White Paper

MySQL Server Optimization with LSI™ MegaRAID® Flash Technologies and Solid State Disk (SSD) Technology

Transactional databases can achieve significantly higher performance than is possible with hard drive-only configurations by coupling MegaRAID CacheCade software or MegaRAID FastPath software with today's SSD technology. These premium features are designed to provide higher performance plus lower cost and reduced power per database transaction.[1]

Cloud data center workload demands are increasing at exponential rates due to the explosive growth of transactional web-based applications — also referred to as Web 2.0. These demands are quickly overwhelming traditional hard disk drive-based storage subsystems and their ability to effectively scale to support the growing demands from business and data centers.

The ability to effectively scale storage performance to meet these new demands by utilizing SSD technology is critical to building a solid solution. This paper will illustrate the potential performance benefits of using MegaRAID CacheCade and MegaRAID FastPath software in MySQL server environments. The MySQL performance benefits discussed in this document are also applicable to other similar database applications.

CacheCade

CacheCade Software enables SSDs to be configured as a secondary tier of high performance cache to maximize transactional I/O performance. The solution is designed to accelerate the I/O performance of HDD-based arrays while minimizing investments in SSD technology. Further, CacheCade software reduces the significant investment in additional hard drives required to attain comparable performance gains as can be achieved by implementing the software.

CacheCade software relies on its dynamic ability to accelerate accesses to data hot spots. The highest achievable performance benefit depends upon the size and performance of the SSD relative to the logical volume being cached, and more specifically the total size of hot data compared to the total SSD capacity.

[1] In the context of a MySQL database and the SysBench benchmark, a transaction is a sequence of database commands starting with the "Begin" command and ending with the "Commit" command. The SysBench Advanced Transactional test mode is used to perform database operations on a MySQL "InnoDB" database supporting transactions. Because of the command line options used, each SysBench transaction consists of the following database commands: point queries, range queries, range SUM queries, range ORDERED BY queries, range DISTINCT queries, UPDATEs on index column, UPDATEs on non-index column, DELETE queries and INSERT queries. Each SysBench transaction is constructed so that the total content of the database does not increase or decrease, in other words the same amount of information deleted is subsequently inserted.

Additionally, this solution allows system builders to install just enough SSD capacity for active data and apply it to one or more HDD storage volumes.

By running a short trace of IO activity on the MySQL database (see Figure 1 below), we can determine that the majority of the database reads and writes are focused in specific ranges of logical block addresses (LBA). CacheCade software's intelligent caching algorithms detect and maintain frequently accessed hot read data in the SSDs, drastically reducing data access response times.

Because the majority of hot data read accesses are serviced by the SSDs, the HDD storage volume performance is also enhanced as a significant portion of the HDD IO workload is off-loaded to the SSDs.

Figure 1: XPERF trace of MySQL database IO activity

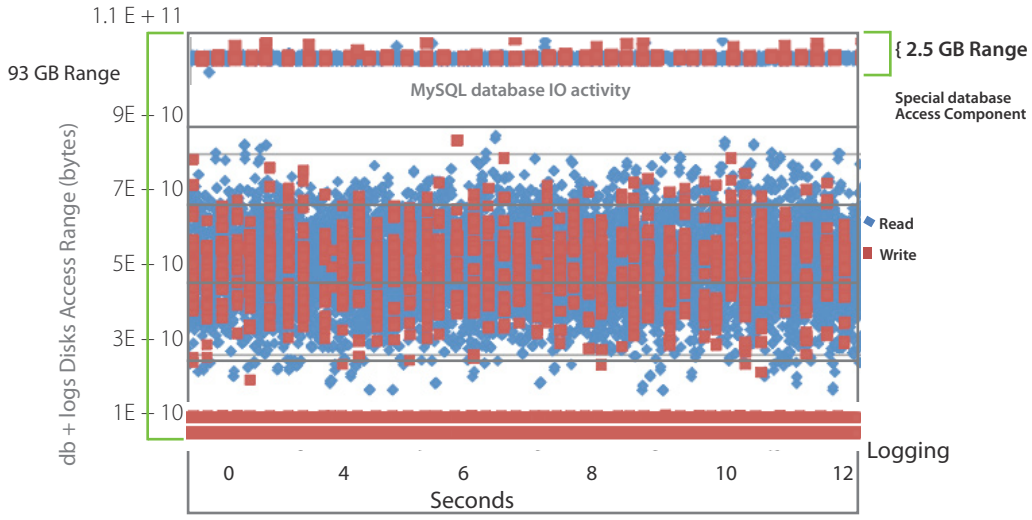

FastPath® Software Summary

FastPath software is a high performance IO accelerator for Solid State Drive (SSD) arrays configured as a virtual disk (or volume) behind a MegaRAID controller. It's a completely re-designed MegaRAID data path, optimized to meet the increasing performance characteristics of SSD technologies. FastPath technology can dramatically boost storage subsystem bandwidth – nearly doubling IOPs in certain IO profiles — providing an application platform that narrows the gap between CPU horsepower and storage horsepower. This means getting more return out of your server purchases. The revolutionary FastPath software's enabled firmware tuned to utilize the extremely low IO latency of SSDs, combined with the acceleration and data protection capabilities of a MegaRAID controller, offers industry leading data protected storage.

Recommendations for CacheCade and FastPath software

In order to optimize server storage performance in various application environments, IT administrators are forced to choose between costly upgrades to host memory capacity (see figure 2), or increasing the number of HDDs to the disk group which is costly and requires additional drive enclosure space and power. CacheCade and FastPath software offer alternatives to this trade-off and are designed to provide increased performance without the additional cost of more host memory or HDDs.

Figure 2: CacheCade software and 64GB SSD upgrade vs. increasing host memory cost comparison

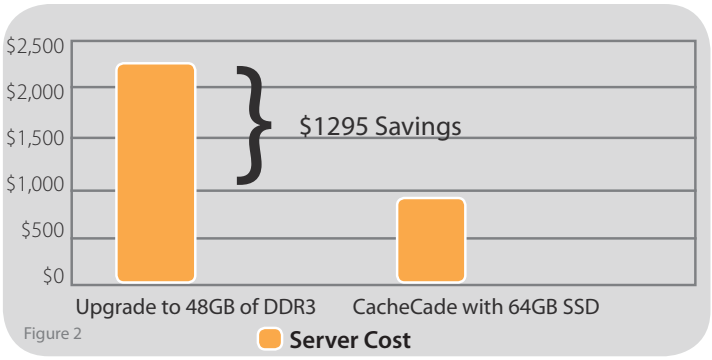

In the above example the base server is configured with a single CPU and 8GB (4 x 2GB) DDR3 memory. The server platform is limited to 6 DDR3 DIMMs per processor, therefore, the main memory upgrade requires substituting 4 x 2GB (8GB) DDR3 with 6 x 8GB (48GB) DDR3. This example also assumes an active working data set or data hot spot of approximately 40 GB. Based on this cost analysis, CacheCade software provides a more cost effective way to accelerate application performance. Individual results may vary depending on the configuration and environment.

Based on certain environmental and application attributes as well as business case requirements, the table below outlines where opting for CacheCade or FastPath software or adding HDD capacity may make the most sense, depending on the configuration. The following table is essential to understanding the advantages of CacheCade and FastPath software storage alternatives compared to short stroked HDDs.

Table 1: Performance optimization solution positioning estimates

Performance Enhancement Solution	SSD Capacity Needed	Existing HDD RAID DATA	Volume(s) Accelerated	Targeted Application Workload	Perfomance Benefits	Cost of Capacity	Cost of Ownership (4)
Add more short stroke HDDs	None	Redistribute on HDDs	One HDD Volume	All Workloads	Low (2)	Low to High (3)	High
CacheCade Software	Hot Data Only(1)	No Charge Required	One or More HDD Volumes	Random Read Intensive	Moderate	Moderate	Moderate
Fast Path Software	All Volume Data	Load or Migrate to SSD	One SSD Volume	All Workloads, Especially Ran-dom	High	High	Low

Table 1 Notes:

1. CacheCade software max. SSD capacity 512GB
2. Depends on the class and number of HDDs added and amount of short stroking
3. Depends on the amount of usable HDD storage when short stroking
4. Power, space, cooling, reliability, etc.

(Individual results may vary depending on the configuration and environment)

MySQL Benchmark and Testing Details

In order to simulate a typical online transactional processing environment, an actual MySQL Server database was loaded on to a server.

Baseline Configuration:
- (16) 15K RPM SAS HDDs

CacheCade software configuration:
- (12) 15K RPM SAS HDDs
- (4) Intel X25-E SLC SSDs

FastPath software configuration:
- (4) Intel X25E SLC SSDs in RAID 0
- (8) Intel X25E SLC SSDs in RAID 5 & RAID10

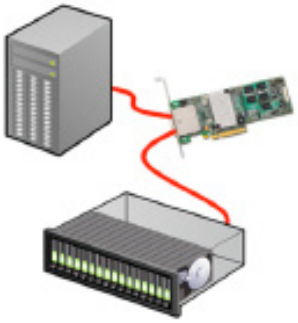

To achieve more optimal write performance based on SSD characteristics and the additional data redundancy operations required by RAID 10 and RAID 5, a larger set of SSDs were chosen for the FastPath software RAID 5 and RAID 10 tests. See table 2 for test configuration details.

A 93 GB database (400,000,000 rows) was loaded on these configurations along with two 190MB log files. The I/O characteristics of the database query requests are 74% reads and 26% writes – with an average transaction size of around 16KB.

Table 2: Test Configuration Details

Server	**1 Intel Xeon L5520** • 4-cores each, 2.27GHz • MySQL Server SysBench **6GB PC3-8500 ECC Regis tered Memory**
Storage	**MySQL Database and Log Files on MegaRAID RAID volume** • RAID 0, 10 and 5 volumes • Seagate 15K.2 Savvio 6G SAS (model ST9146852SS) • Intel X25-E 3G SATA (model SSDSA2SH032G1GN) **MegaRAID 9280-8E Controller** • Standard HDD Volume – 2.70.03-0862 (GCA) - Standard HDD VD: WB,DIO,ARA,DWCE,64KB Strip Size • CacheCade software - 2.60.03-800 -High performance SSD Caching and Tiering -CacheCade software: WB,CIO,NRA,DWCE,64KB strip size • FastPath software – 2.70.03-0862 (GCA) -High performance SSD volume -WT, DIO, NRA, DWCE, 64KB strip size • LSI CTS2600-24 2U, 6Gb/s SAS drive enclosure
Software	**Win2K8-SP2 32-bit, MySQL Server v5.1.45** **Industry Standard Transactional Benchmark** • SysBench v0.4.12 32-bit for Windows • OLTP Advanced Transactional Execution • "innodb" MySQL database engine used

LSI Test Results

In order to verify consistency of results, two separate benchmark runs were conducted for each configuration at each RAID level. Significant performance scaling in transactions per second is evident with CacheCade and FastPath software in all three RAID level configurations tested. See figures 3, 4, and 5.

The CacheCade software performance improvement gains range from 1.5x to almost 2x that of the hard drive only configuration, while FastPath software exhibited a 2x to 2.7x transactional performance improvement during the testing.

Figure 3: MySQL SysBench RAID 0 transactions per second

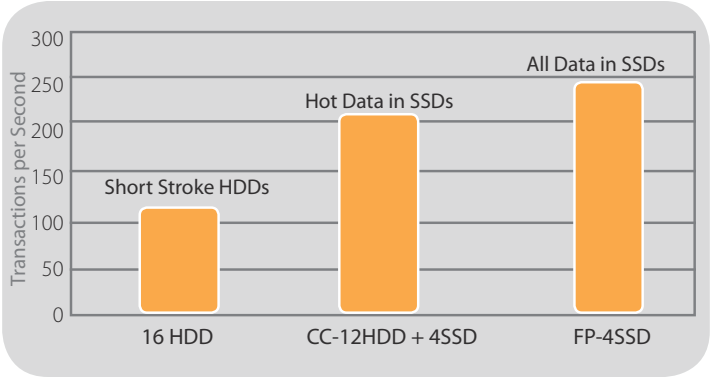

Figure 4: MySQL SysBench RAID 10 transactions per second

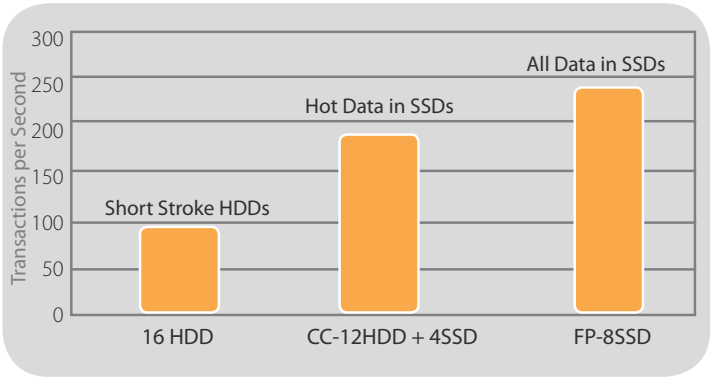

Figure 5: MySQL SysBench RAID 5 transactions per second

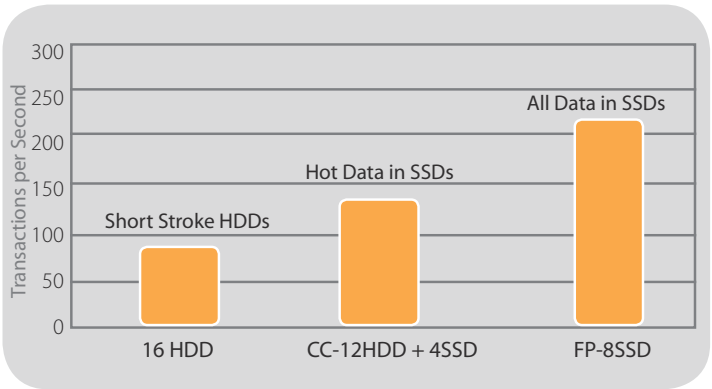

Reduced cost of ownership by decreasing the cost per transaction

The increased performance advantage offered by Flash-based storage, coupled with MegaRAID advanced software options, can reduce the total cost of ownership (TCO). Lower power utilization by deploying SSDs translates into operating cost savings.

Table 3 below represents the base costs for the three test configurations. For easy comparison, the on-going cost for running the system, or OPEX, costs are equal in this table. However, the storage capacities are not identical. HDD-based storage requires much higher device counts to achieve the same level of performance.

The disadvantage of HDDs is the time delay that occurs each time the read/write heads are required to move, and becomes most prevalent for random read and write operations. This can be minimized to some extent by a technique called 'short stroking' or using only a fraction (typically 10% to 20%) of the HDDs' available capacity on the outer sectors of the HDD platters.

When this is done the cost per GB for SSDs is not that different from HDDs. For example, a typical 147GB enterprise HDD costs $280. When short stroked, its cost per GB increases to $9.50 [$280/(20%*147GB) = $9.50/GB]; compared to a 60GB enterprise SSD at $11.40 per GB ($686/60GB = $11.40/GB).

Table 3: System base costs and performance comparison

DESCRIPTION	MSRP	16HDD	12HDD+4SSD & CACHECADE SOFTWARE	4SSD & FASTPATH SOFTWARE	8SSD & FASTPATH SOFTWARE
CacheCade and FastPath software Key	$270	$0	$270	$270	$0
Intel X25-E	**$380**	**$0**	**$1,520**	**$1,520**	**$3,040**
Seagate 15K.2 147GB Savvio HDDs	$280	$4,480	$3,360	$0	$0
System Base Cost	**$15,000**	**$15,000**	**$15,000**	**$15,000**	**$15,000**
OPEX (4 Year Power & Cooling Estimate)	$6,200	$6,200	$6,200	$6,200	$6,200
Total		**$25,680***	**$26,350***	**$22,990**	**$24,240**
Average Performance Improvement		1.0x	1.7x (1)	2.0x	2.5x

In the majority of scenarios tested internally at LSI, CacheCade software can reduce total system cost because it can significantly enhance performance (see figure 7) and reduce response time (see figure 6) **(1) of not just one but multiple HDD volumes.**

CacheCade software demonstrates more efficient data transaction processing behavior resulting in reduced overall system IO request response times and improved transaction processing, allowing for more efficient server and application productivity. Based on the test results, the same amount of work can be accomplished with fewer servers and at much lower total cost of ownership.

Figure 6: Average IO request response time (in milliseconds) HDDs vs. CacheCade software

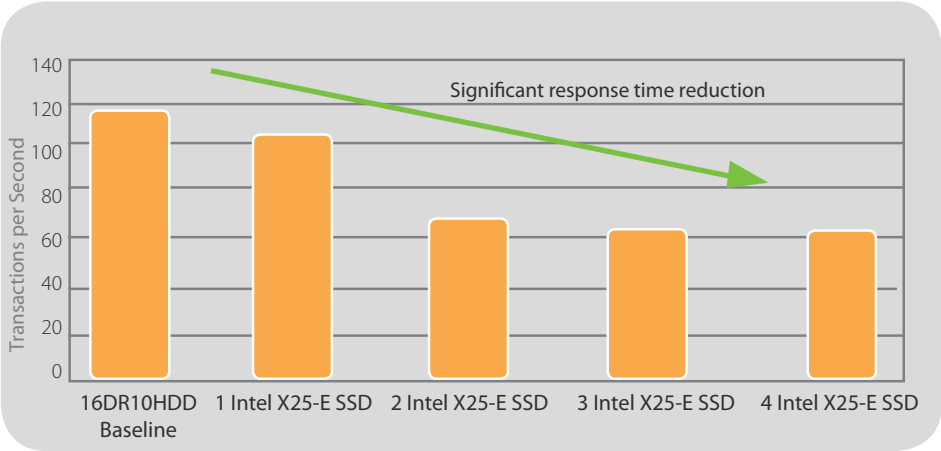

Figure 7: Transaction per Second HDDs vs. CacheCade Software

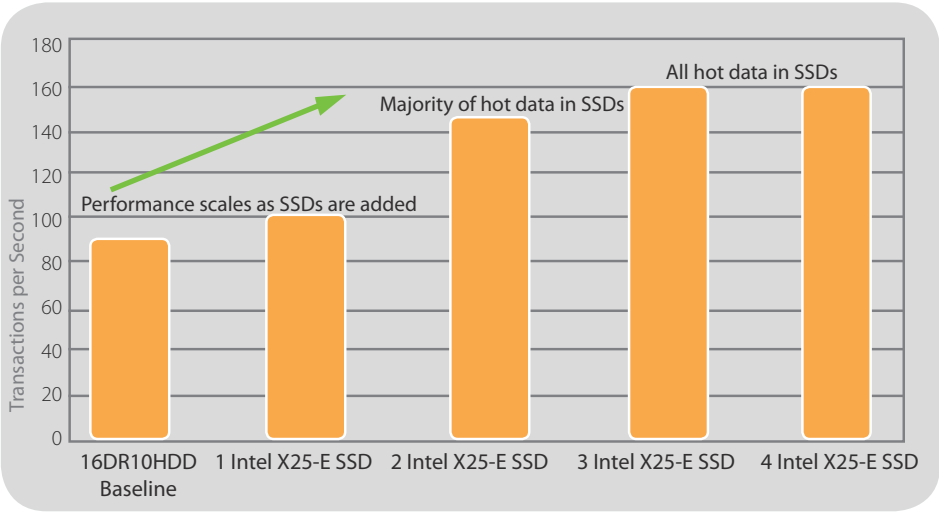

By taking the system cost totals in Table 2 and dividing by the average transaction time (transactions per second), the cost of server performance level per transaction can be determined. This is illustrated in Figure 8: Figure 8 compares RAID 0, RAID 10 and RAI D 5 cost per transaction for 16 HDDs only volume, CacheCade software 12 HDDs plus 4 SSDs volume and FastPath software 4 and 8 SDDs only volumes.

Figure 8: Solution cost per transaction

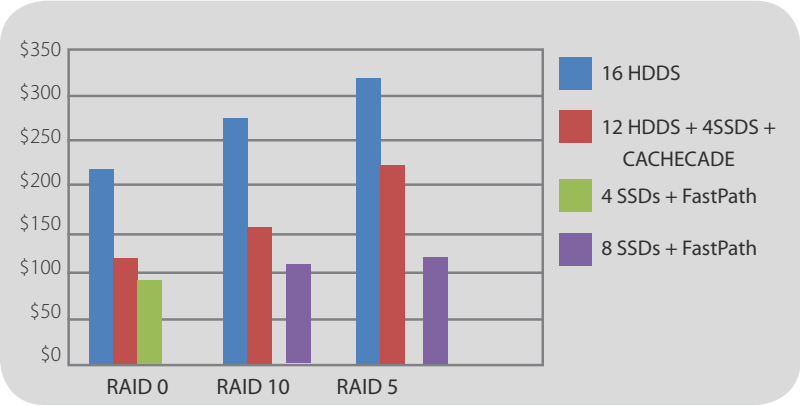

Conclusion

MegaRAID CacheCade software and MegaRAID FastPath software are designed to deliver performance scalability, reduced ownership costs and transactional efficiencies critical to meeting the demands of today's rapidly expanding transactional web 2.0 ecosystem.

CacheCade software provides a cost conscious option to integrate and realize SSD performance, reduce transactional cost, and improve overall application performance without extensive investment in SSDs. This is particularly evident in random read-intensive workloads where adding a few SSDs to an HDD array is feasible and maintaining existing HDD capacity, whereas adding additional short stroked HDDs is not. CacheCade provides an additional advantage in that it does not require any migration of user data off of the existing HDD volume.

MegaRAID FastPath software is a viable alternative to hard drive only disk arrays, particularly when overall read and write performance is the priority, and total storage capacity and cost are of lesser concern.

For more information and sales office locations, please visit the LSI web sites at: lsi.com/channel