

EMC XTREMCACHE ACCELERATES VIRTUALIZED ORACLE

EMC XtremSF, EMC XtremCache, EMC Symmetrix VMAX,
EMC FAST VP, VMware vSphere, Oracle Database 11g

- XtremCache improves Oracle performance in targeted applications
- Symmetrix VMAX protects data
- FAST VP automates storage placement in the array

EMC Solutions Group

Abstract

This white paper describes how EMC® XtremCache™ and EMC FAST™ VP on EMC Symmetrix® VMAX® storage with VMware vSphere accelerates online transaction processing (OLTP) performance in a virtualized Oracle environment. Though the testing was performed on VMAX 40K, the capabilities and benefits of VFCache are applicable to VMAX 10K and VMAX 20K.

December 2013

vmware[®]
PARTNER
TECHNOLOGY
ALLIANCE

Copyright © 2013 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: H10664.1

Table of contents

Executive summary	4
Business case.....	4
The storage performance challenge.....	4
The solution.....	5
Technology components	6
EMC XtremCache.....	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
EMC FAST VP.....	8
EMC Symmetrix VMAX.....	8
EMC Symmetrix VMAX 40K series with Enginuity 5876.....	8
VMware vSphere.....	9
XtremCache and FAST VP in this solution	10
Working together.....	10
Solution architecture	11
Introduction.....	11
Physical architecture.....	11
XtremCache with virtualized Oracle Database: testing and validation	12
Introduction.....	12
Workload profile.....	12
Recommended configurations.....	13
XtremCache.....	13
Oracle.....	13
VMware vSphere.....	13
Performance characteristics.....	14
Test results.....	14
Conclusion	16
References	17
White papers.....	17
Product documentation.....	17

Executive summary

Business case

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

- Growth in the numbers of active users
- Time-critical applications and escalating service-level agreements (SLAs)
- Increased complexity of business processes and new analytic workloads
- Multiple databases with high concurrent access

Businesses can use virtualized server platforms to consolidate workloads from multiple departments. As the workloads are consolidated, businesses may find that the preceding performance factors are affecting one or more of the workloads, which impairs these businesses' ability to meet performance goals or SLAs.

Businesses need to consider new approaches to performance challenges in order to meet individual workload demands cost-effectively and without sacrificing data protection. A solution that enables businesses to achieve high performance in targeted workloads, without any change to applications, can address these needs.

The storage performance challenge

The latest servers with multi-core processors may find that a conventional storage subsystem represents a potential performance bottleneck. 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. Fortunately, flash technologies have emerged with innate solid-state memory capabilities to fill the performance gap between disk drives and server processing power.




Figure 1. CPU and disk-drive performance improvements

In a traditional architecture, as shown in Figure 2:

- The storage array services read and write I/O requests.
- Performance varies depending on the back-end array's media, workload, and network.




Figure 2. Traditional architecture

What if you could improve your application performance by decreasing latency and increasing IOPS and transactions per minute (TPM)? This solution enables you to do that.

The solution

One of the ways to improve performance is to focus flash technologies on high-demand applications and data.

EMC® Symmetrix® VMAX® Fully Automated Storage Tiering™ for Virtual Pools (FAST VP) dynamically tiers data across a storage pool of flash drives and other drive types. FAST VP moves the most frequently accessed data to the flash tier and less frequently accessed data to the slower storage tiers. This improves application performance and increases IOPS and transactions per minute, while reducing latency.

EMC XtremCache™ is a server flash caching solution that extends flash technology into the server, which adds another tier of intelligence and performance to the I/O stack. XtremCache automatically identifies frequently accessed read data and promotes that data to a flash cache in the server. This greatly reduces I/O service times and alleviates I/O load on the array. XtremCache also provides the ability to selectively accelerate specific workloads, which enables the system manager to adjust performance to meet critical SLAs.

For the solution described in this white paper, EMC tested the ability of XtremCache to improve the performance of a targeted database in order to meet aggressive SLA goals. With both XtremCache and FAST VP enabled, we observed a performance improvement of 330 percent for the targeted database. The performance of the other databases also improved significantly.

Technology components

EMC XtremCache EMC XtremCache is a server flash caching solution that reduces latency and accelerates throughput to dramatically improve application performance by using intelligent caching software and PCIe flash technology.

XtremCache accelerates reads, and protects data by using a write-through cache to the networked storage to deliver persistent high availability, integrity, and disaster recovery.

XtremCache coupled with array-based FAST VP provides the most efficient and intelligent I/O path from the application to the data store. The result is a storage infrastructure that is dynamically optimized for performance, intelligence, and protection for both physical and virtual environments.

Server-side flash caching for maximum speed

XtremCache software caches the most frequently used data on the server-based PCIe card (EMC XtremSF™), which puts the data closer to the application, as shown in Figure 3. This reduces the need to access data across the storage area network (SAN) from the storage array, which decreases response time and increases throughput.




Figure 3. XtremCache accelerates I/O performance within the server

XtremCache caching optimization 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 4, the XtremCache advanced architecture combines accelerated performance with data protection:

- XtremCache services reads for performance
- Writes are passed through to the storage array for protection
- XtremCache helps write performance by offloading the majority of the read workload from the array to the server




Figure 4. Advanced architecture with EMC XtremCache

XtremCache provides better performance for read-intensive applications. XtremCache works with applications as diverse as databases, analytics, enterprise application servers, email, and web servers to give them the performance boost they need. XtremCache can be deployed to accelerate all applications on a server, selected applications, a single application, or specific LUNs to meet application performance goals.

XtremCache 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 business.

Write-through caching to the array for total protection

XtremCache protects data by using a write-through algorithm, which means that writes persist to the back-end storage array. EMC networked storage, such as the Symmetrix VMAX series and the EMC VNX™ family of storage arrays, protect data with advanced data services, which include high availability, data integrity, reliability, and disaster recovery.

Application agnostic

XtremCache is transparent to applications. When XtremCache is present, applications do not have to be rewritten, recertified, or retested.

While this white paper focuses on Oracle Database 11g, the XtremCache architecture can directly enhance the performance of many other applications and indirectly enhance write-intensive applications that now have greater access to SAN resources.

Shareable and scalable

XtremCache offloads a lot of the read I/O from the storage array, which enables it to allocate greater processing power to other applications. While one application is accelerated with XtremCache, the array's performance for other applications is enhanced. As XtremCache 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 increasingly capable of processing more transactions.

For more information about XtremCache, refer to *Introduction to EMC XtremCache*.

EMC FAST VP

The Symmetrix VMAX family is designed to take advantage of the latest innovations in flash drive technology. The combination of flash drives and hard-disk drives (HDDs) delivers improved performance and efficiency while minimizing cost per gigabyte (GB).

FAST VP manages the dynamic movement of data between tiers to provide optimal drive utilization and efficiency. Based on customer-defined policies, FAST VP's software algorithmically promotes and demotes user data between specific drive tiers, based on how frequently the data is accessed. More frequently accessed data is moved to higher performing drive tiers, such as flash or 10K/15K Fibre Channel (FC) and SAS drives. Infrequently accessed data is moved to modestly performing, high-capacity drive tiers as needed, such as 7.2L SATA or NL-SAS drives.

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

Note: FAST VP is a shared resource. Performance improvements may depend on the demands from other applications running on the array.

EMC Symmetrix VMAX

The EMC Symmetrix VMAX series of storage arrays provide high performance and scalability for demanding enterprise storage environments. Built on the strategy of simple, intelligent, modular storage, VMAX incorporates a highly scalable Virtual Matrix Architecture™ that enables the storage environment to grow seamlessly and cost-effectively from an entry-level configuration into the world's largest storage system.

VMAX supports flash drives, NL-SAS drives, SAS drives, FC drives, and SATA drives within a single array, as well as an extensive range of RAID types

EMC Symmetrix VMAX 40K series with Enginuity 5876

EMC Symmetrix VMAX 40K builds on the EMC Symmetrix storage array environment for powerful, trusted, smart storage that provides higher levels of performance, availability, and intelligence in the virtual data center.

Symmetrix VMAX 40K is the industry's highest-performance, highest-capacity storage system and is designed for the hybrid cloud and other demanding data center environments. Symmetrix VMAX 40K offers:

- Twice the performance, capacity, and global memory (cache) of Symmetrix VMAX 20K
- High-density configurations (using 2.5" SAS drives)
- System-bay and storage-bay dispersion that provides flexibility for space-constrained data centers

Symmetrix VMAX 40K and the Enginuity™ 5876 operating environment offer management-automation and Data Mobility technologies that are critical to the delivery of cloud-based infrastructure and IT services.

Enginuity 5876 technology for networking, sharing, and tiering storage helps data centers to consolidate applications and deliver new levels of efficiency through higher utilization rates, improved mobility, and simplified storage management.

VMware vSphere

VMware vSphere uses the power of virtualization to transform data centers into simplified cloud computing infrastructures, and enables IT organizations to deliver flexible and reliable IT services. vSphere virtualizes and aggregates the underlying physical hardware resources across multiple systems and provides pools of virtual resources to the data center.

As a cloud operating system, vSphere manages large collections of infrastructure (such as CPUs, storage, and networking) as a seamless and dynamic operating environment, and also manages the complexity of a data center.

XtremCache and FAST VP in this solution

Working together With XtremCache, EMC extends its FAST architecture into the server, which adds another tier of intelligence and performance to the I/O stack. XtremCache helps to identify the data that is most frequently accessed by the application and moves that data one step closer to the application inside the server.

While XtremCache is server-based, FAST 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. XtremCache automatically identifies frequently accessed read and write data and promotes a copy of that data to the flash on the XtremSF card inside the server. XtremCache offloads the majority of read workload from the array, which accelerates application performance. Because a copy of frequently accessed data is located in the cache inside the server, I/O service times are greatly reduced.

Both technologies are designed to work together to ensure that the most frequently accessed information is served with the lowest latency. As shown in Figure 5, XtremCache and FAST VP create the most efficient and intelligent I/O path from the application to the data store. Used in combination, XtremCache and FAST VP 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 5. XtremCache and FAST automatically tier hot and cold data

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. Database consolidation benchmark setup

The solution architecture consists of an Oracle Database 11g single instance, three Cisco servers, and a four-engine Symmetrix VMAX storage array. The solution supports a variety of EMC storage arrays, such as VMAX 10K or VNX. In this solution, we used VMAX 40K.

The three servers used for this test configuration were Cisco UCS C-460/M2 rack mounts with four Xeon processors and 40 cores each. The servers were connected to the array through two SAN switches. One server was deployed with a targeted database and included a VFCache card.

The system was virtualized using vSphere 5 and run on Red Hat Enterprise Linux 5. The virtual machines each had 24 vCPUs and 160 GB of memory. Additional virtual machines, if configured, could share the same XtremCache flash cache.

XtremCache with virtualized Oracle Database: testing and validation

Introduction

Virtual server environments represent an opportunity to increase server scalability and utilization while centralizing and streamlining system management.

EMC tested the capabilities of XtremCache with FAST VP to accelerate online transaction processing (OLTP) performance in a virtualized Oracle Database 11gR2 environment using VMware vSphere 5. This section presents the results of the testing.

Note: Benchmark results are highly dependent on 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

Multiple workloads were used during testing of this solution, as described in Table 1.

Table 1. Workload profiles

Configuration	Custom DB	HR DB	ERP DB
Database size	700 GB	1.2 TB	1.2 TB
Concurrent users	30	20	50

The three databases resided on the same set of disks (two separate pools: Data and Log).

The use case assumed that the “ERP DB” database was mission-critical for the customer. FAST VP on the array and XtremCache on the server were applied only to those LUNs hosting ERP DB. The other two databases continued to run from the HDDs.

We ran several test scenarios, as described in Table 2.

Table 2. Test scenarios

Test scenario	Description
Baseline	Took a performance baseline to validate the performance characteristics of each workload in the environment.
With FAST VP on ERP DB	Enabled FAST VP on ERP DB, then measured the TPM and application-level transaction latency to show the performance contribution of FAST VP to the application.
With XtremCache and FAST VP on ERP DB	Enabled both XtremCache and FAST VP on ERP DB, then measured the TPM and application-level transaction latency. This shows the performance improvements that can be achieved when both technologies are used.

Recommended configurations

XtremCache

In a virtualized environment, multiple virtual machines on the same server may share the performance advantages of a single XtremCache solution. However, as this solution shows, all the resources of the XtremCache solution can be directed at a single application to maximize its performance boost. As shown in Figure 7, the XtremSF card resides on the server host, while XtremCache caching software (indicated as small blue icons) is installed on each of the virtual machines that will be accelerated by XtremCache. The Virtual Storage Integrator (VSI) plug-in for XtremCache, which resides on the vCenter client, is used to manage XtremCache.

Note: Figure 7 represents a more general case in which multiple virtual machines are deployed per server and each server configuration includes an XtremCache installation.




Figure 7. XtremCache implementation in a virtualized environment

We enabled XtremCache on all the data LUNs of the targeted application. However, XtremCache was not enabled on the log LUNs because they are dominated by writes. With this configuration, XtremCache uses its caching algorithms to automatically maintain a copy of the hottest data for immediate access.

XtremCache copies and stores the hottest read data on the XtremSF card within the server. The application can repeatedly read this data without sending the I/O request to the storage array. The target workload exhibited a 70/30 percent read/write mix.

Oracle

No specific tuning was required for Oracle Database.

VMware vSphere

Figure 8 shows the virtual machine disk (VMDK) layout of the physical disks. EMC recommends that you isolate the XtremCache VMDK to its own virtual Host Bus Adapter (vHBA).

In general, ESX servers can be configured with one XtremSF card supporting multiple virtual machines on the server. In this solution, we configured one virtual machine per server.




Figure 8. VMDK layout of physical disks

Performance characteristics

Our test environment employed three workloads, deployed on three different servers, that were connected to a single VMAX 40K array with FAST VP enabled.

Test results

We executed the workloads simultaneously, and we took measurements to establish a baseline and then measured the performance improvement that was delivered with XtremCache and FAST VP.

In our testing, ERP DB represented the workload with the most demanding SLA requirement, and we focused the XtremCache resources on this application.

Figure 9 shows the same results as relative OLTP transactions per minute. ERP DB had a significant performance improvement of 3.3 times the baseline results. HR DB also had a big performance improvement of 2.2 times the baseline results as VMAX resources were freed up when the back-end load moved away from the storage array for ERP DB. Custom DB had a performance improvement of 1.2 times the baseline results for the same reason.




Figure 9. Relative OLTP transactions per minute improvement

Transaction latency on ERP DB, the critical database, decreased by 66 percent, while the transaction throughput increased threefold.

Note: Customers can 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 XtremCache
- Tuning of the Oracle database
- Sharing of XtremCache and FAST VP with other applications

Conclusion

EMC's testing with a virtualized Oracle OLTP workload compared a baseline configuration with consolidated databases to the same system with FAST VP enabled and focused on the most critical database. We then compared these configurations to the system with FAST VP and XtremCache focused on the most critical database.

When we applied FAST VP to a single database, performance improved significantly. The performance of the other databases also improved as the load was taken off the HDDs and promoted to the flash drives.

Similarly, when XtremCache was also applied to the same single database, there was a dramatic throughput increase on the targeted database. There was also a measureable increase in the other databases because resources were freed up when the back-end load moved away from the storage array for the targeted database.

Table 3 shows the throughput improvements achieved on each of the databases.

Table 3. Throughput improvements

Workload	With Fast VP System throughput increase (%)	With FAST VP and XtremCache System throughput increase (%)
Custom DB	120	120
HR DB	120	220
ERP DB*	220	330

* Critical database

No changes to the databases were required to achieve these throughput increases. At the same time, XtremCache and FAST VP maintained the integrity of the data.

References

White papers

For more information, see the following white papers:

- *Introduction to EMC XtremCache*
- *EMC XtremCache Accelerates Oracle—EMC XtremSF, EMC XtremCache, EMC Symmetrix VMAX and VMAXe, Oracle Database 11g*
- *EMC XtremCache Accelerates Oracle—EMC XtremSF, EMC XtremCache, EMC VNX, EMC FAST Suite, Oracle Database 11g*
- *EMC XtremCache Accelerates Microsoft SQL Server—EMC XtremSF, EMC XtremCache, EMC VNX, Microsoft SQL Server 2008*
- *EMC Mission-Critical Infrastructure for Microsoft SQL Server 2012 (Accelerated with XtremCache)—EMC Symmetrix VMAX 10K, EMC FAST VP, SQL Server AlwaysOn Availability Groups, VMware vSphere*

Product documentation

For more information, see the following product documents:

- *EMC XtremCache Data Sheet*
- *EMC Symmetrix VMAX Data Storage System Specification Sheet*
- *VMware vSphere 5 Data Sheet*
- *VMware vSphere Basics ESXi 5.0*