

EMC VFCACHE ACCELERATES ORACLE

EMC VFCache, EMC VNX, EMC FAST Suite, Oracle Database 11g

- VFCache extends Flash to the server
- FAST Suite automates storage placement in the array
- VNX protects data

EMC Solutions Group

Abstract

This white paper describes how EMC VFCache™ and EMC® FAST Suite on EMC VNX™ storage accelerates Oracle Database 11g online transaction processing (OLTP) performance.

February 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.

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

Part Number: H10501.1

Table of contents

| | |
|--|-----------|
| Executive Summary | 4 |
| Business case..... | 4 |
| The storage performance challenge..... | 4 |
| EMC solution..... | 5 |
| The VFCache solution | 6 |
| Server-side Flash caching for maximum speed | 6 |
| Write-through caching to the array for total protection | 7 |
| Application agnostic..... | 7 |
| Shareable and scalable | 8 |
| Solution architecture..... | 9 |
| Introduction | 9 |
| Physical architecture..... | 9 |
| EMC VNX..... | 9 |
| EMC FAST Suite | 9 |
| EMC FAST Cache | 10 |
| EMC FAST VP | 10 |
| VFCache technology with Oracle Database: testing and validation | 11 |
| Introduction | 11 |
| Workload profile | 11 |
| Recommended configuration | 11 |
| Performance characteristics | 11 |
| Conclusion | 14 |
| References | 15 |
| White papers | 15 |
| Product documentation..... | 15 |

Executive Summary

Business case

In an increasingly competitive environment, businesses are driven to automate business processes and to improve service, while lowering IT costs. Meeting these new requirements has become critical to the success of organizations. Consequently, operational and revenue-generating applications are experiencing dramatic demands on performance, driven by:

- Growth in the numbers of active users
- Escalating service-level agreements
- Increased complexity of business processes and new analytic workloads
- Multiple databases with high concurrent access

Businesses need to consider new approaches to performance challenges in order to meet these demands cost-effectively and without sacrificing data protection.

The storage performance challenge

The latest servers with multi-core processors represent a potential performance bottleneck for the storage subsystem. As processing capacity and heavier workloads are added, the storage system is challenged to keep pace with the growing I/O demands. While CPU performance improves 100-fold every decade, magnetic disk remains relatively flat, as shown in Figure 1.




Figure 1. CPU performance versus disk drives

In a traditional architecture, as shown in Figure 2:

- Reads and writes are serviced by the storage array
- Performance varies depending on the back-end array’s media, workload, and network

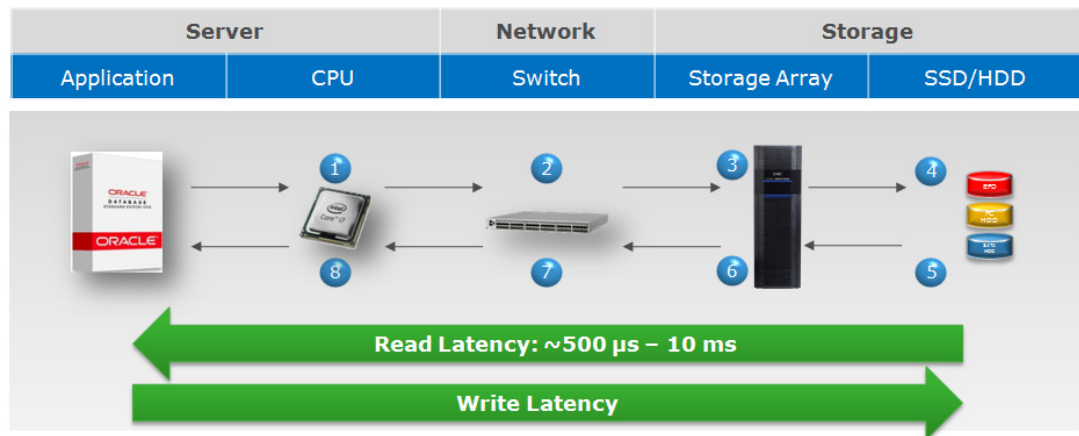


Figure 2. Traditional architecture

What if you could double your application performance by decreasing latency and increasing IOPS and transactions per minute?

EMC solution

EMC VFCache™ is a server Flash caching solution that uses intelligent caching software and PCIe Flash technology to reduce latency and increase throughput, which dramatically improves application performance. With the arrival of VFCache, EMC extends its FAST (Fully Automated Storage Tiering) technologies into the server, adding another tier of intelligence and performance to the I/O stack.

While VFCache is server-based, the FAST software dynamically stores and serves data in the storage array from the most cost-effective and highest performing drive type necessary for the data's activity. Because VFCache automatically identifies frequently accessed read data and promotes that data to the FLASH on the PCIe card, VFCache alleviates I/O pressure from the array, which accelerates application performance. Since frequently accessed data now is located on dedicated storage inside the server, I/O service times are greatly reduced.

VFCache operates autonomously from VNX™ FAST Cache and VNX FAST VP (Fully Automated Storage Tiering for Virtual Pools). VFCache is a dedicated server-side read-only cache that complements FAST Cache's shared array-side cache, which supports both reads and writes. While both caches are transitory, FAST VP moves data between tiers in the array-based storage pool based on predefined customer policies.

All three technologies are designed to work together to ensure that the most frequently accessed information is served with the lowest latency. As shown in Figure 3, VFCache and the FAST Suite, which includes FAST Cache and FAST VP, create the most efficient and intelligent I/O path from the application to the data store. Used in combination, VFCache and the FAST Suite further streamline data delivery across the data continuum, which provides an end-to-end tiering solution that optimizes application capacity and performance from the server to the storage.




Figure 3. VFCache and FAST automatically tier hot and cold data

The VFCache solution

Server-side Flash caching for maximum speed

VFCache software caches the most frequently used data on the server-based PCIe card, thereby putting the data closer to the application, as shown in Figure 4. This reduces the need to access data across the network from the storage array, which both decreases response time and increases performance.




Figure 4. VFCache accelerates I/O performance within the server

VFCache automatically adapts to changing workloads by determining which data is most frequently referenced and promoting it to the server Flash cache. This means that the “hottest” data automatically resides on the PCIe card in the server for faster access.

As shown in Figure 5, the VFCache advanced architecture combines accelerated performance with data protection:

- Reads are serviced by VFCache for performance
- Writes are passed through to the storage array for protection

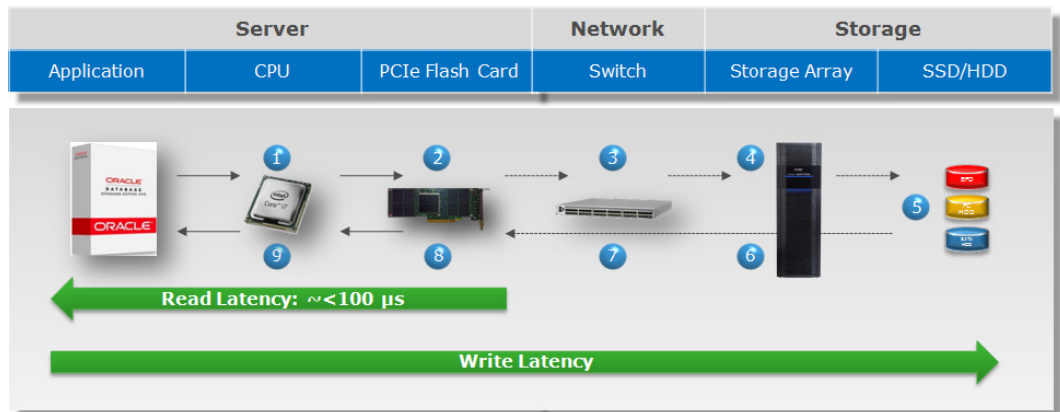


Figure 5. Advanced architecture with EMC VFCache

VFCache is particularly beneficial for read-intensive applications, such as the Oracle Database 11g online transaction processing (OLTP) application used to validate this solution. In the solution tests, VFCache demonstrated read throughput improvements of up to 210 percent and reductions in read latency by as much as 60 percent.

VFCache enhances both virtualized and bare-metal applications so you can smoothly migrate your data center to a private cloud at a pace that makes sense for your business¹.

Write-through caching to the array for total protection

VFCache protects data by using a write-through algorithm, which means that writes persist to the back-end storage array. EMC trusted networked storage, such as the EMC[®] Symmetrix VMAX[™] and the EMC VNX[™] family of storage arrays, protects data and provides additional performance benefits with advanced data services, which include high availability, data integrity, reliability, and disaster recovery.

Application agnostic

VFCache is transparent to applications, so no rewriting, retesting, or recertification is required to deploy VFCache in your environment.

While this white paper focuses on Oracle Database 11g, VFCache works with applications as diverse as databases, analytics, enterprise application servers, email, and web servers to give them the performance boost they need.

¹ VFCache is not supported in shared disk environments or active/active clusters.

Shareable and scalable

While directly enhancing the performance of read-intensive applications, VFCache can indirectly enhance write-intensive applications that now have greater access to SAN resources.

VFCache offloads much of the read traffic from the storage array, which allows VFCache to allocate greater processing power to other applications. While one application is accelerated with VFCache, the array's performance for other applications is maintained or even slightly enhanced. As VFCache is installed on more servers in the environment, the result is a highly scalable I/O processing model. The environment as a whole, including the servers and the storage system, is capable of processing increasingly more IOPS.

Note FAST Suite is a shared resource and performance improvements may depend upon the demands from other applications running on the array.

For more information about EMC VFCache, see the white paper *Introduction to EMC VFCache*.

Solution architecture

Introduction

This section provides an overview of the physical architecture of this solution.

Physical architecture

Figure 6 shows the physical architecture for this solution.




Figure 6. EMC VFCache with Oracle Database and EMC VNX5500

The solution architecture consists of an Oracle Database, Cisco server, and an EMC VNX5500™ storage array. The solution can be configured with various storage arrays, such as Symmetrix VMAXe or VNX. In this solution, EMC used VNX5500. The server used for this test configuration was a Cisco UCS C-250 rack-mount server with two Xeon processors and 12 cores. The server configuration included the VFCache server-based Flash cache and two 8 Gb/s connections to the SAN switches.

EMC VNX

The EMC VNX™ family delivers industry-leading innovation and enterprise capabilities for file, block, and object storage in a scalable, easy-to-use solution. This storage platform combines powerful and flexible hardware with advanced efficiency, management, and protection software to meet the demanding needs of today's businesses.

VNX and VNXe are supported by VFCache when used for block storage access.

EMC FAST Suite

The VNX series has been expressly designed to take advantage of the latest innovation in Flash drive technology. The combination of Flash drives and hard disks deliver improved performance and efficiency while minimizing cost per GB.

The EMC FAST Suite contains the necessary software, specifically FAST Cache and FAST VP, to improve performance and maximize storage efficiency on the VNX. With only a few Flash drives and the FAST Suite, VNX customers can benefit from a “Flash 1st” data strategy, which ensures that that highly active data is stored on and

served from Flash drives for optimal performance, while less active data is tiered on hard disk drives (HDDs) for the lowest total cost of ownership (TCO).

EMC FAST Cache

A caching tier is a large-capacity secondary cache that uses Flash drives positioned between the storage processor's DRAM-based primary cache and HDDs. FAST Cache is a non-disruptive, read/write cache that extends the VNX's existing cache by up to 2 TB. FAST Cache monitors incoming I/O for access frequency and automatically copies frequently accessed data in 64k chunks from the back-end drives into the cache. FAST Cache is easy to administer and cost-effectively provides immediate performance benefits to the system.

EMC FAST VP

FAST VP manages the dynamic tiering of data across a storage pool made up of more than one drive type, which enables optimal disk utilization and efficiency. Based on customer-defined policies, FAST VP's software algorithmically promotes and demotes user data within the pool based on how frequently the data is accessed. More frequently accessed data is moved to higher performance tiers, such as Flash or SAS drives. Infrequently accessed data is moved to modestly performing, high-capacity tiers as needed, such as SAS or NL-SAS drives. Over time, the most frequently accessed data resides on the fastest storage devices, and infrequently accessed data resides on economical bulk storage.

FAST VP provides both capital expenditure and operational expenditure benefits by enabling you to purchase a mixed-drive allocation that results in lower power and cooling costs, a smaller data footprint, and decreased administration time.

Note For this solution, only FAST Cache was used during testing.

VFCache technology with Oracle Database: testing and validation

Introduction

EMC tested the capabilities of VFCache with FAST Cache to accelerate online transaction processing (OLTP) performance in an Oracle Database 11gR2 environment. This section presents the results of the testing.

Note Benchmark results are highly dependent upon workload, specific application requirements, and system design and implementation. Relative system performance will vary as a result of these and other factors. Therefore, this workload should not be used as a substitute for a specific customer application benchmark when critical capacity planning and/or product evaluation decisions are contemplated.

Workload profile

EMC's testing used a standard TPC-C-like OLTP workload, with a 1.2 TB database and a 70/30 percent read/write mix. EMC executed the test using the following steps:

1. Took a performance baseline to validate the performance characteristics of the environment.
2. Enabled VFCache and took measurements of transactions per minute (TPM) and application-level transaction latency to demonstrate the performance contribution of VFCache to the application.
3. Enabled FAST Cache on VNX and took measurements of TPM and application-level transaction latency to demonstrate the performance improvements that can be achieved when both technologies are used together.

Recommended configuration

VFCache is enabled on all of the data LUNs. However, VFCache is not enabled on the log LUNs because they are dominated by writes. The rotating drives can easily deliver the required performance for log LUNs.

No specific tuning is required for Oracle Database.

With this configuration, VFCache uses its caching algorithms to automatically maintain a copy of the hottest data for immediate access.

Performance characteristics

Figure 7 compares the overall system throughput (TPM) of the baseline and VFCache-enabled environments. With the hot data available in the server's VFCache, the TPM was 2.5 times the baseline. As TPM performance improved, there was a 30 percent reduction in read latency.

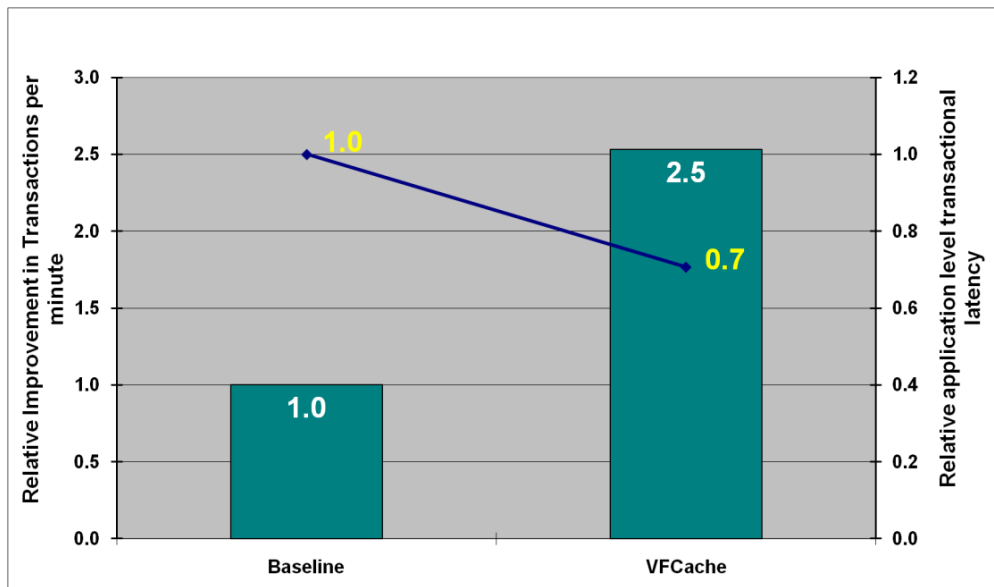


Figure 7. Relative OLTP transactions per minute (TPM) improvement

The addition of FAST Cache resulted in higher performance and lower transaction latency. As shown in Figure 8, TPM was 7.8 times the baseline, while relative application-level transaction latency was 80 percent less than the baseline.

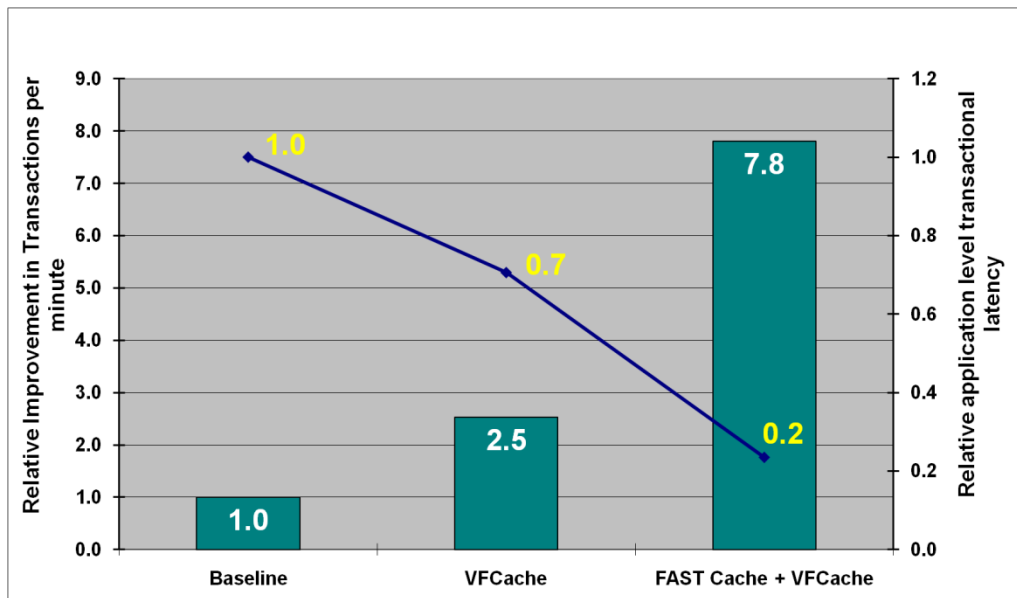


Figure 8. Performance impact of FAST Cache with VFCache

It is important to note that individual customers might see different results. Improvements in application performance depend on a variety of factors, including:

- I/O read-to-write ratio
- Inherent scalability of the workload
- Existing constraints within the storage subsystem, before deploying VFCache
- Tuning of the Oracle database
- Sharing of VFCache, FAST Cache, and FAST VP with other applications

Conclusion

EMC's testing with an Oracle OLTP workload compared a baseline configuration to a system equipped with VFCache and a system equipped with both VFCache and FAST Cache. The tests confirmed the following performance advantages:

- Configured with VFCache alone, system throughput, measured in transactions per minute, was 2.5 times the throughput of the baseline, while the latency was reduced to 30 percent.
- Configured with both VFCache and FAST Cache, system throughput, measured in transactions per minute, was 7.8 times the throughput of the baseline, while latency was reduced to 20 percent.
- No changes to the database were required.
- VFCache and FAST Cache maintained the integrity of the data.

References

White papers

For additional information, see the white papers listed below:

- *Introduction to EMC VFCache*
- *EMC FAST Cache - A Detailed Review*
- *EMC VFCache Accelerates Microsoft SQL Server - EMC VFCache, EMC VNX, Microsoft SQL Server 2008*
- *EMC VFCache Accelerates Oracle - EMC VFCache, EMC Symmetrix VMAX and VMAXe, Oracle Database 11g*
- *EMC VFCache Accelerates Virtualized Oracle - EMC VFCache, EMC Symmetrix VMAX and VMAXe, VMware vSphere, Oracle Database 11g*

Product documentation

For additional information, see the product documents listed below:

- *EMC VFCache* - data sheet
- *EMC VNX Family* - data sheet
- *EMC VNX Series Total Efficiency Pack* - data sheet