

BOSTON

Servers · Storage · Solutions



LSI MegaRAID CacheCade Pro 2.0 Software Evaluation

By Boston Limited

Table of Contents

ABOUT BOSTON LIMITED	3
INTRODUCTION	4
TEST CONFIGURATION.....	5
TESTING METHODOLOGY	5
BENCHMARK RESULTS.....	6
CONCLUSIONS	8

© 2012 Boston Limited. Errors & Omissions Excepted.
All brands, names and trademarks remain the property of their respective owners.

About Boston Limited

With over 20 years of trading within the distribution and OEM marketplace, Boston Limited continues to lead the way in providing the latest high performance power optimised technologies into the HPC, ISP, military, VFX, enterprise and broadcast markets with our multi-award winning server, storage, workstation and clustered solutions.

Our experience in the assembly, testing and validation of high-performance bespoke solutions combined with our ability to provide client specific branding, documentation, packaging and global on-site maintenance packages, allows us to provide a genuine range of value added services to our partners.

Headquartered in the UK, Boston has offices in the USA, Germany and India to support its growing global operations. For more information, visit our website at www.boston.co.uk

Introduction

LSI® MegaRAID® CacheCade® Pro 2.0 software promises to be the answer to IT managers' serious pain points with storage, those related to random I/O performance. This new functionality is supported on several LSI SAS 6Gb/s RAID controller lines and adds SSD caching functionality to the already impressive feature set to help improve performance and negate some of the pitfalls of traditional hard disk technology.

Today it is possible to purchase magnetic hard disks with capacities of up to 3TB for under £300 but they struggle to provide 200 IOPS, while some SSDs easily achieve 50,000 IOPS but cost 5-10 times that for only a few hundred gigabytes of capacity.

LSI CacheCade Pro 2.0 software and other similar hybrid storage pools enable IT managers to take the main benefit of magnetic disks, namely the £ /GB, and marry it with the main benefit of SSD, IOPS performance. It's a no brainer.

Implementation of the new cache feature coined CacheCade™ Pro 2.0 software is very simple. The RAID controller even continues to operate with the same functionality, drivers and toolset as before. The latest firmware update is required to enable the CacheCade Pro 2.0 software so existing users will need to upgrade firmware in advance to utilize the feature. Customers must purchase the CacheCade Pro 2.0 software license and apply it to their controller or purchase a controller with this preinstalled. If the key is to be purchased separately it is available as a software key or as a physical hardware key.

Each controller can then be equipped with up to 512GB of standard S-ATA or SAS SSDs from LSI's validated selection and assign it as a virtual caching drive; this can be as simple as a single SSD or up to as many as 32 SSDs in a RAID volume. Whilst RAID protection or performance may not be required for everyone, for critical applications at least two SSDs are likely to be used in a RAID 1 mirror alongside a standard BBU to protect the controller cache.

New or existing RAID arrays utilizing standard magnetic/spinning hard disks can be enabled with CacheCade Pro 2.0 software by associating these RAID volumes with the new SSD accelerator volume. From here the array data is automatically cached using this volume for reads and optionally for write operations.

The controllers officially supporting CacheCade Pro 2.0 software are as follows

- LSI MegaRAID 9260 Series
- LSI MegaRAID 9265 Series
- LSI MegaRAID 9280 Series
- LSI MegaRAID 9285 Series

Additionally the Supermicro range of AOC-U/SAS2H8 controllers can be enabled for CacheCade Pro 2.0 with the appropriate soft or hardware key.

Test Configuration

- Supermicro 846E16 Storage Chassis
- Supermicro X8DTH-6 Motherboard
- Dual Xeon X5670 CPUs
- 8GB DDR-3 Registered ECC Memory
- LSI 2108 RAID Controller - Firmware 12.13.0-0104 - Driver 5.1.112.64 - CacheCade Pro 2.0 software Enabled
- 4 x Western Digital WD2002FYPS 2TB Green Power S-ATA HDDs
- 1 x Toshiba MK1001GRZB (0106) SSD for CacheCade Pro 2.0
- 2 x Toshiba MK1001GRZB (0106) SSDs for CacheCade Pro 2.0 in RAID 0 for CacheCade Pro 2.0 Stripe
- Windows 2008 R2 x64 Standard
- IOMeter 2006.7.27 Tests performed with 8 Workers, 8 outstanding I/O's per worker

Testing Methodology

For simplicity we will focus on the four corners of disk performance testing although extensive testing has also been concluded. The four cases are:

4K Random Read Simulating random small block I/O similar to database / OLTP applications with a significant read bias.

4K Random Write Simulating random small block I/O similar to database / OLTP applications with a significant write bias.


1MB Sequential Read Simulating high throughput applications with a read bias such as video playback.

1MB Sequential Write Simulating high throughput applications with a write bias such as data or video capture.

A RAID 5 using 4x WD2002FYPS drives was created and tests were conducted in 4 different caching scenarios to help understand the benefits of CacheCade Pro 2.0 software.

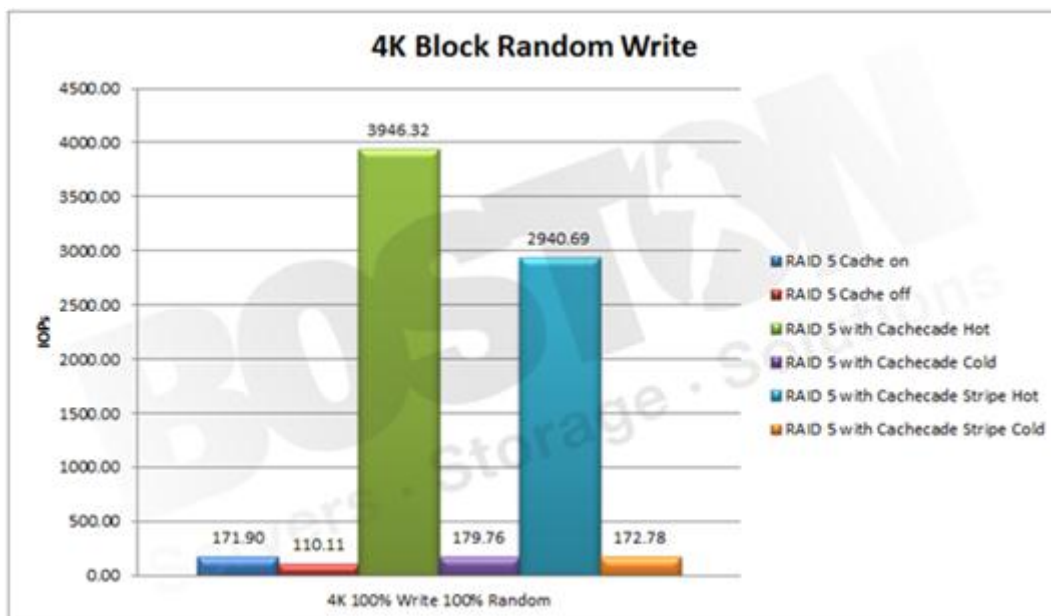
1. Cache on - The controller's DRAM cache, read ahead, IO cache and disk cache were enabled.
2. Cache off - The controller's DRAM cache, read ahead, IO cache and disk cache were disabled.
3. CacheCade Pro 2.0 Software - Hot - The controller's DRAM cache, IO cache and disk cache were enabled, read ahead disabled and CacheCade Pro 2.0 software was enabled in Write Back mode (read and write). Additionally IOMeter was instructed to only use the first 62 million sectors of the RAID 5 volume to limit the target areas to around 30GB to simulate a hot spot of activity. 3 benchmarks were run to fill the cache in advance and results were collected on the 4th. This is simulating a working dataset which is small enough to reside permanently in the SSD cache.
4. CacheCade Pro 2.0 Software - Cold - The controller's DRAM cache, IO cache and disk cache were enabled, read ahead disabled and CacheCade Pro 2.0 software was enabled in Write Back mode (read and write). IOMeter was instructed to use the whole disk and spread the transfers across any of the sectors in the whole 5.45TB RAID 5 volume - simulating a dataset much greater than the SSD cache. 3 benchmarks were run in an attempt to fill the cache in advance and results were collected on the 4th.

Benchmark Results

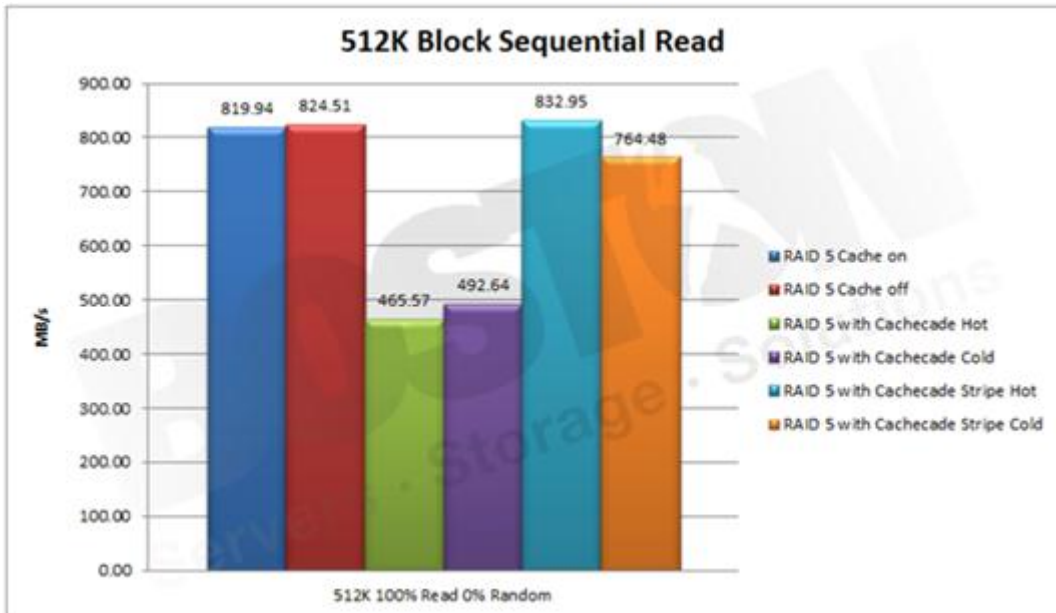


The benefits of CacheCade Pro 2.0 software are clear to see on this first test. Random reads are a very difficult operation as you cannot predict exactly which data you will need next. Read ahead algorithms can predict data which has a pattern or is sequential but not completely random access. Where the dataset is already hot and in the cache we can see the benefits of SSD straight away, more than 150 times the IOPs of a non CacheCade Pro 2.0 software volume.

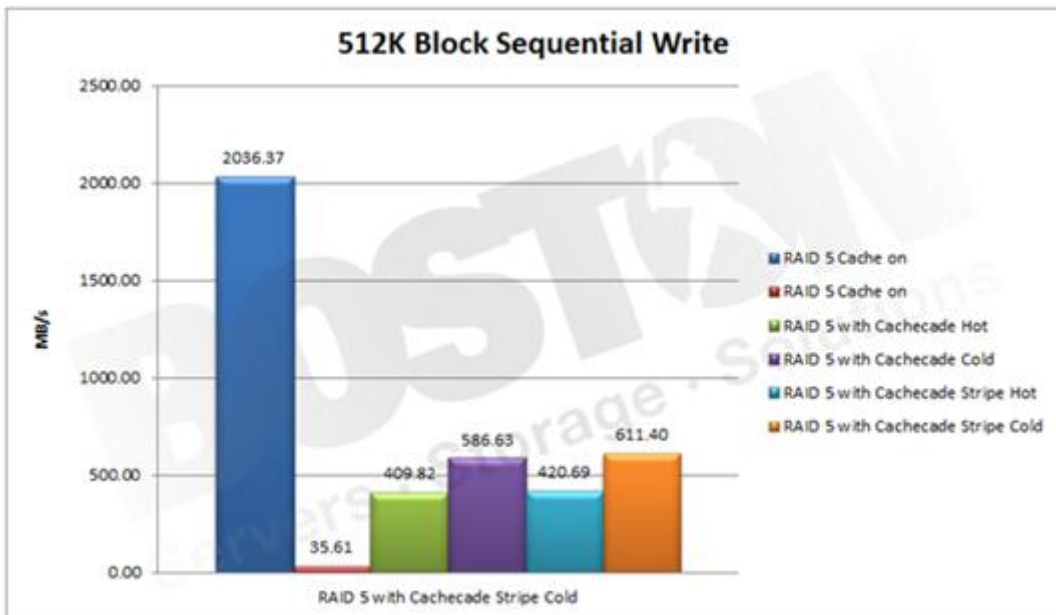
Interestingly, where the dataset is much larger than the SSD cache there is little or even no improvement in performance as the number of cache hits is low due to the unpredictable access pattern.



Again, CacheCade Pro 2.0 software shows impressive performance benefits over standard disk arrays for hot cache data which shows an astonishing more than 20 times improvement in performance due to the high IOPs and low latency provided by the SSD cache. What is achieved is a significant gain for a relatively modest investment.



Interestingly the read throughput tests show a different story when using a single SSD - CacheCade Pro 2.0 software appears to hinder these types of workloads. This is simply explained, we have moved our data from several hard disks which are actually suited well to sequential workloads onto a single SSD - as a result the performance is now somewhat reliant on the characteristics of that SSD. One SSD's throughput is lower than that of 4 magnetic disks working together. One way to improve the performance and gain the best of both worlds is to add more SSD's and create a RAID 0, or better still a RAID 10.



Our last test really shows the benefits of having a write cache, it's well known that with no write caching at all write performance is painfully low as every write is committed to disk and confirmed before the next transaction starts. The standard RAID controller and disk caches give excellent an excellent boost to performance uncached array - we can see in a synthetic test how the DRAM controller cache can handle more than 2GB/s of throughput a significant figure. CacheCade Pro 2.0 software is able to show a benefit over an uncached array however again the performance is slower than without CacheCade Pro 2.0 software for these purely sequential operations. Notably performance is better with a cold cache, possibly as more of the cache is available for write operations.

Conclusions

LSI CacheCade Pro 2.0 software is not suited to every workload; however in those cases where gains can be had the benefits are significant. Our benchmarks are synthetic and won't exactly replicate a real world environment, but they can help us to generalize and make assumptions based on different access patterns.

If random reads or writes in specific hot spots of data are a pain point for your application then CacheCade Pro 2.0 software is a simple addition to your server which can reap performance gains of 20 to 150 times. All this is available for what is essentially a modest investment which consumes little power and requires no specialist knowledge to deploy.

On the other hand, if your workload is biased towards sequential disk access you may find that SSD caching does not enhance performance. In those cases, striping several SSD's together to achieve better throughput could help achieve the benefits of both worlds.

In reality, most workloads have a mixture of both random and sequential access which means that there is likely a benefit to be had for most applications.

For the price tag the level of performance delivered by CacheCade Pro 2.0 makes it a simple choice for the IT manager dealing with I/O hotspots.

High Availability Storage

High Performance Computing

Turn-key Cloud Solutions

Award Winning Products

Enterprise Class Liquid Cooling



UK Headquarters

Tel: +44 (0) 1727 876 100

E-mail: sales@boston.co.uk

www.boston.co.uk



India

Tel: +91 (0) 22 39530897

E-mail: sales@bostonindia.in

www.bostonindia.in



Germany

Tel: +49 (0) 89 9090 199-3

E-mail: info@bostondeutschland.de

www.bostondeutschland.de