LSI™ MegaRAID® CacheCade™ Performance Evaluation in a Web Server Environment

Evaluation report prepared under contract with LSI Corporation

Introduction

Interest in solid-state storage (SSS) is high, and IT professionals are beginning to explore various ways to take advantage of the increased performance and reduced power consumption that SSS offers. Various enterprise applications across the spectrum of size and complexity may be good candidates for SSS.

SSS can be deployed directly as primary storage where the users decide what data to place on the storage, or it can be deployed as a cache in front of traditional storage, where the storage controller determines what data is placed on the SSS or solid state disk (SSD). Each implementation has advantages.

LSI Corporation commissioned Demartek to evaluate MegaRAID CacheCade software in a web server environment. CacheCade software is optimized to enhance read cache performance with features such as hot spot detection and retention. As the operating system makes read requests to a CacheCade-managed volume, the logical block is written to the SSS and subsequent read requests for this cached data are serviced by the SSS. Hot spot detection and retention ensures that only frequently accessed data is placed into the SSD storage.

Evaluation Environment

These tests were conducted in the Demartek lab in Arvada, Colorado using Demartek servers, disk drives and networking infrastructure. LSI supplied the MegaRAID controller card and disk enclosure. The web server content was created by Demartek and run on Windows Server 2008 R2 with IIS 7.5.

Evaluation Summary

The MegaRAID CacheCade software provides outstanding performance improvements to the web server tested here with minimal configuration needed by an administrator. With one and two 32GB SSDs in the configuration, we observed more than 3x and 5x performance improvements, respectively, in overall throughput and even higher improvement in web server response time. These tests show that the MegaRAID CacheCade solution with SSDs can improve the performance of existing web servers or can provide the same performance with fewer web servers, lowering the total cost of ownership. The performance was significantly improved even when all of the web server content did not fit into the SSD cache.

1 – Test Environment

These tests were conducted using a relatively simple web server design that represents a web hosting server with many websites containing primarily static, read-only data. CacheCade software was designed to take advantage of web server workloads, among others.

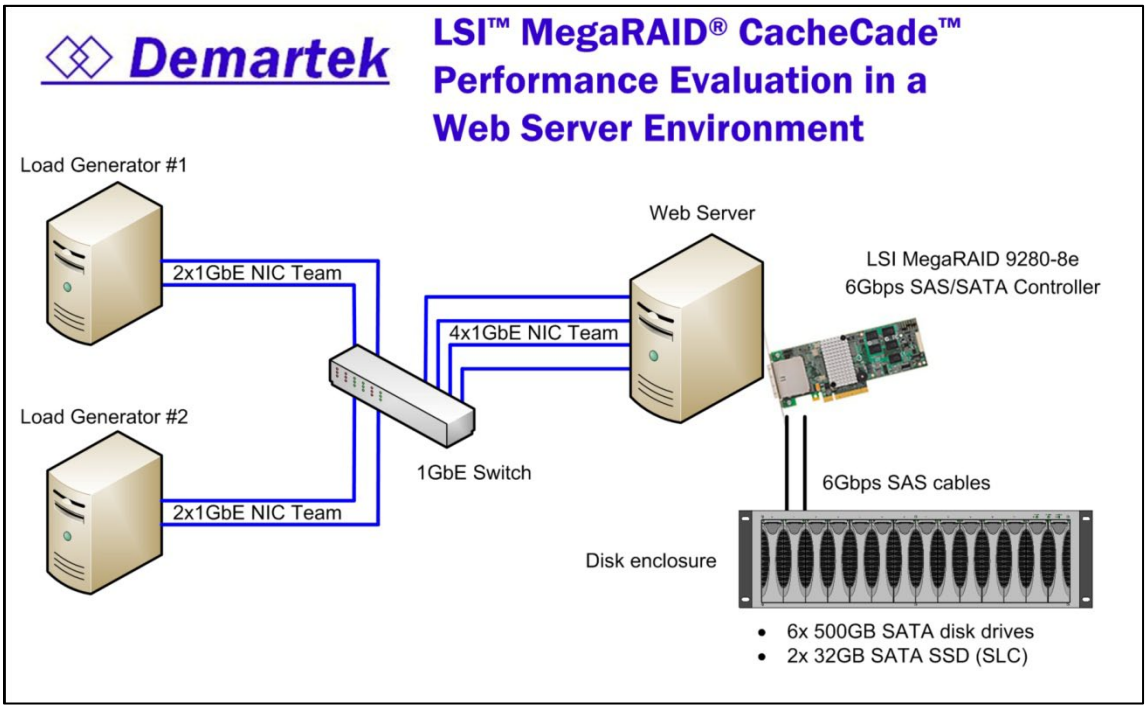

Web Server Configuration

- SuperMicro X8DTH-6F, PCIe 2.0
- Dual Intel Xeon E5540, 2.53GHz, 8 total cores, 16 logical processors
- 8GB RAM
- 2x Motherboard 1GbE NIC for general LAN traffic
- 2x Intel PRO/1000 PT dual-port 1GbE Server Adapters (4x1GbE NIC Team) dedicated to web server test traffic
- LSI MegaRAID 9280-8e 6Gbps SAS/SATA RAID controller
  - Firmware 2.100.03-0921
  - Driver: megasas2.sys, 4.30.0.64, July 15, 2010
- Disk enclosure containing
  - 6x Seagate Barracuda 7200.11 500GB SATA disk drives configured as RAID10. Two 100GB logical volumes, one for the operating system and one for the web server content.
  - 2x Intel X25-E 32GB SLC SSDs allocated to the LSI MegaRAID controller
- Windows Server 2008 R2 Enterprise Edition with IIS 7.5

Load Generator Configuration (x2)

- Dell PowerEdge 2900

- Dual Intel Xeon E5345, 2.33 GHz, 8 total cores
- 48GB RAM
- 2x Motherboard 1GbE NIC for general LAN traffic
- Intel PRO/1000 PT dual-port 1GbE Server Adapter (2x1GbE NIC Team) dedicated to web server test traffic
- Windows Server 2008 R2 Enterprise Edition
- NeoLoad 3.1 web load generator software

Network Switch

- Dell PowerConnect 2748, 48-port 1Gb Ethernet switch

Network Configuration

In order to ensure that the network was not the bottleneck, teamed, dual-port 1GbE NICs were installed into the load generators and into the web server to provide increased network bandwidth for the web server test traffic. The teamed NICs were configured with Static Link Aggregation. The web server test traffic was dedicated to the teamed NICs and any general LAN traffic was handled by the 1GbE motherboard NICs. The client load generators and the web server were connected to the same 1Gb Ethernet switch, and the appropriate switch ports were configured into Link Aggregation Groups (LAG). The network latency in this configuration was less than 1 ms between the load generators and the test web server.

Web Server Client Load Generator

A web server load tool, NeoLoad, from Neotys software, was used to generate the client load on the web server. The NeoLoad load generators were installed onto two servers that provided the client workloads. The NeoLoad console was installed onto one of the servers acting as a load generator.

Web Server Content

The web server had the following characteristics:

- Approximately 23GB of content on the operating system volume, including the paging file
- Approximately 40GB of content on the web server data volume
- Approximately 1.48 million files of web content data
  - 80,000 HTML text pages
  - Approximately 1.4 million graphic images (JPEG and PNG)

Web Server Activity

The website requests and responses had the following characteristics:

- Requests randomly referenced the entire 40GB and 1.48 million files approximately evenly over the entire duration of the 90-minute tests.
- Each request for an HTML page returned unique HTML text and eleven or twelve unique graphic images.
- Web server workload tuned to reproduce typical throughput.

2 – SSS Technology as a Cache

Deploying SSS as a cache is one of the ways to deploy SSD technology. This has the advantage of letting the caching system observe the data access patterns and determine what data to place on the SSS in order to get the most benefit. In most cases, the user or administrator does little or nothing to achieve the performance benefits. The caching system puts as much frequently accessed data into its SSD cache as possible, regardless of the capacity of spinning disks behind the cache or the amount of data stored on the spinning disks. The only real task for the administrator or user is to decide how much SSD technology to deploy and then to configure the caching system to use that amount of SSD.

In order to test this caching implementation, we deployed a web server with direct attached storage (DAS) using a MegaRAID controller with the CacheCade feature that could be enabled. The DAS storage included a combination of spinning disk drives and solid state drives. Using a web load generator application, we sent at least 100,000 requests to the web server over 90-minute periods of time, and took several measures of web page and web server performance. During the tests, we expected that many of the pages of web content would be accessed multiple times. We ran this set of tests with the following storage configurations on the web server:

- The operating system and web content on spinning disk drives and no SSDs
- The operating system and web content on spinning disk drives and one SSD assigned to CacheCade software
- The operating system and web content on spinning disk drives and two SSDs assigned to CacheCade software

The SSDs were each 32GB in raw storage capacity. With the web content requiring 40GB of storage capacity, one SSD would not be enough to cache all the web content data, but two SSDs would be enough to cache the entire web content.

We expected the caching implementation to gradually increase the web server performance over time, as more and more of the repeatedly accessed web server content data was loaded into the storage cache, based on the hot spot data detection and retention algorithms.

Demartek also evaluated deploying solid state storage as primary storage in a similar environment. These reports are part of a series of reports evaluating the LSI MegaRAID SSD management functions. These reports are available at the following locations.

- Demartek Publishes Evaluation of LSI MegaRAID CacheCade Performance
- Demartek Publishes Evaluation of LSI MegaRAID FastPath Performance

3 – Test Results

Our tests consisted of 90-minute runs with the client load generators making at least 100,000 web server requests, using different storage configurations. These client requests accessed the entire 40GB of web content randomly and approximately evenly. Many of the web pages were accessed multiple times during each test.

We disabled the default compression of data in IIS 7.5 so that as much data and I/O as possible would be pushed through the storage subsystem. We disabled and removed all "8dot3" short filenames on the web server to improve NTFS metadata processing.

Caching Everything is Better

We found that the best performance was obtained when both the operating system logical volume and the web content logical volume were under the control of CacheCade software. Although none of the web server content was located on the boot volume, there was enough operating system activity on the boot volume, including the paging file activity, that CacheCade software was able to provide a performance boost by caching the operating system volume in addition to caching the web server content volume. In the tests where we enabled caching only for the operating system volume or only for the web server content volume, the performance was lower than when we enabled caching for both volumes.

Caching Improved Over Time

As expected, performance increased over time as more and more of the data was accessed multiple times and placed into the cache. We chose 40GB of web server content so that only one 32GB SSD would not be able to cache the entire web content, but two 32GB SSDs would be able to hold the entire web content in addition to the approximately 23GB of operating system volume data, including the paging file and the IIS log files, that might be cached. While most of the web content volume was frequently accessed during the tests, only the frequently accessed portion of the operating system volume was placed into the SSD cache.

General Performance Overview

There were significant performance improvements with a single SSD above the performance obtained by the spinning hard drives only, and additional, significant improvements when the second SSD was added. When we added a third SSD to this configuration with this data, there was no significant improvement.

Total Pages

There were 80,000 unique HTML pages of web content. The client load generators accessed these pages randomly and evenly throughout the tests.

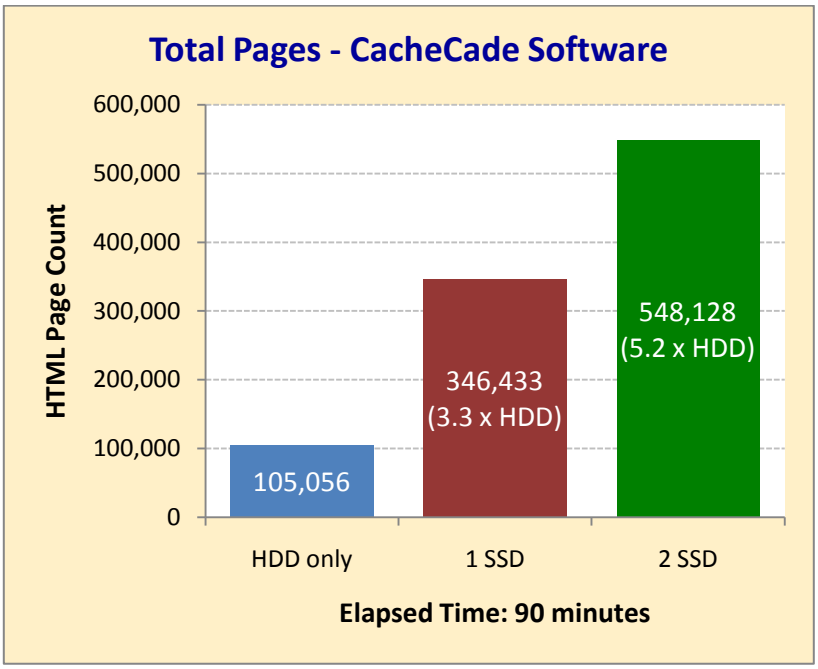

Total Hits

A web "hit" is an access of any file on the web server. For these tests, each HTML page request resulted in the access of one page of text and eleven or twelve graphic images, resulting in twelve or thirteen "hits" per request.

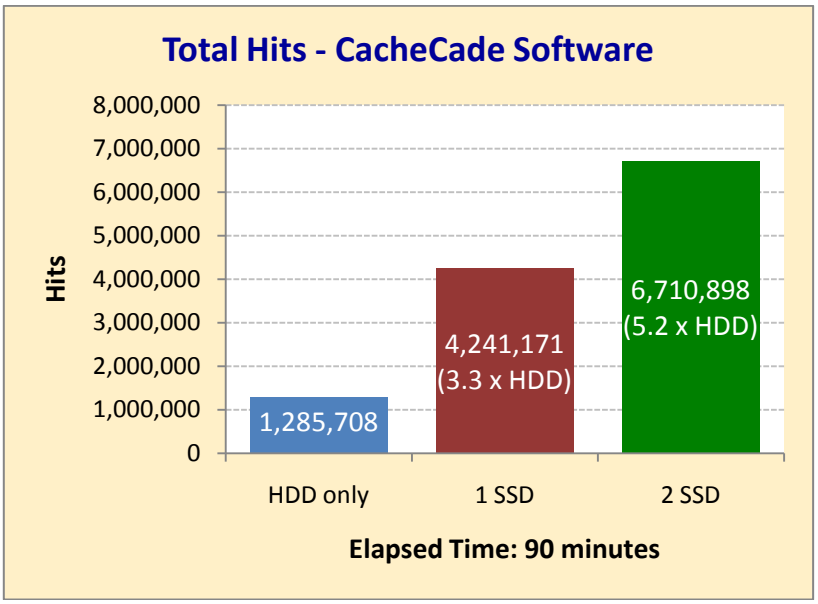

Total Throughput

Throughput is a measure of the total number of Megabytes transferred from the web server to the clients over the duration of the test.

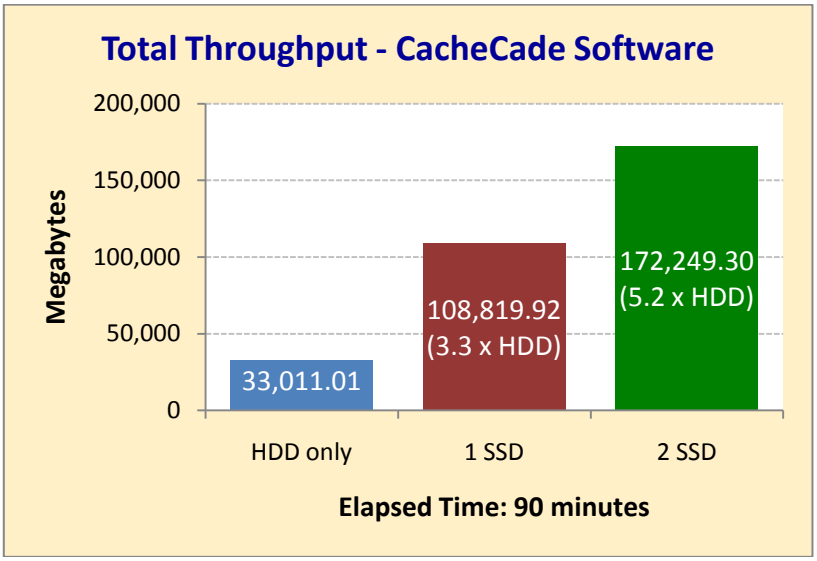

Performance Ramp-up

Before each test run, the web server and load generators were rebooted to insure that there was no data leftover in their system and storage caches. At the beginning of each test run, the system and storage caches were empty but began to accumulate data as the test made progress and data began to be accessed repeatedly.

The following charts show various measures of the performance characteristics over the duration of the test runs. These performance rates are measured at the clients and include the full end-to-end transaction.

- **Throughput** – Megabits per second of data transfer from the web server to the clients as a results of the client requests
- **Page Hit Rate** – Hits per second of the HTML web content files accessed by the clients. This total counts the HTML text page and all the graphics images for that page as one page hit.
- **Average Page Response Time** – This is the average time, in seconds, to return the complete HTML page from the web server to the client.
- **Average Time to First Byte** (TTFB) – This is the average time, in seconds, for the client to receive the first byte of the response from the web server.

For the hard-drive-only configuration, a small amount of data made its way into the RAID controller DRAM cache and system DRAM cache, but the hard-drive-only configuration achieved an almost flat level of performance after the first few minutes of the tests.

The 1-SSD configuration performance climbed steadily for at least the first hour and then achieved an almost flat performance level at the end of 90 minutes. The 2-SSD configuration performance level was increasing throughout the 90-minutes of the test, indicating that its cache was not full and that it had not yet achieved its maximum performance. The hot data detection algorithm utilized by CacheCade software is specifically designed to load only frequently accessed data into the SSD cache which improves application performance.

These tests show that adding SSDs with CacheCade software improves the website user experience and supports more concurrent users with fewer web servers or extends the performance of existing web servers, resulting in a reduced total cost of ownership.

Server Performance

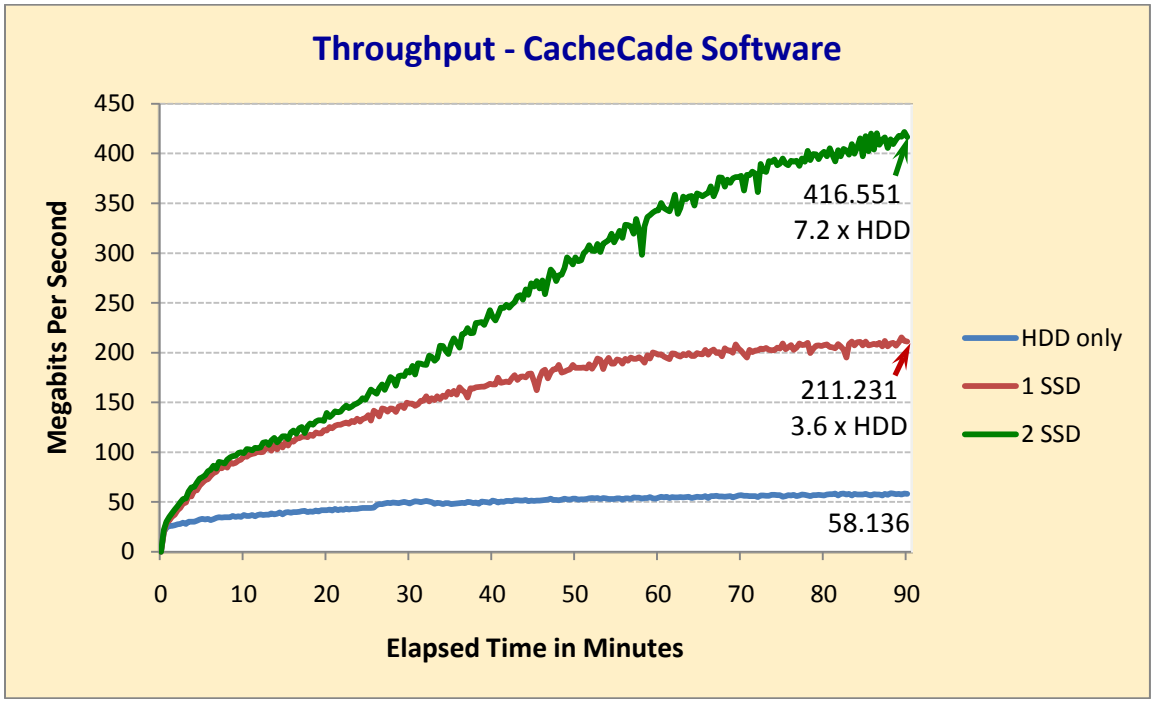

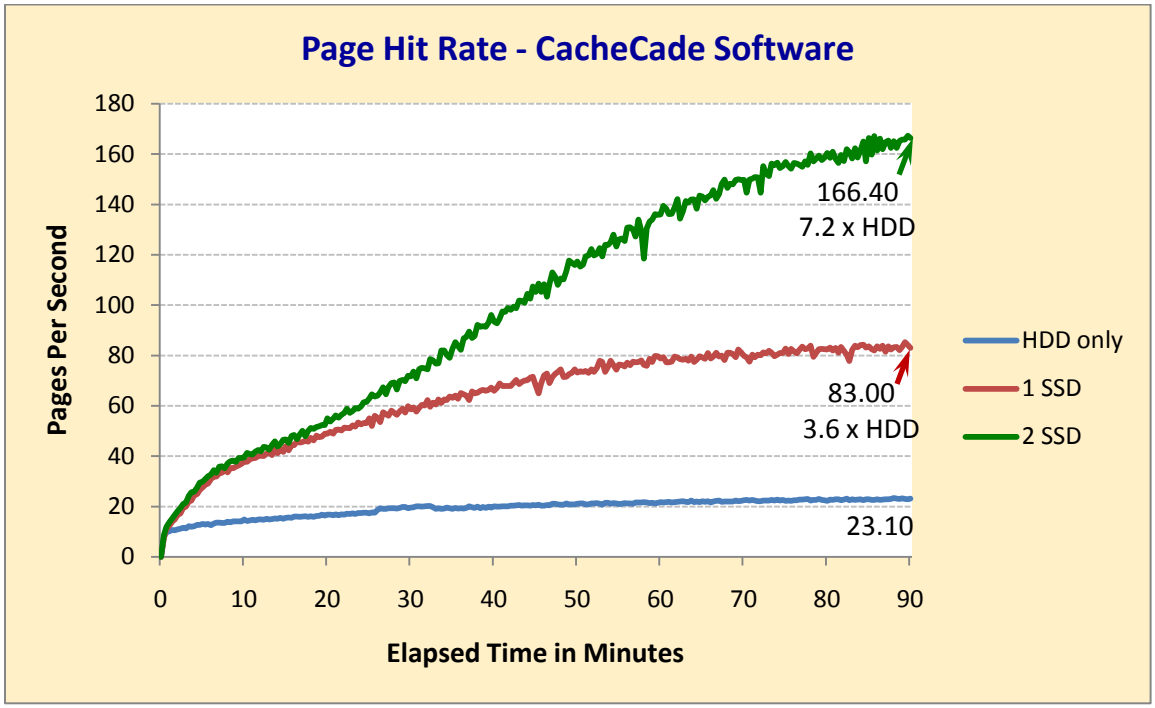

Server Responsiveness

Various studies have shown that even slight improvements in web site responsiveness improve the user's perception of the quality of the web site. CacheCade software with SSDs provide dramatic improvement in responsiveness by reducing the average response time to less than 0.1 seconds.

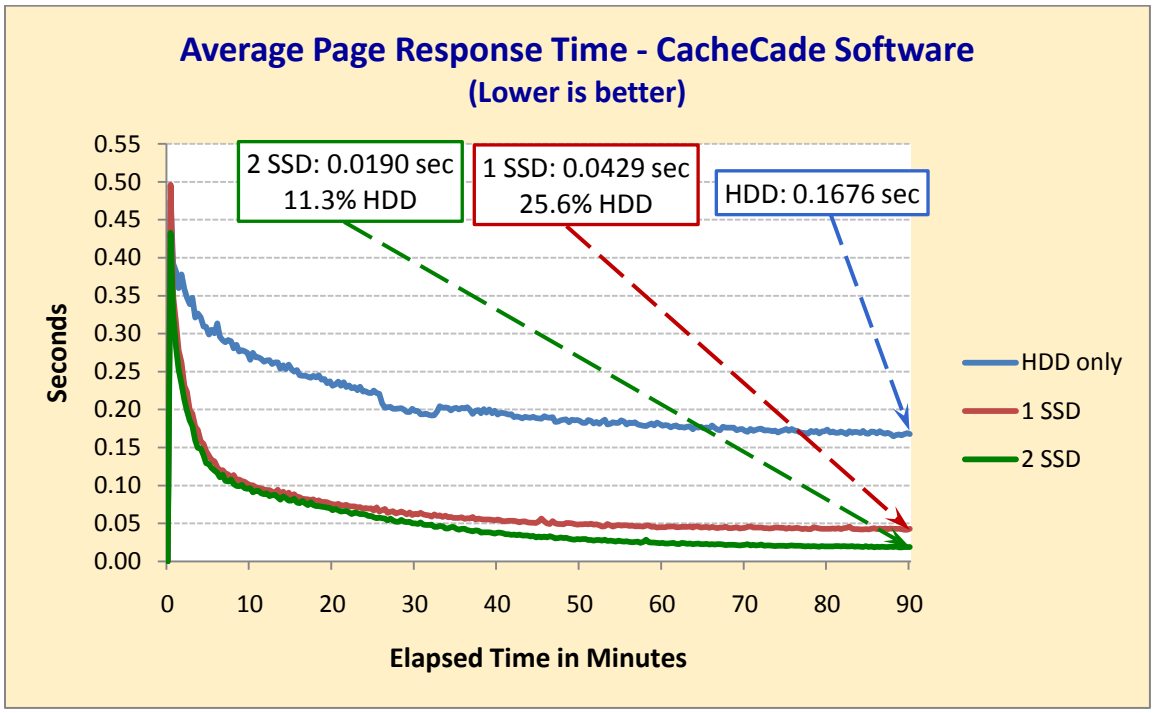

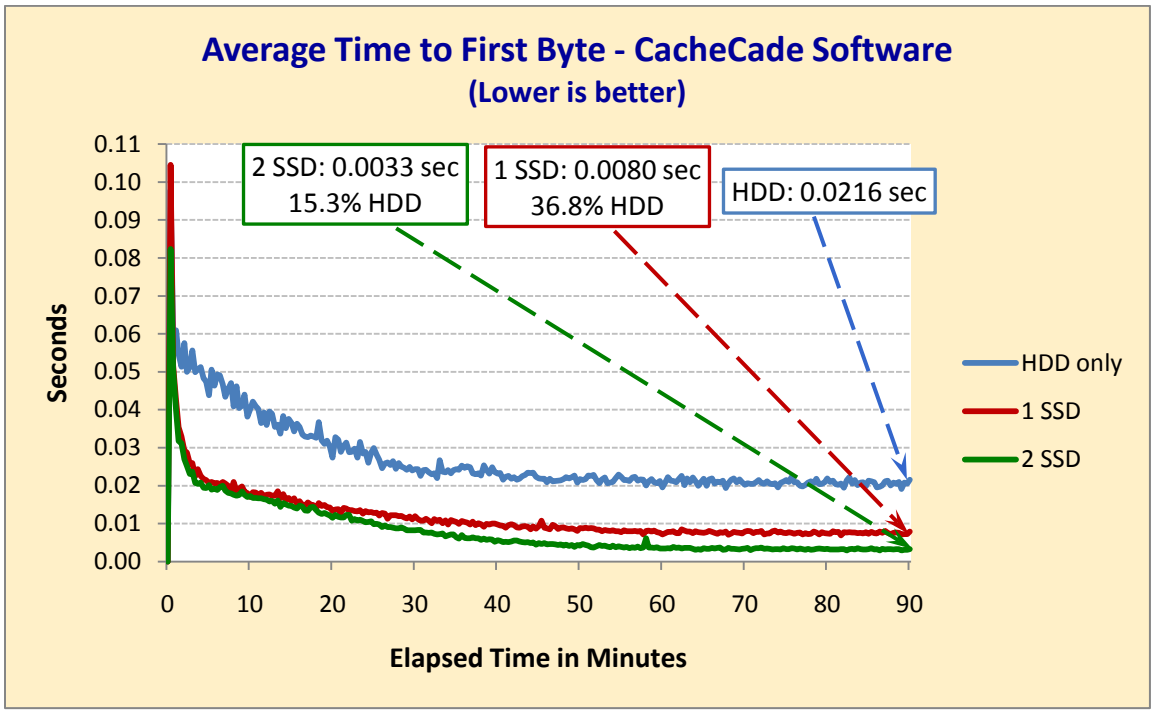

4 – Total Cost of Ownership

Total cost of ownership (TCO) calculations should be included when making technology acquisition decisions. TCO includes not only the initial capital expense (CapEx) but also includes ongoing operational expenses (OpEx).

CacheCade and CapEx

Examining the prices of DDR3 ECC DRAM and SLC SSDs, we can see the price difference of adding DDR3 ECC DRAM vs. adding the same capacity of SLC SSD to a server.

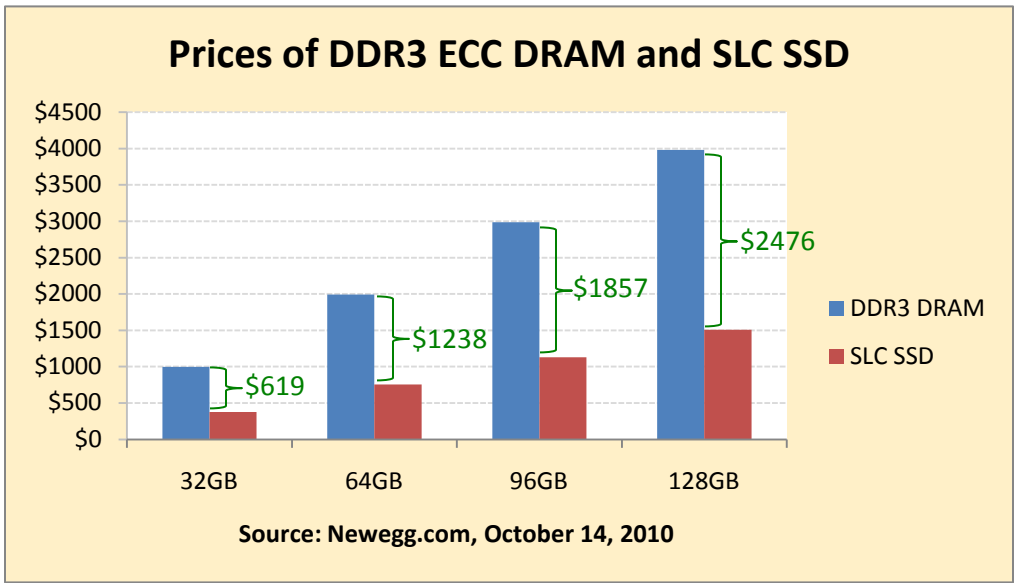

With CacheCade software and the SSDs, we were able to configure the test web server with 8GB of RAM and use the SSDs to provide the caching benefit rather than depending on large amounts of DDR3 server RAM to provide caching benefits. The CapEx benefits of using CacheCade software and SSDs include:

- The hardware expense savings increase as more caching capacity is added with SSDs rather than system memory. These savings are magnified when extended to multiple servers.
- SSDs allow the increase of storage cache size beyond what might be physically possible with system memory.
- A web hosting service can provide increased web server performance without adding additional web servers, further reducing CapEx.
- For a new web hosting installation, the total number of web servers required can be reduced while maintaining performance levels.

CacheCade and OpEx

Operational expenses are also reduced by using CacheCade software and SSDs for caching. Each web server in this type of configuration consumes less electric power than it would using the equivalent capacity of server memory. Power consumption can be further reduced by using fewer servers each configured with CacheCade software and SSDs.

Summary and Conclusion

The LSI MegaRAID CacheCade software for SSD caching feature provides outstanding performance improvements with minimal interaction required by an administrator. In addition, the SSDs that we added to this configuration are less expensive than adding DRAM memory in order to achieve similar performance benefits.

We observed at least 5x performance improvements in all measurements that we collected when comparing the hard-drive-only performance to the 2-SSD caching performance and at least 3x performance improvements when we added only one SSD. This performance improvement was observed in raw throughput and in response time.

This report is part of a series of reports evaluating the LSI MegaRAID SSD management functions. These reports are available at the following locations.

- Demartek Publishes Evaluation of LSI MegaRAID CacheCade Performance
- Demartek Publishes Evaluation of LSI MegaRAID FastPath Performance