LSI™ WarpDrive™ Acceleration Card Performance Evaluation in Access-Intensive Web Server Environments

Evaluation report prepared under contract with LSI Corporation

Introduction

When users have a choice of websites to do their business, providing acceptable response times is critical in retaining customers. As the client base grows, the high demand for serving pages may cause users to experience slow page delivery. To solve this problem, businesses may spread the web page demand across scores of systems each capable of providing web pages. This report will evaluate whether using solid state storage will allow fewer systems to be deployed.

Background

Solid state storage uses a robust version of the flash technology found on camera memory cards and is being used to augment or replace hard disk drive (HDD) storage. Since solid state storage has no spinning disks and movable read/write heads, it can provide from 10 to 1000 times more transactions per second than traditional HDD storage. Initial implementations of solid state storage have been solid state disks (SSD) which emulate the form factor and interface characteristics of a HDD.

LSI Corporation commissioned Demartek to evaluate their WarpDrive Acceleration Card, a solid state storage implementation that plugs into the PCI-Express® (PCIe) bus in a server. The purpose of this evaluation was to compare the performance of a web server using WarpDrive for web content storage versus HDD. WarpDrive is able to provide 300GBs of storage and by interfacing to PCIe, more transactions per second and higher throughput can be supported compared to HDD.

Evaluation Environment

These tests were conducted in the Demartek lab in Arvada, Colorado using Demartek servers, disk drives and networking infrastructure. LSI supplied the WarpDrive Acceleration Card solid state storage 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 WarpDrive Acceleration card provided outstanding performance with the web server compared to using HDD. By moving the web server content data to WarpDrive, we observed that 3 to 17 times more page views per minute were supported, with lower response time, over using a small number of high performance hard drives. These tests show that using WarpDrive solid state storage card can improve the performance of existing web servers or can provide the same performance with fewer web servers, lowering the total cost of ownership.

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.

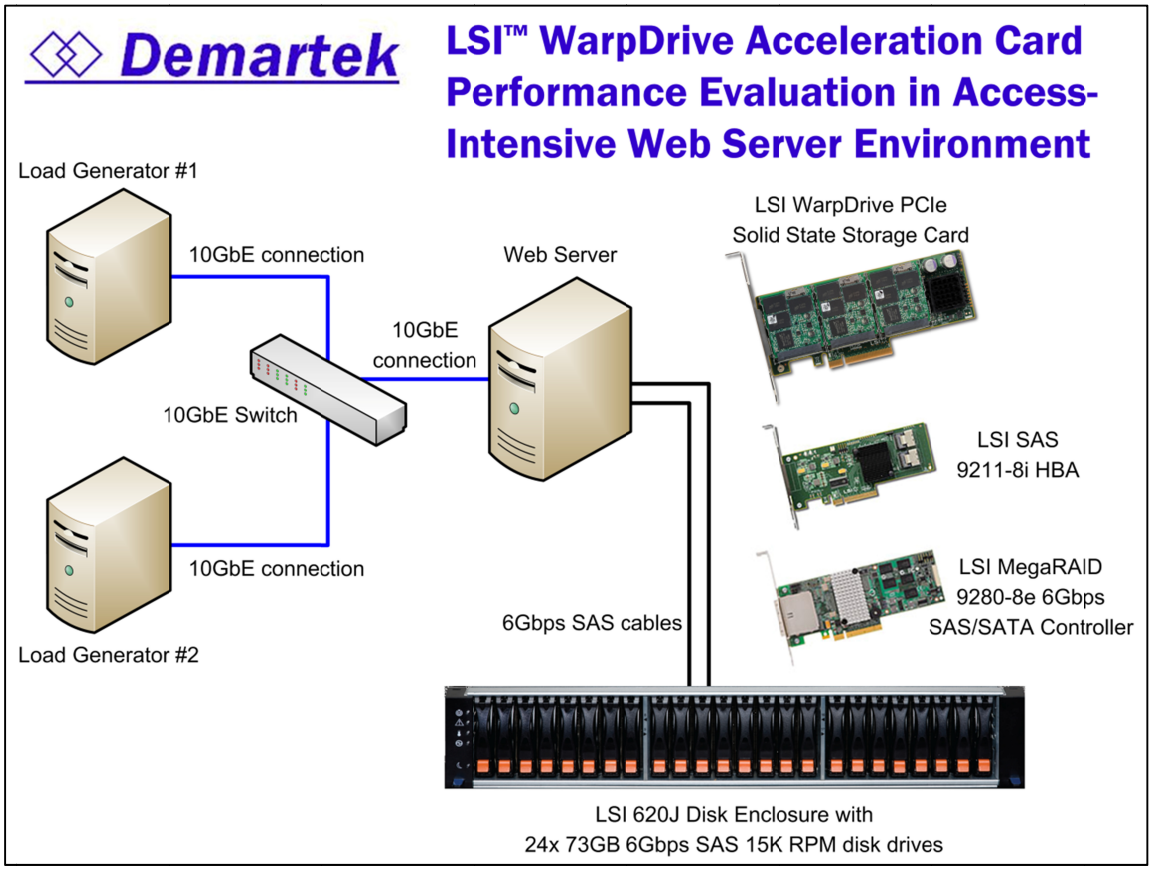

Web Server Configuration

- IBM System x3650 M2, PCIe 2.0
- Dual-processor, quad-core Intel Xeon X5570, 2.93GHz, 8 total cores, 16 logical processors
- 16GB RAM
- 2x Motherboard 1GbE NIC for general LAN traffic
- Intel® Ethernet Server Adapter X520-SR2 dual-port 10GbE Adapter dedicated to web server test traffic
- IBM ServeRAID-MR10i SAS/SATA Controller
  - Driver: megasas.sys, 4.5.1.64, 6/10/2009 2:37 PM
  - Internal hard drives: 5x Seagate Savvio 10K.3, 300GB 6Gb/s SAS disk drives
- LSI WarpDrive Acceleration Card PCIe solid state storage card
  - Firmware: 06.80.00.00
  - Driver: lsi_sss.sys, 2.10.34.0, 10/21/2010 6:07 PM
  - The WarpDrive is low-profile and does not require any extra power connectors.
- LSI SAS 9211-8i HBA (and external converter connector card)
  - Driver: lsi_sas2.sys, 2.0.2.71, 7/13/2009 3:59 PM

- LSI MegaRAID 9280-8e RAID Controller
  - Driver: megasas2.sys, 4.31.1.64, 8/9/2010 7:58 AM
- LSI 620J disk enclosure containing
  - 24x Seagate Savvio 2.5-inch, 15K.2, 73GB, 6Gb/s SAS disk drives.
- Windows Server 2008 R2 Enterprise Edition with IIS 7.5

Load Generator Configuration (x2)

- Dell PowerEdge 2900
- Dual-processor, quad-core Intel Xeon E5345, 2.33 GHz, 8 total cores
- 48GB RAM
- 2x Motherboard 1GbE NIC for general LAN traffic
- Intel® 10GbE XF SR Server Adapter dedicated to web server test traffic
- Windows Server 2008 R2 Enterprise Edition
- NeoLoad 3.1 web load generator software

Network Switch

- Cisco Nexus 5020, 40-port 10Gb Ethernet switch

Network Configuration

During initial testing, we found that our 1GbE network became the bottleneck when testing the high-performance LSI WarpDrive Acceleration card, even when using teamed 1GbE NICs. In order to remove the network bottleneck, we moved the test configuration over to our 10GbE network using 10GbE NICs in the servers, which allowed us to achieve significantly higher web server performance. 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 31GB 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 three unique graphic images.

2 – Solid State Storage Technology as Primary Storage

When deploying solid state storage as primary storage, web content is placed directly on the solid state media. Solid state storage provides significant performance improvement over spinning disks for all read and write activity. Typically, pairs of solid state storage cards should be deployed for redundancy, but in test cases and many actual web applications, only a single card may be used because the data is read-only. A backup or staging copy can be stored on HDD.

In order to test this primary storage solid state storage implementation, we deployed a web server with direct attached storage (DAS) using internal disk drives for the operating system. We deployed three storage targets for the web server content: a DAS storage enclosure partly and completely full of 2.5-inch, 73GB, 15K RPM SAS disk drives and the LSI WarpDrive Acceleration Card PCIe solid state storage card. Using a web load generator application, we sent at least 1,000,000 requests to the web server over 90-minute periods of time, and took several measures of web page and web server performance.

Test Configurations

We ran three tests, varying the storage configuration for each test. For these tests, the operating system was deployed on two internal 10K RPM SAS disks configured as a RAID1 volume of 278 GB using the on-board IBM ServeRAID-MR10i SAS/SATA Controller.

- **Test 1**: A single NTFS-formatted volume of 201GB consisting of a six-drive RAID10 configuration of 73GB 15K RPM hard disk drives was used for the web server content data. This volume was located in the LSI 620J external disk enclosure controlled by the LSI SAS 9211-8i HBA. A SAS pass-through card that provides external connections was used to connect to the external disk enclosure.
- **Test 2**: A single NTFS-formatted volume of 814GB consisting of a 24-drive RAID10 configuration of 73GB 15K RPM hard disk drives was used for the web server content data. This volume was located in the LSI 620J external disk enclosure controlled by the LSI MegaRAID 9280-8e RAID Controller. The 9280-8e controller was used because it can configure all 24 disk drives into a single RAID group and the 9211-8i HBA cannot.
- **Test 3**: A single NTFS-formatted volume of 279GB using the LSI WarpDrive Acceleration Card was used for the web server content data. To put additional load onto the WarpDrive Acceleration Card, the web server logs were also directed to this volume for this test. The LSI 620J disk enclosure and its disk drives were not connected to the server during this test.

Power Consumption

We also measured the power consumption during separate 15-minute tests for the server and storage solution. For tests 1 and 2, this included the total power consumed by the server, add-in HBA or RAID controller, and the LSI 620J disk enclosure and the 15K RPM disk drives. For test 1, there were only six disk drives in the disk enclosure. For test 3, this included the power consumed by the server including the LSI WarpDrive Acceleration Card. For test 3, no add-in HBA or RAID controller was present in the server and the disk enclosure was not connected to the server or powered on.

3 – Test Results

Our tests consisted of 90-minute runs with the client load generators making at least 1,000,000 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.

General Performance Overview

As expected, there were significant performance improvements when the web server content data was moved to the WarpDrive Acceleration Card. We also noticed that the WarpDrive Acceleration Card "ramped-up" very quickly to full performance.

To increase the load on the WarpDrive Acceleration Card, we also directed the web server logs to it for this test. The web server writes a log record for each page or file request from the web server clients. The number of records written to the web server log is approximately equal to the number of "hits" on the web server.

Moving the Bottleneck

This web server application workload using high-performance solid-state storage technology, when pushed to a heavier load, showed an example where ***the storage was no longer the bottleneck***, but the 1GbE network was the bottleneck, which required us to move to the 10GbE network. We believe that as more solid-state storage is deployed and pushed to heavier loads, I.T. professionals will discover similar new bottlenecks.

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.

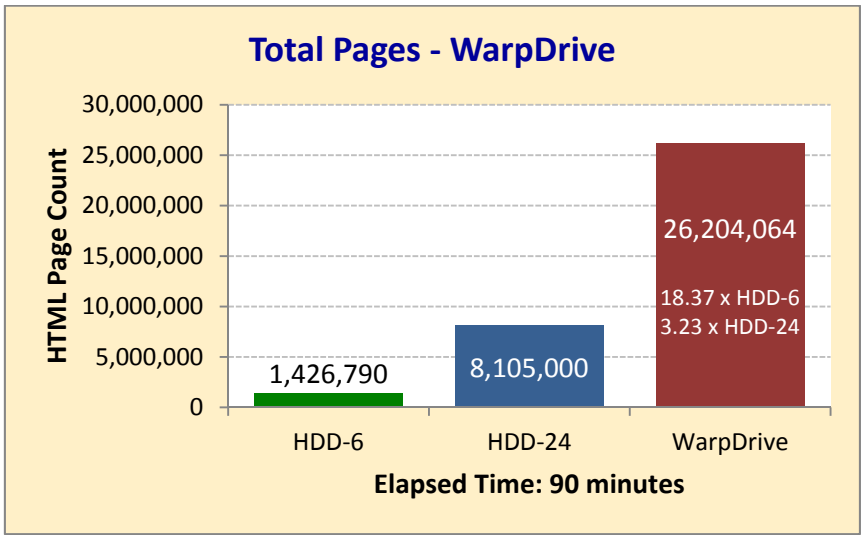

Total Hits

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 three graphic images, resulting in four "hits" per request.

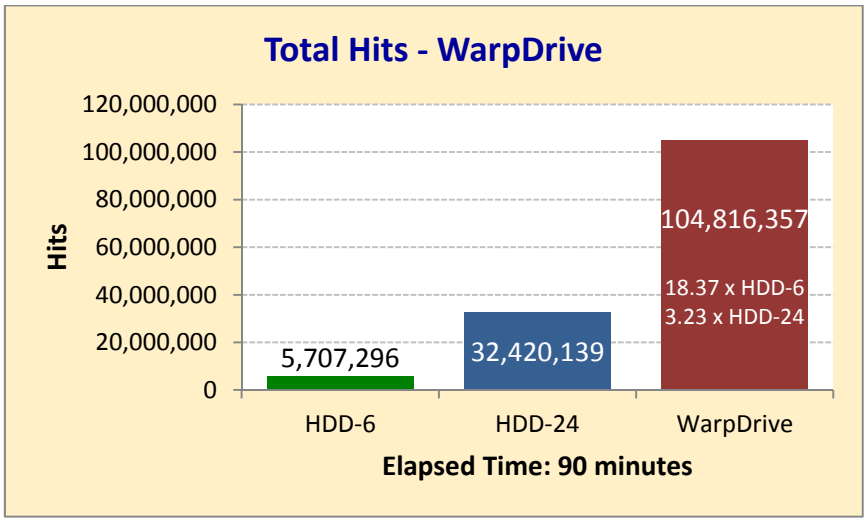

Total Throughput

Throughput is a measure of the total number of Gigabytes (GB) transferred from the web server to the clients over the duration of the test.

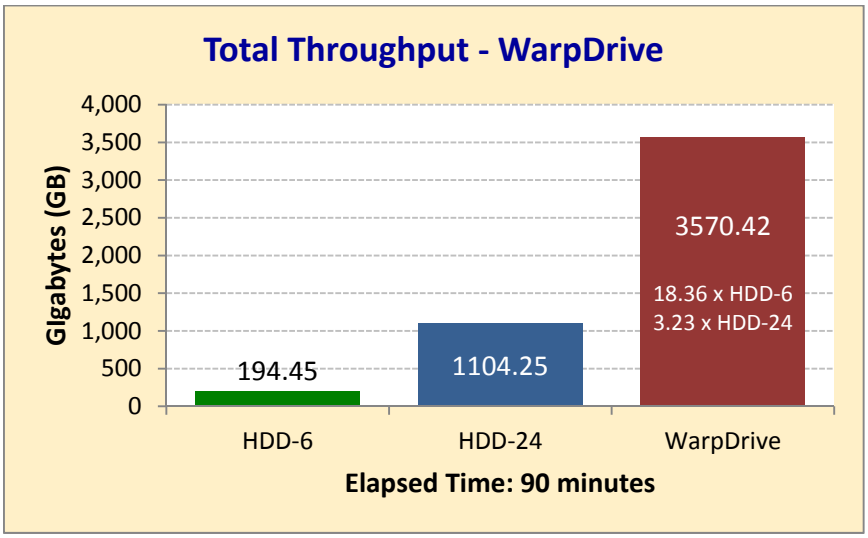

Performance Ramp-up

Before each test run, the web server and load generators were rebooted to ensure that there was no data leftover in their system and storage caches.

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.

- **Page Hit Rate** – This is the page hits per second of the HTML web content files accessed by the clients. This total counts the HTML text page and all the images (JPEG and PNG) for that page as one page hit.
- **Throughput** – This is the megabits per second of data transfer from the web server to the clients as a results of the client requests.
- **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.

The difference in performance for these web server tests was dramatic. The solid state storage configuration performance climbed to its maximum performance levels within the first two minutes and sustained those rates for the duration of the test. The added write load on the WarpDrive contributed to the variability of its performance. By moving the web server logs to the WarpDrive, we were able to achieve slightly higher overall web server performance than when the web server logs were written to spinning hard disks.

These tests show that adding the WarpDrive Acceleration Card PCIe solid state storage card significantly improves the website user experience and supports more concurrent users with fewer web servers or significantly extends the performance of existing web servers.

Server Performance

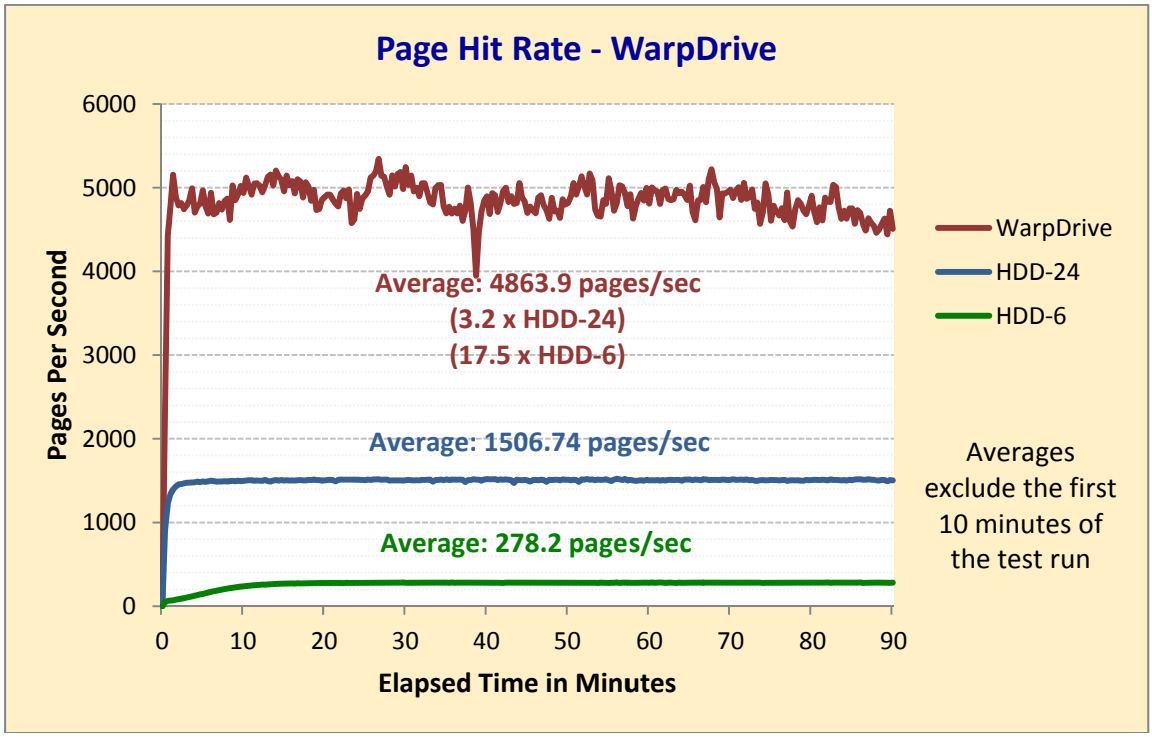

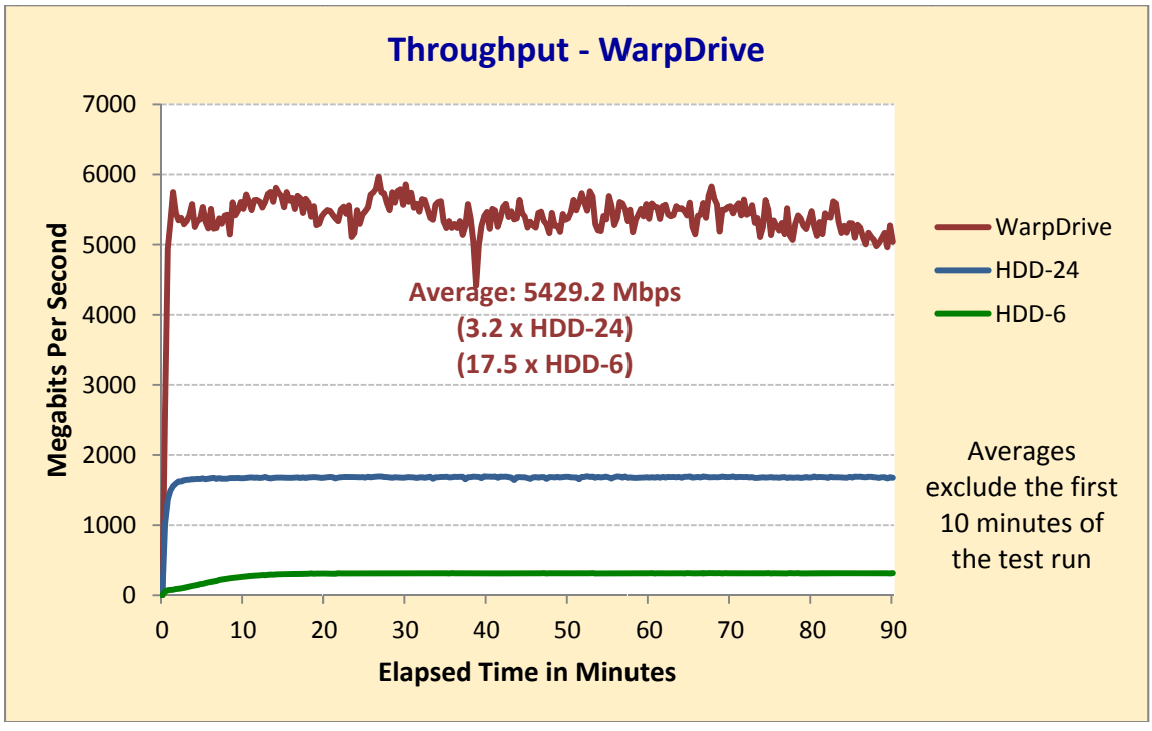

Server Responsiveness

Various studies have shown that even slight improvements in web site responsiveness improve the user's perception of the quality of the web site. The WarpDrive Acceleration Card provided dramatic improvement in responsiveness by reducing the average response time to less than 0.02 seconds.

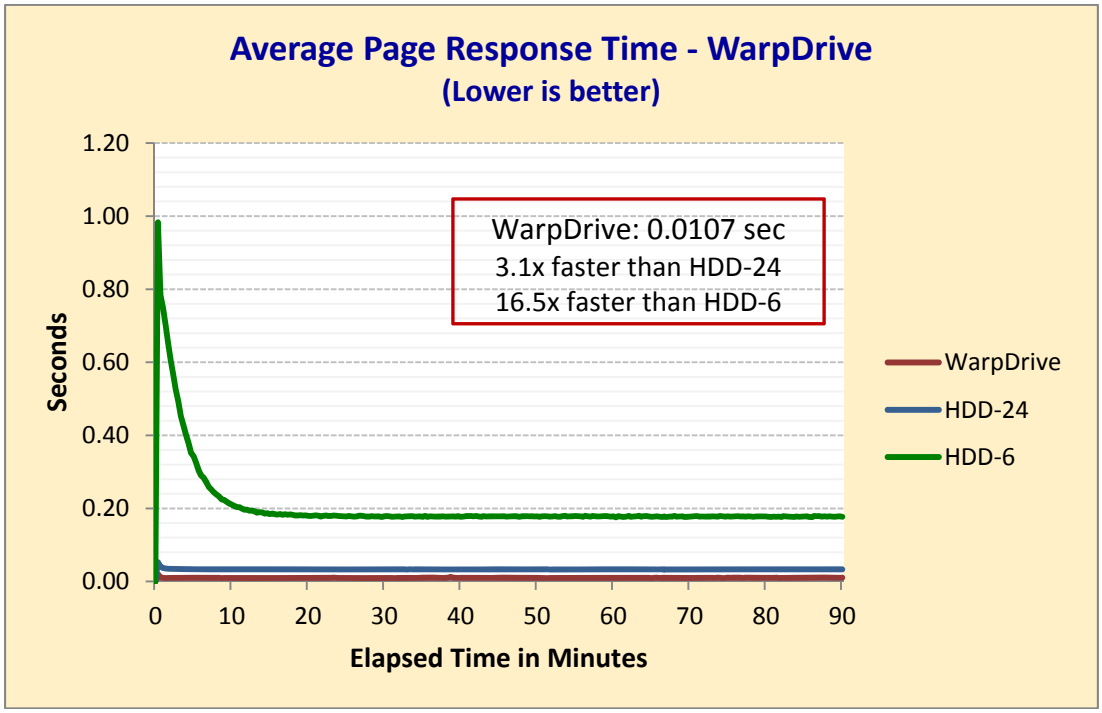

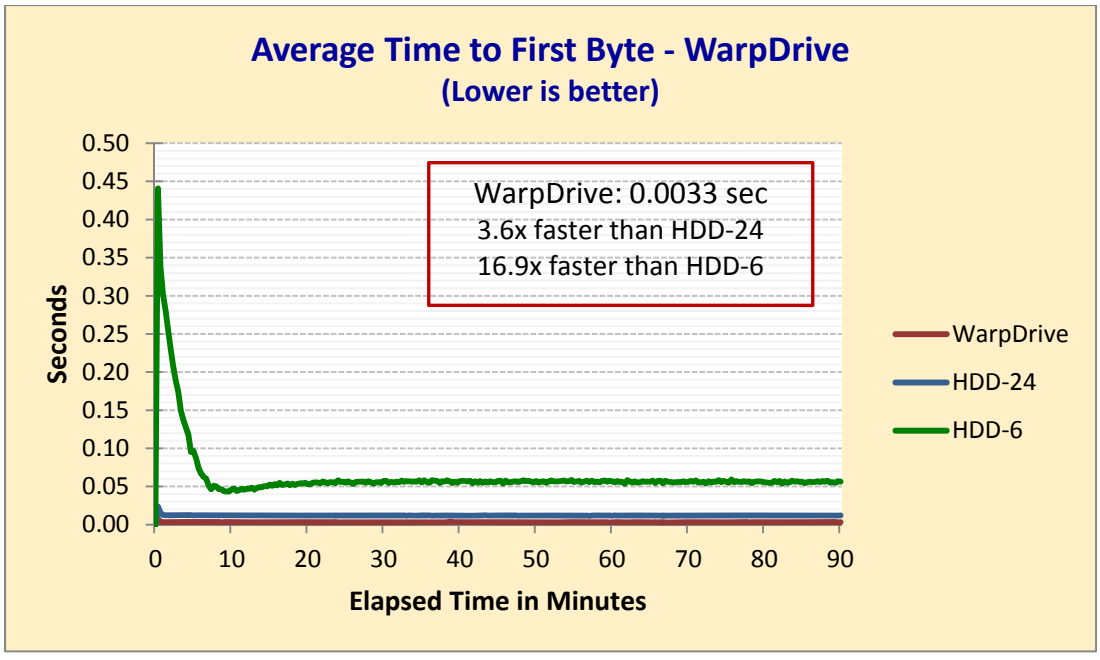

Power Consumption

Solid state storage brings with it the promise of higher performance while consuming less electrical power. We repeated our web server tests for 15 minutes each while measuring the power consumption using the AEMC PQL-120 Power Quality Logger, an inline electrical power measurement and logging device. For each test, we measured the total power consumption in watts for the web server solution. This includes the server, all the add-in cards and, for the hard disk configurations, the external disk enclosure and its disk drives. We measured the power consumption for two minutes before and after the 15-minute test periods while the web server was idle in order to get a steady-state measurement. In addition, we measured the power consumption of the server while it was idle, without any of the tested storage adapters and drives to get a baseline power consumption measurement.

The average power consumption was calculated over minutes 5 – 15 of the test period to allow for ramp-up time. Each solution reached approximately the same performance during minutes 5 – 15 that it did for the longer tests.

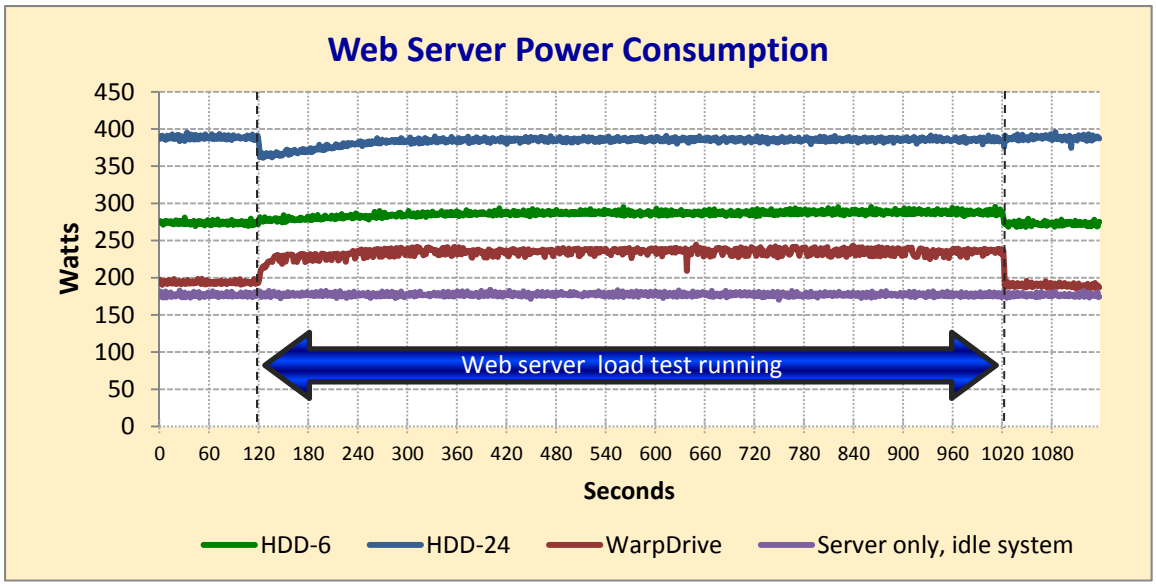

	Server only	HDD-6	HDD-24	WarpDrive
Average Watts consumed while idle	177.0	272.5	388.0	189.5
Difference above server only while idle	–	+95.5	+211.0	+12.5
Average Watts consumed during test	177.0	288.1	385.8	235.1
Difference above server only during test	–	+111.1	+208.8	+58.1

Of the configurations using the tested storage, the WarpDrive consumed the least amount of electrical power while providing the highest performance. As expected, the 24-drive configuration with the external disk enclosure consumed the most electrical power during the test period.

Given the observed performance of the 24-drive solution, and assuming a linear increase in performance by adding multiple sets of 24-drive disk enclosures, a solution with 72 drives in three

enclosures would almost equal the performance of the WarpDrive solution. Estimates of a 72-drive configuration are also included in the table below.

	Watts	Watts % of WarpDrive	Watts Factor		Hits/sec.	Hits/sec. % of WarpDrive
HDD-6	288.1	122%	1.2x		278	5.7%
HDD-24	385.8	164%	1.6x		1507	30.9%
HDD-72 *	*811.0*	*345%*	*3.5x*		*4521*	*92.9%*
WarpDrive	235.1	–			4864	–

* Estimated

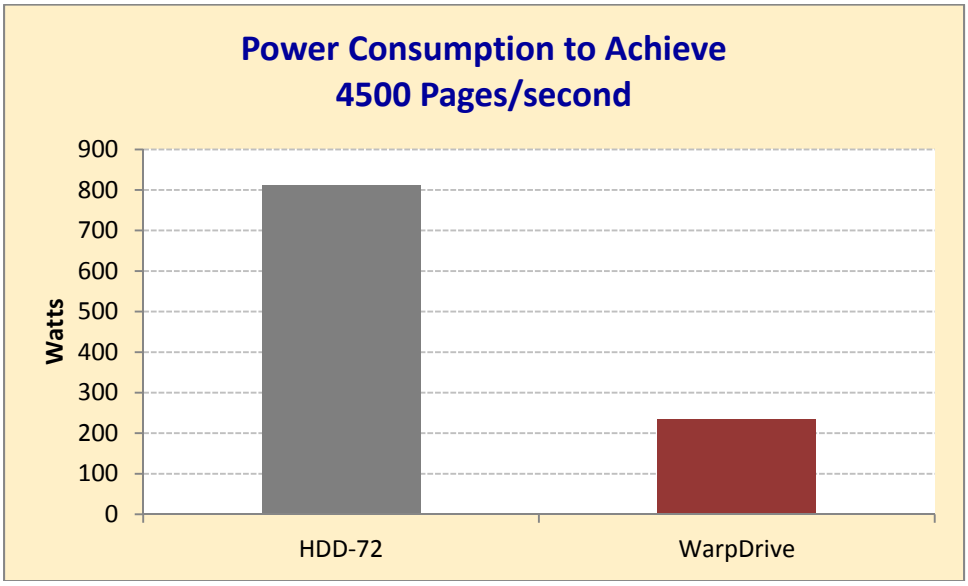

Rack Space

The WarpDrive solution consumes 2U of rack space, the 24-drive complete solution consumes 4U of rack space and a 72-drive complete solution would consume 8U of rack space.

4 – Total Cost of Ownership

An analysis of the LSI WarpDrive Acceleration Card would not be complete without discussing the cost of the solution. Two sets of cost figures should be examined:

1. The cost of the configurations tested
2. The estimated cost of an HDD configuration that would match the WarpDrive configuration

These two cost figures should be examined because in the first case, the basic costs of the hardware are required for budgeting reasons. The second case should be examined in order to show the cost to meet a specific performance level.

List prices are provided except where noted and are in US dollars. Pricing may vary by supplier, quantity and any negotiated discounts. Except for the LSI WarpDrive, list prices were obtained from CDW.com on December 10, 2010.

Costs

The cost of the server hardware is approximately $10,000 and is used for all the configurations. This price includes the processors, memory, on-board controller, 5 internal 10K SAS disk drives, internal DVD/CDROM drive and the Intel 10GbE adapter. The server is 2 rack units (2U) in height and the 620J disk enclosure is 2 rack units (2U) in height.

Configuration	HDD-6	HDD-24	HDD-72	WarpDrive
Server	$10,000	$10,000	$10,000	$10,000
LSI 620J disk enclosure	$2,279	$2,279	$6,837	—
SAS external cables (@$180)	$360	$360	$1,080	—
LSI 9211-8i HBA	$274	—	—	—
LSI 9280-8e RAID	—	$795	$795	—
73GB, 15K RPM disks (@$206)	$1,236	$4,944	$14,832	—
WarpDrive	—	—	—	$11,500
TOTAL	**$14,149**	**$18,378**	**$33,544**	**$21,500**
Pages/Sec.	**278**	**1507**	***4521**	**4864**

* Estimated

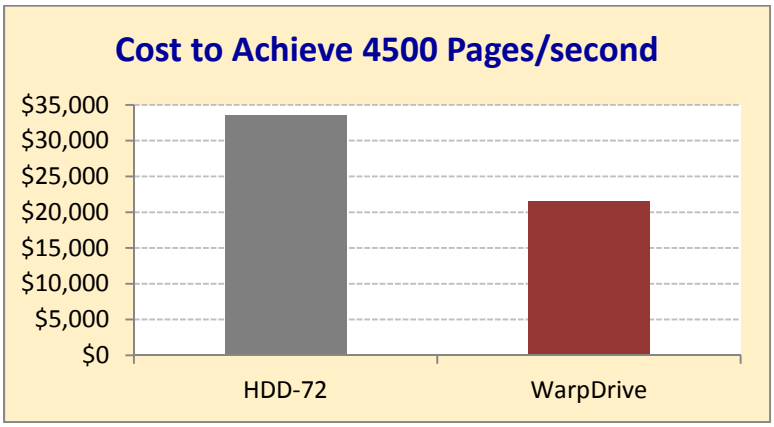

Summary and Conclusion

The LSI WarpDrive Acceleration Card PCIe solid state storage card used as primary storage provides outstanding performance improvements for the web server application that we tested.

We observed significant improvements in all measurements that we collected when comparing the WarpDrive Acceleration Card solid state storage card to hard drive configurations for the web server content data:

- More than 3x greater performance than the 24-drive, 2.5-inch 15K RPM hard-disk drives
- More than 18x greater performance than the six-drive, 2.5-inch 15K RPM hard disk drives
- Response time more than 3x faster compared to the 24-drive HDD configuration
- Response time more than 17x faster compared to the six-drive HDD configuration
- Lower power consumption
- Lower rack space consumption

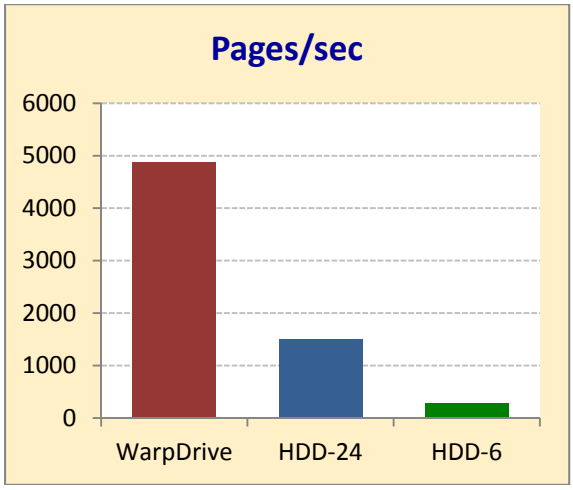

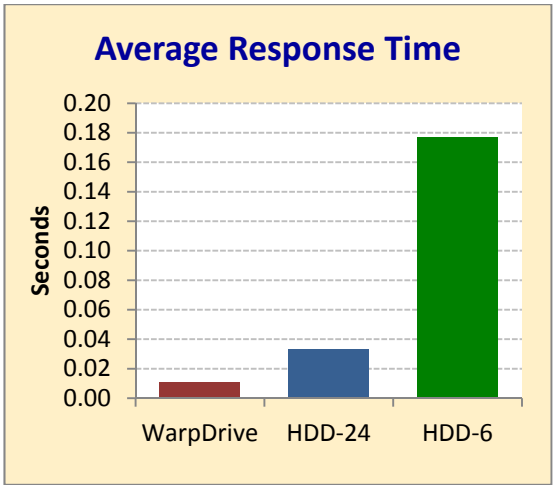

For high-traffic web sites, using WarpDrive cards inside of web servers would provide significantly higher performance than equivalent servers with HDD. In addition, the WarpDrive solution provides this performance benefit with lower power consumption and smaller rack space consumption, requiring less overall infrastructure to serve web pages. Because of the improved response time, web pages can be served with fewer servers if they are equipped with WarpDrive cards.

By deploying WarpDrive cards in web servers and consolidating servers, there can be tremendous economic benefit over similar web servers with HDD configurations.

This report is available at http://www.demartek.com/Demartek_LSI_Warpdrive_Performance_Evaluation_2011-01.html.