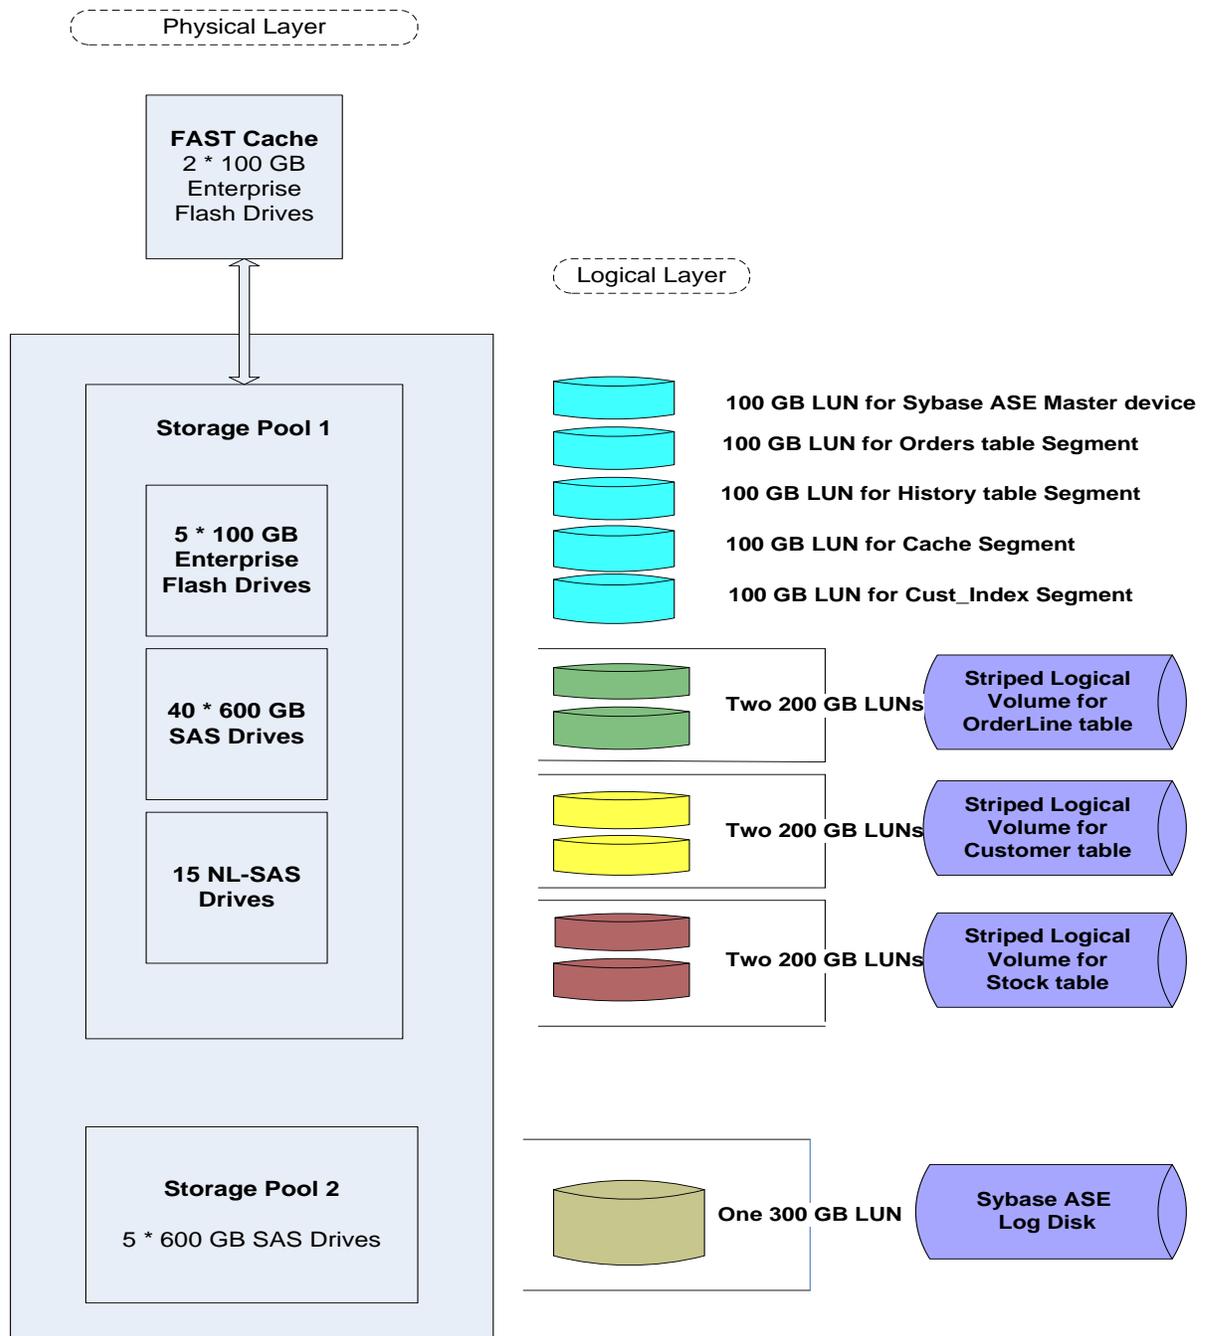


Figure 8. Use Case 3 (FAST Suite) disk and database layout

The five Flash drives in the storage pool hold the hot data for the Sybase OLTP application and two Flash drives act as FAST Cache (read/write) to cache the intraday changes to the hot working dataset. In this configuration, five Flash drives were in the pool and the majority of the hot data should have migrated to these Flash drives during the manual migration step discussed earlier. The small FAST Cache handles sudden intraday spikes over small data ranges. This solution achieves approximately 4.9 times the number of transactions compared to the baseline configuration after relocation of the hot/cold data among the storage tiers and after the warmup of FAST Cache.

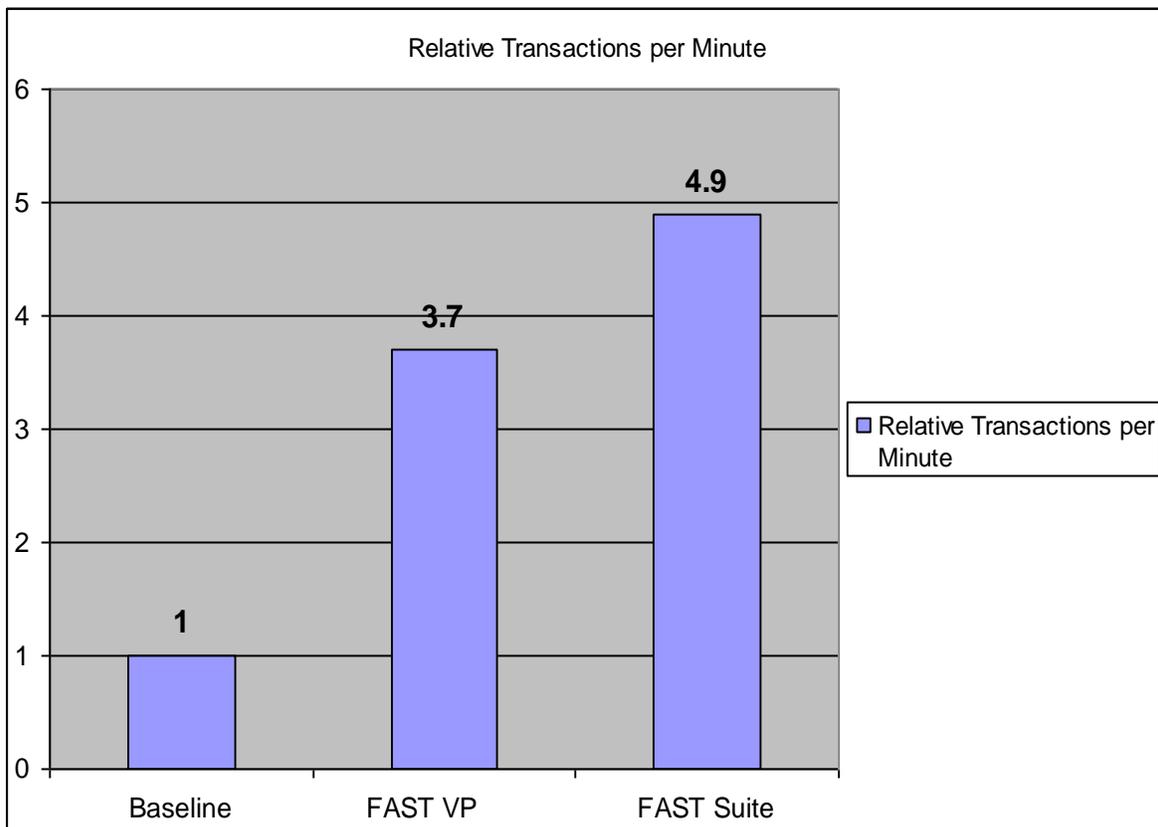


Figure 9. FAST Suite results

Results summary

Figure 10 represents a comparison of TPM rates for the baseline, FAST Cache, FAST VP, and FAST Suite use cases.

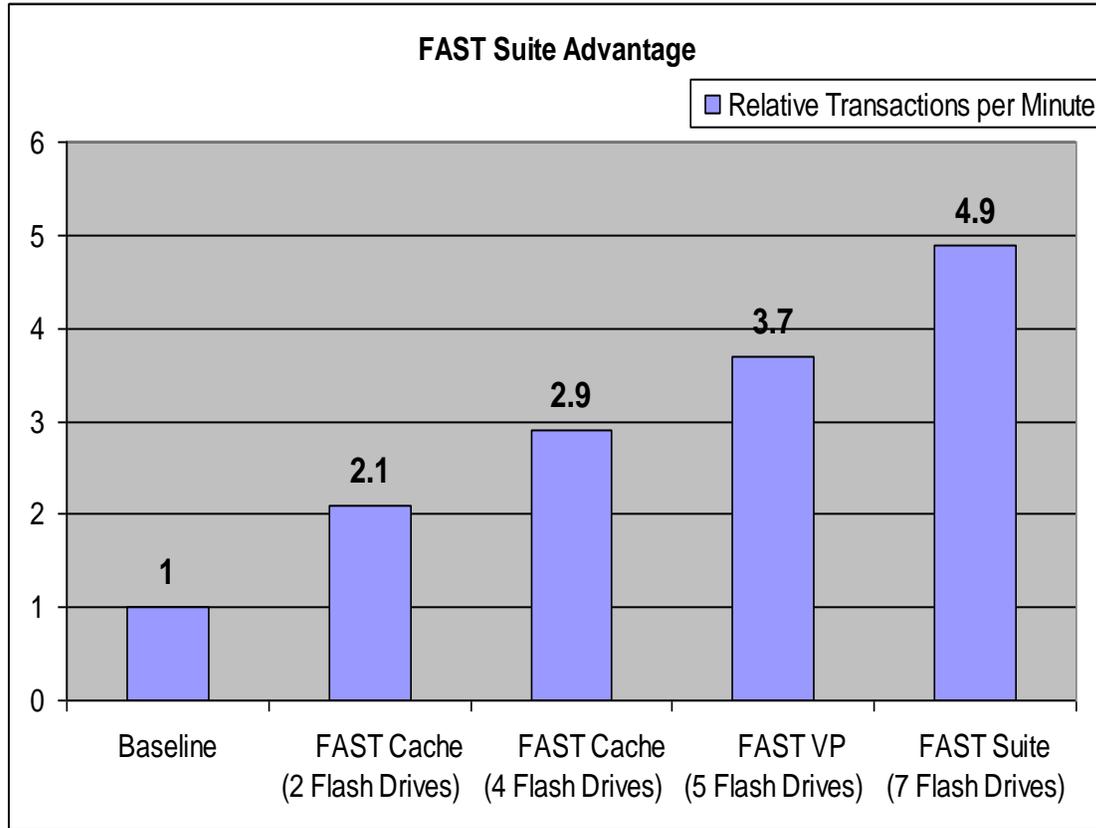


Figure 10. FAST Suite advantage

The figure indicates that once the hot data is cached by FAST Cache, the TPM rate is almost 2.9 times the baseline scenario using four 100 GB Flash drives. In another use case configuration, adding five 100 GB Flash drives to a tier storage configuration and relocating hot data to the Flash tier improves the TPM rate by 3.7 times as compared to the baseline configuration. The best TPM scenario is to add five Flash drives to the pool for holding hot data and to enable FAST Cache on the data pool using two Flash drives to handle sudden intraday spikes over small data ranges. Thus, FAST Suite helps in achieving almost 4.9 times the TPM rate as compared to the baseline. The increased transactions and reduced latency allow the application to scale much better and improve the end-user experience significantly.

Storage considerations for using FAST Suite

Before you choose your system configuration, consider the following:

- To expand the storage pool add as many drives as possible because it minimizes the likelihood of temporarily skewed I/O activity clustering. When you add new

drives to the pool, the system usually assigns the new data coming in to the newly added drives first to ensure that a balanced amount of data is stored over all drives in the pool over time.

- Use thick LUNs with FAST VP for the Sybase Data LUNs and enable FAST Cache (read/write) on the storage pool containing Data LUNs.
- Use a homogeneous pool with only SAS drives for Sybase Log LUNs.
- Include the appropriate tiers in your storage pool for guaranteed performance (for example, Flash and SAS drives for high performance, SAS and NL-SAS drives for high cost savings).
- EMC suggests using these FAST Suite proportions in terms of storage capacity for Sybase deployments:
 - For FAST VP
 - 2 percent Flash drives
 - 16 percent SAS drives
 - 80 percent NL-SAS drives
 - For FAST Cache
 - Up to 2 percent Flash drives in FAST Cache in terms of storage capacity

These recommendations ensure good performance in the case of large data allocation (table expansion, bulk loads) and periodic maintenance (batch jobs, reporting). The tier proportion recommendation may vary with each system depending on the working dataset size, number of active users and locality of the data.

Impact of data placement and locality of data on FAST Suite performance

FAST Cache continuously ensures that the hottest data is served from high-performance Flash drives. FAST VP automatically optimizes storage pools based on relocation schedule, ensuring that only active data is being served from the Flash drives, while cold data is moved to lower-cost disk tiers. High locality of data is important to realize the benefits of FAST VP technology. FAST VP recalculates the activity temperature of all slices in a pool periodically and proposes the movement of data based on this calculation. The relocation of data can be done manually or in a scheduled time window. When the Sybase OLTP application is deployed for the first time, it takes several relocation cycles until all the relevant data is relocated to the appropriate tiers. In addition, initially the data relocation may need to be executed more frequently to speed up the data warm-up process. Since it may take some time before all the data is allocated to an appropriate tier, you will see improved performance with time. However, using the same system for totally different types of work during the day and night will change the locality of data and may have an impact on system performance. FAST Cache can be conveniently enabled/disabled at a

storage pool level depending on the application performance requirements and can be shared across all storage pools in the system.

Conclusion

FAST Suite technology significantly increases the performance of Sybase enterprise applications by optimizing the investment of the Flash tier and improving the capacity utilization of the rotating drives. FAST Suite provides the highest performance and capacity efficiency, and applies to any application with a low I/O latency requirement. Even though FAST Suite can be used with almost any application with any type of I/O pattern, it is especially well-suited for applications that access data with high locality of reference with random I/O and relatively small working sets compared to the total database size. While the performance improvements observed with our scenarios may not be representative of all Sybase OLTP customer scenarios, it nevertheless illustrates the potential for significant improvements in the processing times and better TCO. Also, we were able to improve the Sybase OLTP application performance without tuning the parameters or adjusting the deployment process.

References

The following white papers and documents are available on EMC.com and the Powerlink® website:

- [Leveraging EMC FAST Cache with Sybase OLTP Applications](#)
- [EMC VNX Family Data Sheet](#)
- [EMC VNX Series Unified Storage Systems Specification Sheet](#)
- [EMC VNX Series Total Efficiency Pack Data Sheet](#)
- [Using EMC FAST Suite with SAP on EMC VNX Unified Storage](#)

The following can be found on the Sybase website:

- *“Monitoring performance with sp_sysmon”*
http://infocenter.sybase.com/help/index.jsp?topic=/com.sybase.dc20022_1251/html/monitoring/X99294.htm
- *“Solid State Disks: Why are you still using metal platters???”*
<http://blogs.sybase.com/database/2010/04/solid-state-disks-why-are-you-still-using-metal-platters/>