



Accelerate Applications Using EqualLogic Arrays with directCache

Abstract

This paper demonstrates how combining Fusion ioMemory products with directCache software in host servers significantly improves the performance of applications and databases that use EqualLogic array storage. The test results show how directCache accelerates read requests and, by removing the read burden off the EqualLogic array and onto host servers, also accelerates write request performance on the array.

Introduction

directCache was designed to improve read performance and reduce read latency by offloading read request from the EqualLogic array onto flash memory. This paper reveals:

- Why adding directCache to your environment can improve the overall performance and health of an EqualLogic array.
- How this solution translates into improved real-world performance of your application.

Solution

Fusion directCache software transforms ioMemory into a transparent, auto-tiering acceleration cache to non-disruptively improve physical storage performance and efficiency. directCache software seamlessly integrates with existing block storage devices requiring no changes to the storage infrastructure, preserving existing SAN and DAS configurations. Host operating systems continue to see and use block devices as-is. Applications receive maximum benefit with minimum disruption since only the most frequently accessed blocks are cached providing instant performance gains. Existing data protection strategies remain in-place without disruption.

ABOUT THE TESTS

A series of benchmarks were run using Flexible IO (FIO) Tester. Each test execution was a variation of different block sizes, threads, and queue depths. For a baseline, all of the test configurations on the EqualLogic array were run without directCache. After the baseline

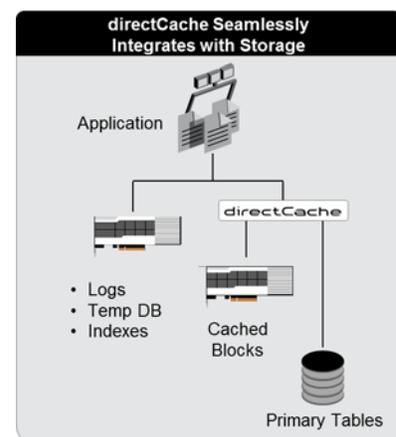


Figure 1: directCache seamlessly integrates with existing storage infrastructure.



performance metrics were captured, directCache was installed on the server and volumes from the EqualLogic Array were bound to the ioDrive2 using the directCache software. After binding the volumes, the baseline test sequence was executed with directCache enabled.

In addition to synthetic benchmarking, “TPC-C-like” queries were used to generate a database workload. TPC-C was chosen due to its transactional nature, which better reflects a real-world workload.

CONFIGURATION

The EqualLogic PS6010VX was configured with the EqualLogic software on the host server to leverage the multipath driver provided by Dell. Jumbo packets on the Dell Force10 S4810 switch and on the Broadcom NICs were enabled to increase network throughput. A quick benchmark shows the array was able to read 800MB/s of sequential reads and 800/MBs of sequential writes. The functionality of the multipath driver was verified by observing that the iSCSI traffic was split between both ports on the Broadcom NIC. The multipath failover policy was set to Least Queue Depth, which is the recommended setting. The 16 - 300GB 15k RPM disks were striped in a RAID 0 configuration for performance. Finally, a logical volume of 1.6TB was created, which is approximately one-third the raw capacity of the EqualLogic array.

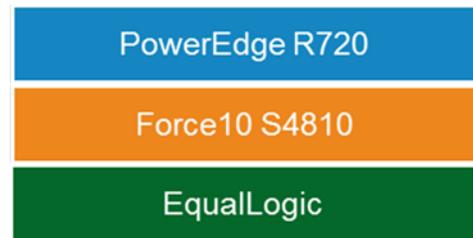


Figure 2: Test software stack configured for maximum performance.

By design, the FIO benchmarking tool enforces the given read/write ratio parameter. This rule can make it difficult to benchmark disks bound to directCache. For example, if you have a single spindle capable of 100MB/s of random read/write requests and set the read/write ratio to 75/25, the best write performance we can achieve under the ratio is 25MB/s and the best read performance would be 75MB/s. Binding the disk to directCache would provide read speeds in excess of 75%. Since the read/write ratio is forced, only 75MB/s of read performance coming from directCache are achieved. As a workaround, the team setup the main benchmarking job to execute 4 child jobs simultaneously: 3 jobs for random reads and 1 job for random writes. Although this technique does not enforce a strict read/write ratio it:

- Allows each job to run independently
- Pushes higher bandwidth and IO requests onto directCache and the EqualLogic array

The single write job was able to achieve bandwidth and IO request numbers close to the baseline tests.

Test Results

The first set of results establishes a performance baseline for the EqualLogic array and directCache. The TPC-C test results demonstrate directCache’s intelligent algorithms and interaction with non-synthetic data.



FUSION-IO®

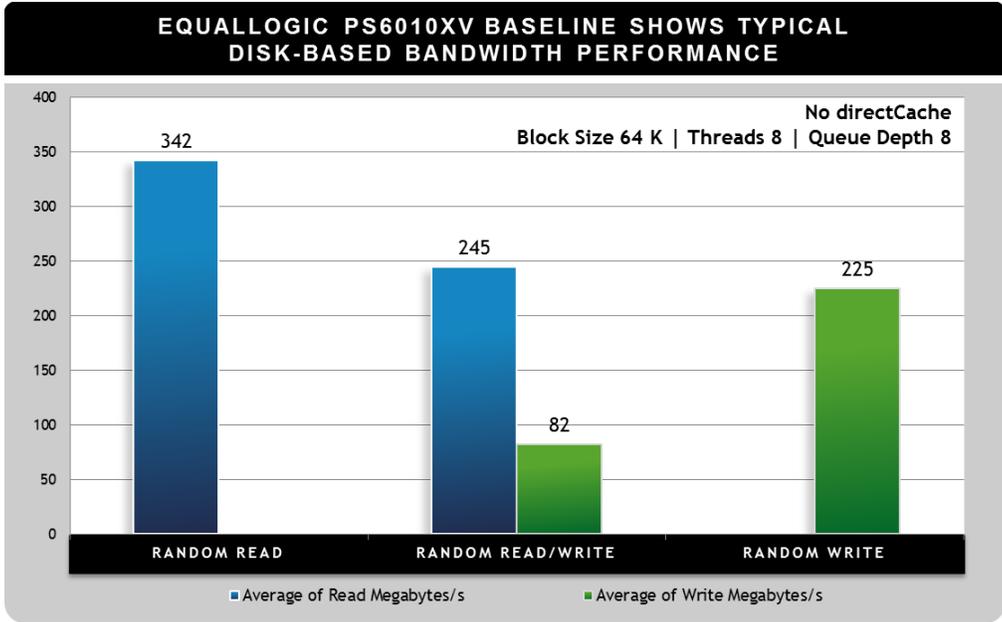


Figure 3: Performance baseline shows typical disk-based performance

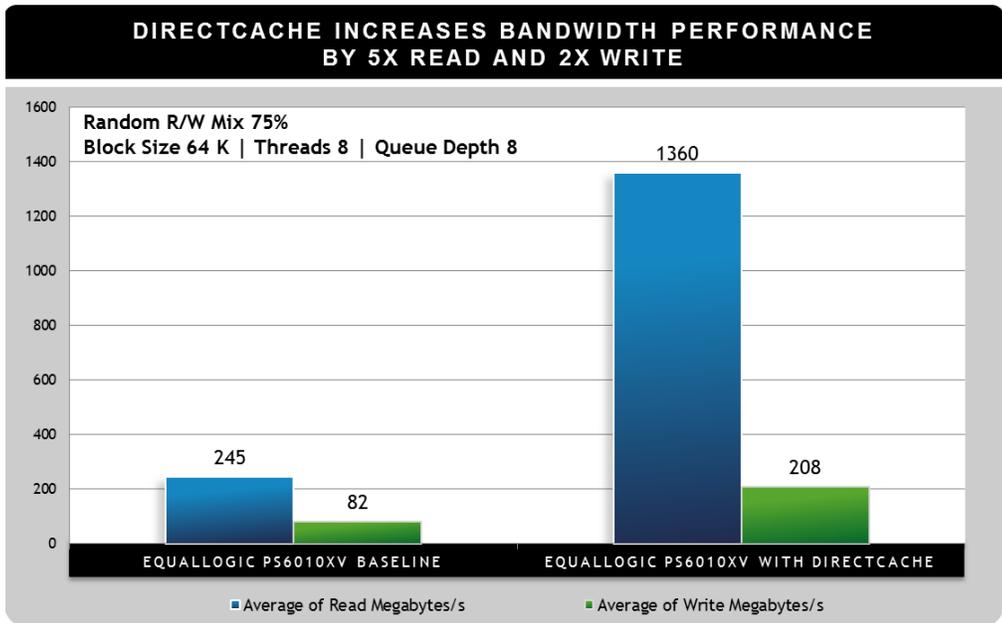


Figure 4: directCache increases read performance by 5x and write performance by 2x over disk-based storage.



FUSION-IO®

DIRECTCACHE INCREASES IOPS PERFORMANCE BY OVER 17X READ AND 2X WRITE

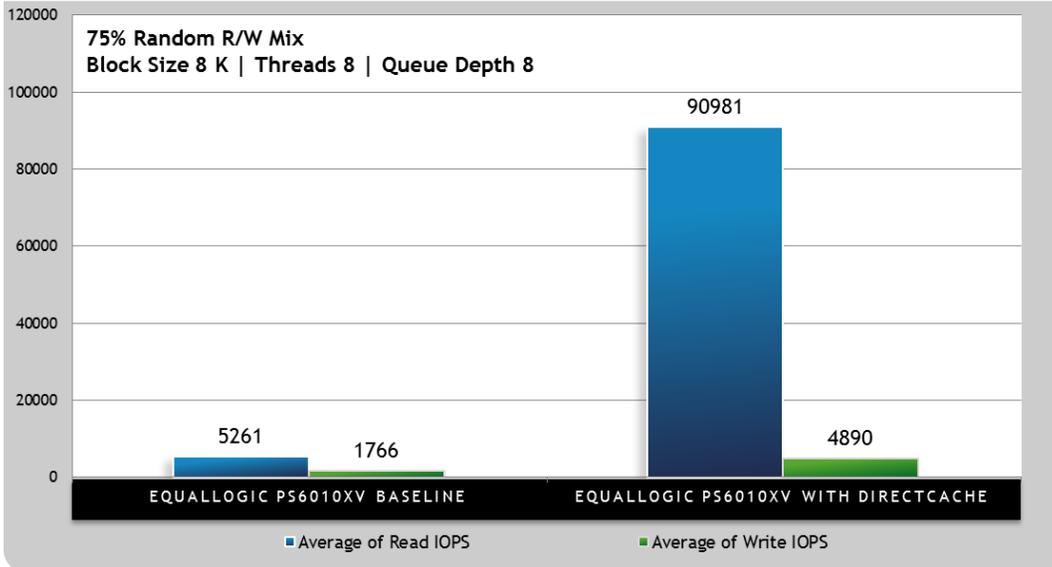


Figure 5: directCache improves IOPS performance more than 17X Read and 2X Write over disk-based systems.

DIRECTCACHE DECREASES LATENCY 22X FOR READS AND ALMOST 2X FOR WRITES

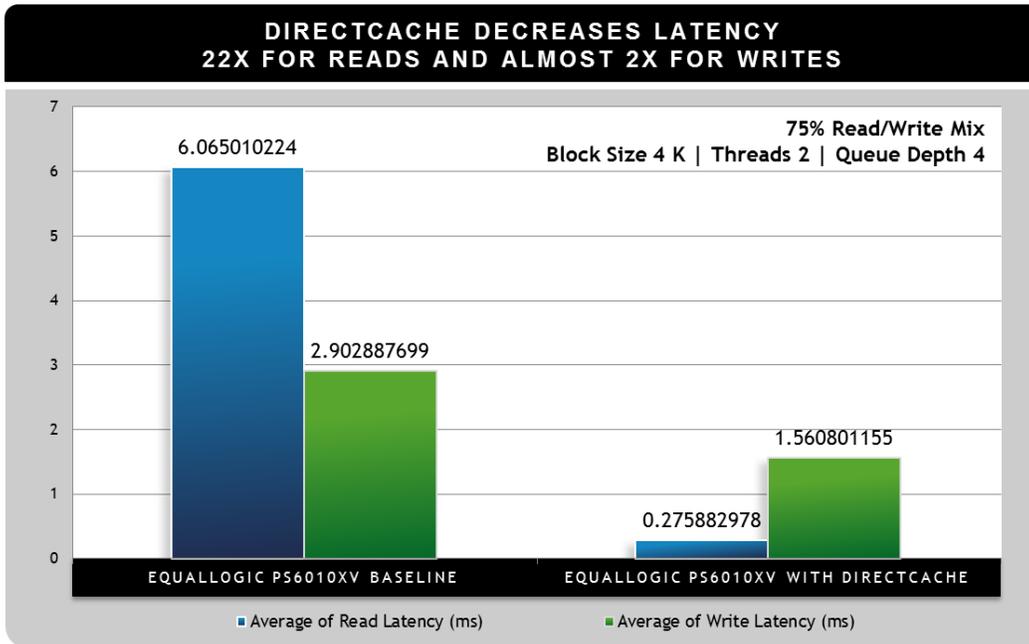


Figure 6: directCache decreases latency performance by 22X for reads and almost 2X for writes.



FUSION-io®

REAL WORLD TEST RESULTS

To further demonstrate the effectiveness of directCache we leveraged HammerDB to simulate a TPC-C workload. HammerDB is an open source load testing tool which works with multiple database platforms. We configured HammerDB with the following test parameters. To learn more about the parameter definitions in HammerDB, visit http://hammerora.sourceforge.net/hammerora_mssqlserver_oltp_v2.7.pdf.

HammerDB Parameters

Name	Value
Duration	45 minutes
Number of Warehouses	200
Virtual Users to Build Schema	100
Virtual Users	500
User Delay (ms)	50
Repeat Delay (ms)	50
Minutes of Ramp Time	1
Keying and Thinking Time	False

Table 1: HammerDB parameters created to simulate TPC-C, a more transactional "real-world" workload.

Database Parameters

Name	Value
TPC-C Database data files	1 per core
Tempdb data files	1 per core
Max Degree of Parallelism	2
Minimum Server Memory (MB)	2000
Maximum Server Memory (MB)	2000
Lock Pages in Memory	Enabled for the SQL Server service account

Table 2: Database parameters set for TCP-C workload.

Note: In a production environment most database servers have much more memory. The SQL Server Minimum and Maximum memory parameter were limited to accelerate testing and to illustrate directCache's benefits on disk I/O.



FUSION-IO®

Storage Configuration

To isolate the TPC-C database data files, three physical volumes were created from the RAID 0 storage pool for the following:

- Tempdb
- TPC-C database transaction log
- TPC-C database data files.

This approach isolated the TPC-C database data files for caching. Caching tempdb is optional and transaction log is not recommended. We did not cache tempdb for this paper.

DATABASE CACHING RESULTS

The following charts demonstrate directCache's ability to intelligently cache active data. HammerDB was run for 45 minutes to capture the results. SQL Server was restarted between each run to clear DMV's and the memory buffer.

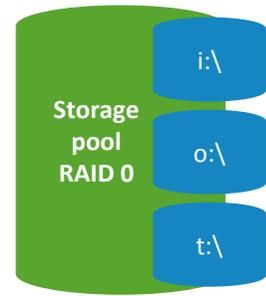


Figure 7: RAID 0 storage pool isolates TPC-C data files for caching.

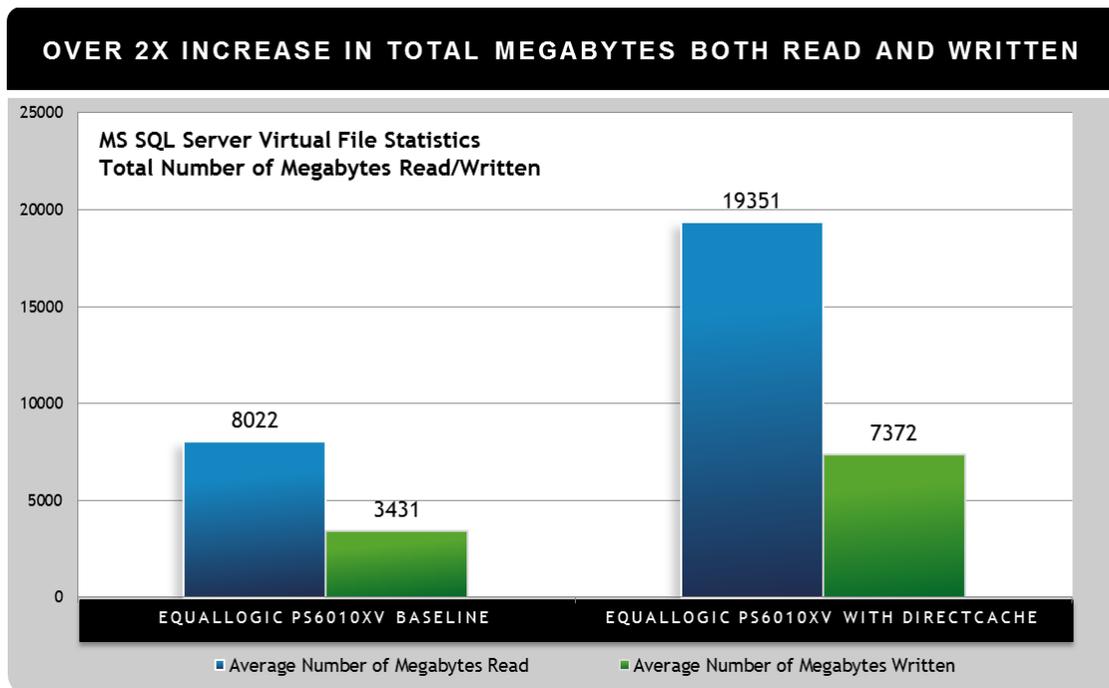


Figure 8: Intelligent caching allows over 2x increase in total megabytes read and written.



FUSION-io®

2X DECREASE IN TOTAL IO STALLS (QUEUING)

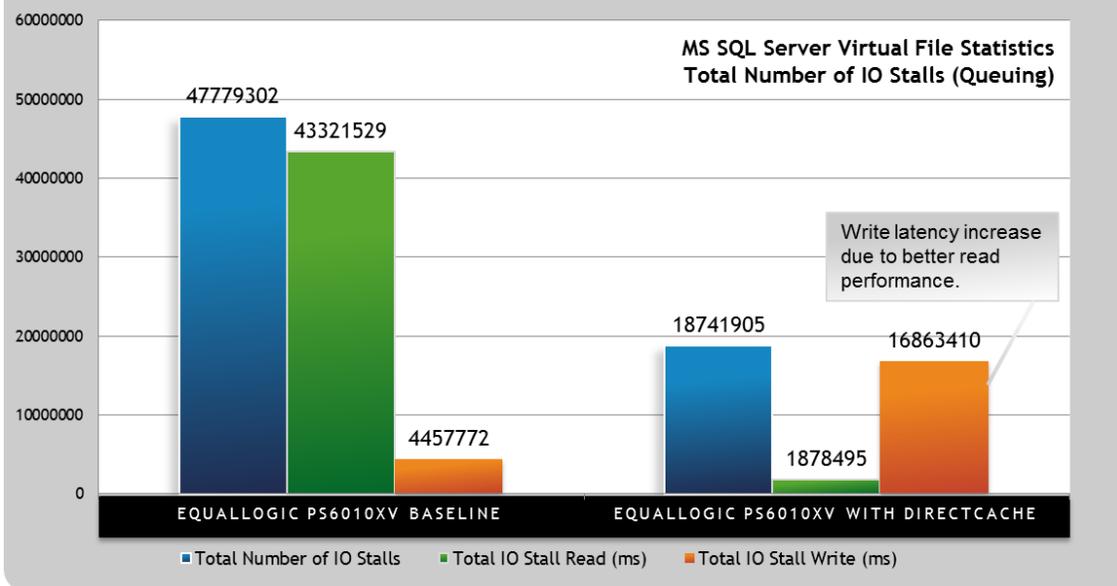


Figure 9: 2X decrease in total IO Stalls (Queuing) achieved by balancing 23X decrease in IO Stalls during reads with 4x increase during writes.

OVER 2X INCREASE READ AND 50% INCREASE WRITE IO OPERATIONS

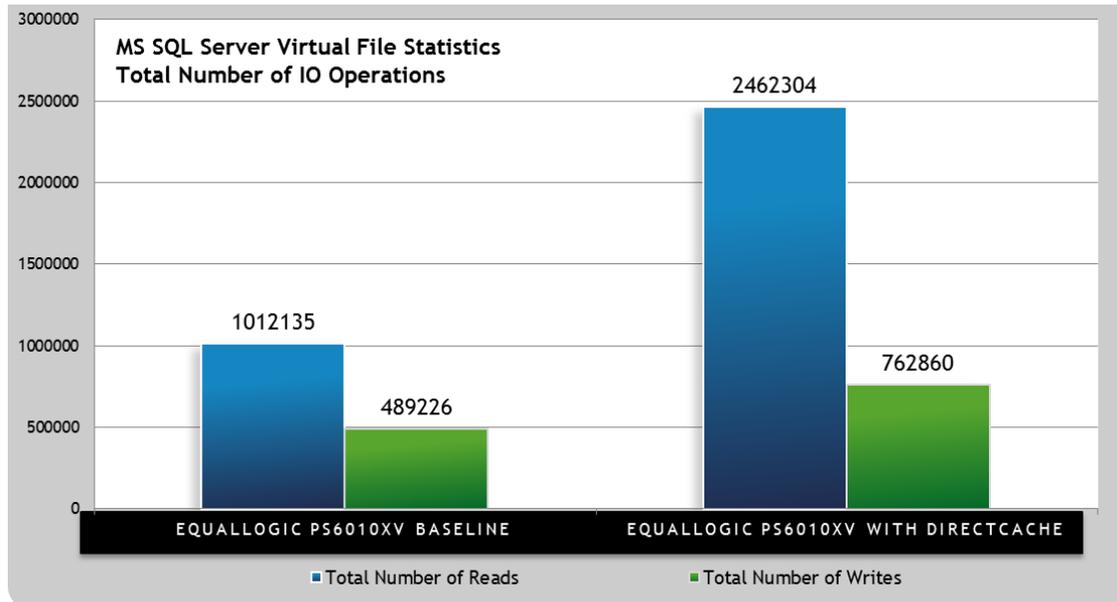


Figure 10: Over 2x increase in read and 50% increase in write IO Operations.



FUSION-io®

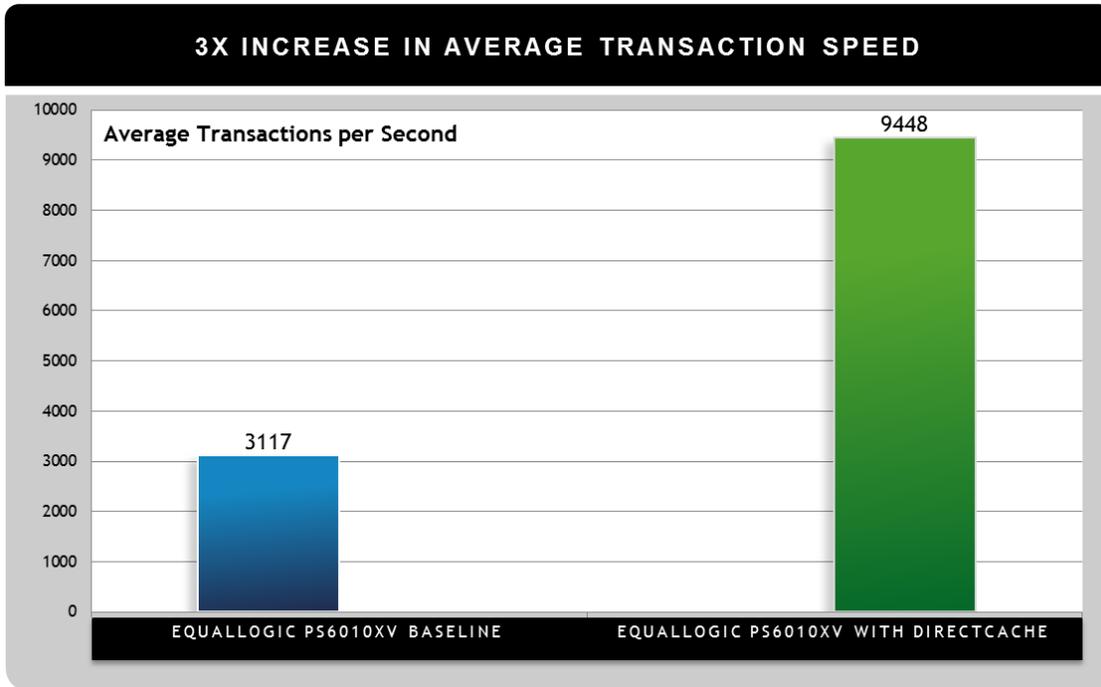


Figure 11: Average Transactions speed increases three times baseline transaction speed.

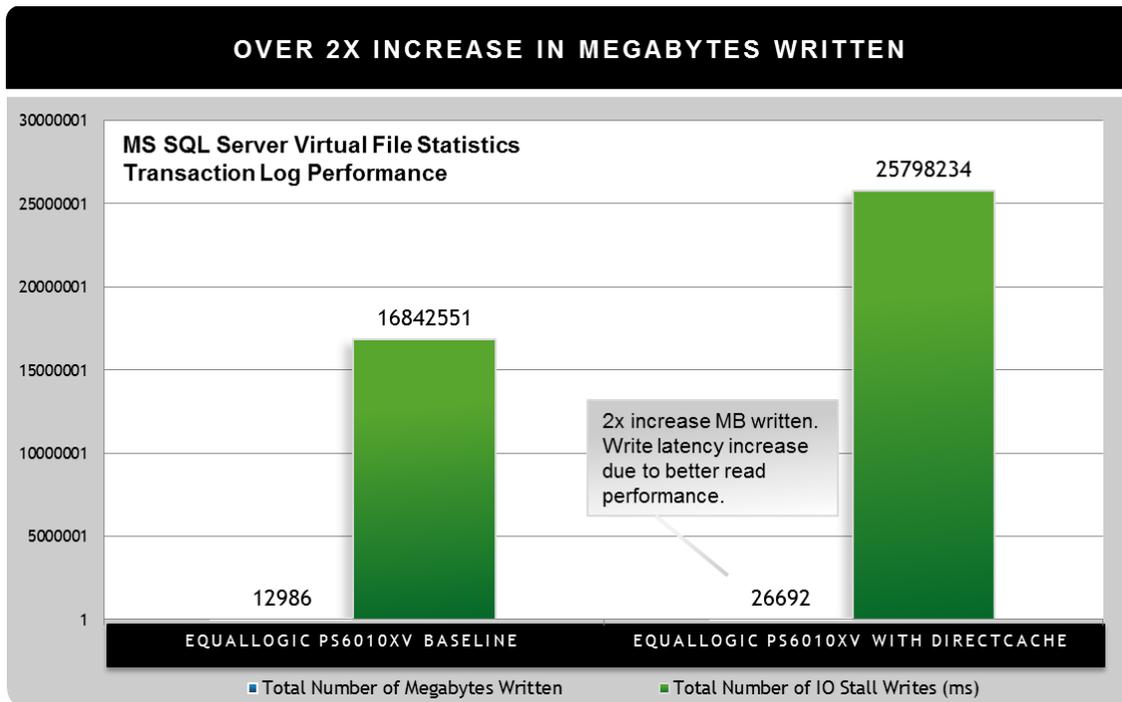


Figure 12: Over 2x increase in megabytes written using directCache



FUSION-IO®

UP TO 2X TOTAL WRITE PERFORMANCE INCREASE ACHIEVED

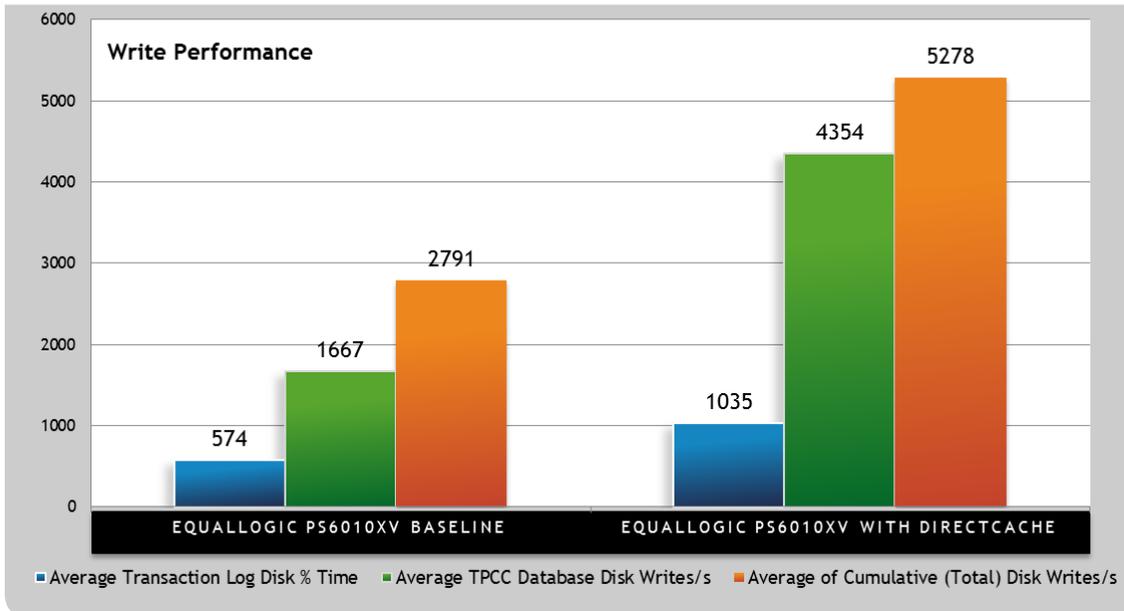


Figure 13: Over 2x write performance increase achieved.

PERFORMANCE INCREASES AND LATENCY DECREASES AS CACHE WARMS

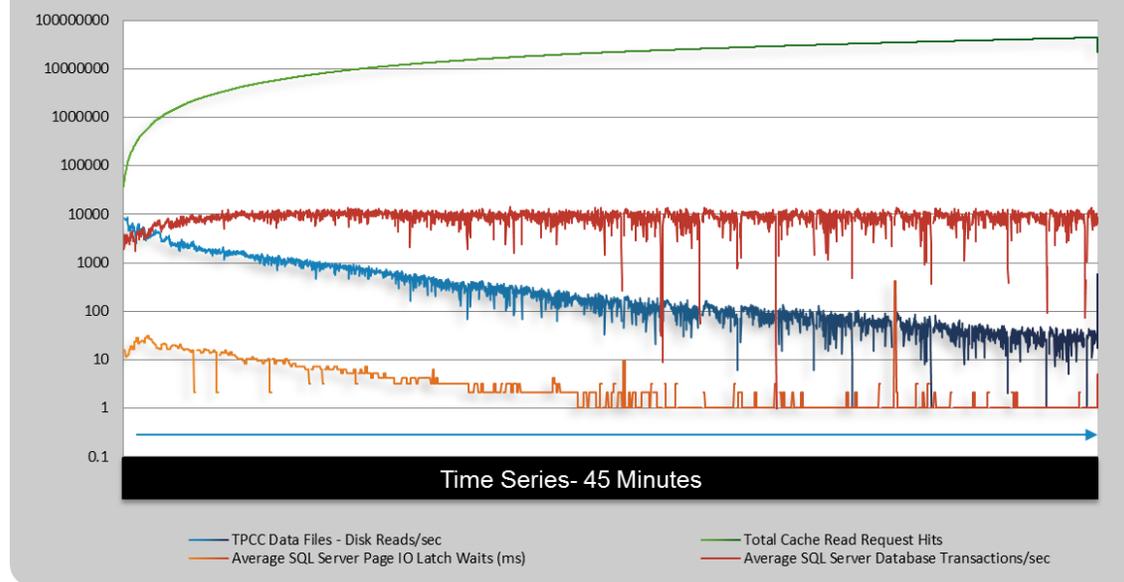


Figure 14: As the cache warms, performance increases and latency decreases.



FUSION-io®

Interpreting the Results

The data proves directCache increased the overall performance of the EqualLogic array by offloading read request onto Fusion's ioDrive2.

Here are the key takeaways from the testing:

- As the cache warms, disk and SQL Server latency decreases and transactions per second increase. Write performance on the EqualLogic array increased as a byproduct.
- The FIO benchmark tests show an 86% improvement in write latency and a 22,000% improvement in read latency. Read bandwidth increased 550% and write bandwidth increased 252% for mixed workloads.
- TPC-C tests show SQL Server transactions per second increased 300% and overall EqualLogic write performance increased 189%.
- The cumulative data movement from the TPC-C database increased 230% and the total number of IO queuing was reduced by 61%.
- The SQL Server latency charts suggest a net increase in write latency when directCache is enabled. However, the write latency increase is due to 284% improvement in write throughput per second and a 189% improvement in write request per second.

Summary

Fusion's directCache provides the perfect solution for those who want to take advantage of the features of an EqualLogic storage array and the performance of Fusion's ioDrive2 flash memory. By combining the two technologies, enterprise applications can achieve the industry's highest performance from Fusion-io as well as maintain all back up, high availability, and de-duplication strategies put in place by an EqualLogic storage array. The two storage devices interact seamlessly to easily deploy new application environments and an uninterrupted experience for environments previously deployed. These test results demonstrate the tremendous impact that Fusion-io and Dell storage can provide.



FUSION-IO®

APPENDIX

Our hardware configuration is as follows:

Server	Storage	Network
		
PowerEdge R720 <ul style="list-style-type: none">• 1 x 785GB Fusion-io ioDrive2.• Intel® Xeon® E5-2650 @ 2.00GHz. 2 processors with 8 cores each.• 64GB RAM• Broadcom BCM57800 NetXtremell Dual Port 10Gb NIC• Fusion directCache (ver. 1.2) and Fusion VSL 3.2.3• Microsoft Windows Server 2012 Data Center Edition• Microsoft SQL Server 2012 with SP1	EqualLogic PS6010XV Array <ul style="list-style-type: none">• 16 x 300GB 15K RPM Disks• Dual Storage Controllers (Active/Passive)• 4 x 10Gbe for iSCSi connectivity	Force10 S4810 Switch <ul style="list-style-type: none">• 48 x 10Gbe ports

TOOLS

To capture performance metrics we used a combination of tools including; SQL Server Profiler, Performance Monitor, several custom PowerShell and SQL scripts. The SQL scripts were used to capture results from several Dynamic Management Views (DMV) queries. The DMVs provide insights into wait statistics from SQL Server. The team used `sys.dm_virtual_file_stats` to capture I/O metrics from SQL Server. The PowerShell scripts were used to execute the TPC-C SQL scripts. The team leveraged FIO (Flexible IO) tester to generate various workloads to simulate disk activity on the EqualLogic array.



FUSION-IO®