Enterprise IO Scalability: LSI® WarpDrive™ Acceleration Card Using Cisco UCS C-Series Rack-Mount Servers

Cisco UCS C460 M1 High-Performance Rack-Mount Server Delivers One Million IOPS

White Paper

February 2011

Contents

1. Introduction

Today's enterprise environments demand a high performance IT infrastructure to be competitive while data center managers are challenged to reduce power consumption due to increased energy cost and concerns over the ecological impact of power generation. In enterprise IT environments the performance of the storage subsystem is critical while it is also one of the top power consumers. So, optimizing the energy consumption of the storage subsystem can have a significant impact on the overall energy bill. In addition, many datacenters have limits to the available power and to cooling capacity. IT managers are continually challenged to support higher transaction loads and more storage capacity while using less power. It is often less costly to increase the power and cooling efficiency of an existing datacenter than to build and move into a new datacenter.

In this paper we look at the performance, scalability and energy efficiency aspects of storage systems based on the LSI WarpDrive Acceleration card SLP-300 in Cisco UCS C-Series rack-mounted servers. The study demonstrates that LSI WarpDrive cards combined with Cisco Unified Computing innovations can help to deliver industry leading performance and scalability at fraction of the power footprint of traditional hard disk drive based solutions.

Sustained throughput of over 1 Million IOPS on Cisco UCS C460 M1 High-Performance Rack-Mount Server (4-RU) and over 690K IOPS on UCS C250 M2 Extended-Memory Rack-Mount Server (2-RU) was achieved under small block random workloads of 70% read 30% writes ratio representing typical transaction processing workloads. The tests were conducted using out-of-the-box server, storage and Windows 2008 R2 operating system settings.

Cisco UCS C-Series Rack-Mount Servers

Cisco UCS C-Series Rack-Mount Servers [1] based on Intel Xeon 5500, 5600 and 7500 series processors, extend Cisco Unified Computing System innovations to an industry standard rack-mount form factor, including a standards-based unified network fabric, Cisco VN-Link virtualization support, and Cisco Extended Memory Technology.

Designed to operate both in standalone environments and as part of the Cisco Unified Computing System, these servers enable organizations to deploy systems incrementally—using as many or as few servers as needed—on a schedule that best meets the organization's timing and budget. Cisco UCS C-Series servers offer investment protection through the capability to deploy them either as standalone servers or as part of the Cisco Unified Computing System. Current generation of UCS C-Series Rack-Mount Servers is depicted in Figure 1 and a high level comparison of features is in Table 1.

One compelling reason that many organizations prefer rack-mount servers is the wide range of I/O options available in the form of PCI Express (PCIe) adapters. Cisco UCS C-Series servers supports spectrum of I/O options, which includes interfaces supported by Cisco as well as adapters from third parties.

Figure 1. From Left to Right, Cisco UCS C200 High-Density, C460 High-Performance, C250 Extended-Memory, and C210 General-Purpose Rack-Mount Servers

Table 1. Comparison of Cisco UCS C-Series Rack-Mount Server Features

	Cisco UCS C200 M1 and M2	Cisco UCS C210 M1 and M2	Cisco UCS C250 M1 and M2	Cisco UCS C460 M1
Ideal for	Production-level virtualization and mainstream data center workloads	Economical, high-capacity, reliable, internal storage; file, storage, database, and content-delivery	Demanding virtualization and large dataset workloads	High-performance, enterprise-critical stand-alone applications and virtualized workloads
Maximum memory	192 GB	192 GB	384 GB	512 GB
Internal disk drive	Up to 4	Up to 16	Up to 8	Up to 12
Built-In RAID	0 and 1 (SATA only)	0 and 1 (5 SATA drives only)		
Optional RAID	0, 1, 5, 6, and 10	0, 1, 5, 6, 10, 50, and 60	0, 1, 5, 6, 10, 50, and 60	0, 1, 5, 6, 10, 50, and 60
Integrated networking	2X integrated Gb Ethernet; 10 Gb unified fabric optional	2X integrated Gb Ethernet; 10 Gb unified fabric optional	4X integrated Gb Ethernet; 10 Gb unified fabric optional	2X Gigabit Ethernet LAN-on-motherboard (LOM) ports; 2X 10 Gigabit Ethernet ports
I/O via PCIe	Two, half-length, PCIe x8 slots: one full height and one low profile slot	Five, full-height, PCIe, x8 slots: two full length and three half length slots	Five PCIe slots: Three low-profile, half-length x8 slots; 2 full-height, half-length, x16 slots	Ten, full-heigh,t PCIe slots: 4 half-length slots, 6 three quarter length slots; 2 Gen 1 slots, 8 Gen 2 slots
Multicore Processors	Up to 2 Intel Xeon 5500 or 5600 Series	Up to 2 Intel Xeon 5500 or 5600 Series	Up to 2 Intel Xeon 5500 or 5600 Series	Up to 4 Intel Xeon 7500 Series

LSI Solid State Storage Technology

Solid-state drive (SSD) technology is being increasingly used to replace or supplement traditional hard disk drives in high performance applications. Benefits of SSDs over traditional hard disk drives include higher IOPS (I/O operations per second) and bandwidth, lower energy consumption, reduced latency and decreased rack space consumption.

As a top tier manufacturer of storage products, LSI is an active proponent of solid state storage with a broad solid state storage product line that includes both software and hardware. In its FastPath™ optimizer and CacheCade™ cache tiering software products, transactional performance of SSD configurations is boosted on MegaRAID 6Gb/s SAS/SATA Adapters.

Offering high performance with low latency and a low CPU burden, the small form factor LSI WarpDrive SLP-300 PCIe solid state storage card [2] is designed to maximize performance of workloads like online transaction processing (OLTP) that demands high IOPS. This technology utilizes industry-standard and widely deployed LSI SAS software and management infrastructure and supports extensive monitoring including health, error rate, and failure monitoring. Built based on the LSI Fusion MPT® Architecture, WarpDrive Cards have a low impact on server CPUs, leaving plenty of processing power for the applications. Multiple flash modules installed on the PCIe card enable both performance and capacity scaling within a card.

The model SLP-300 supports up to 300 GBytes storage capacity using 6 flash modules of 50 GBytes each. The flash modules are designed to provide enterprise-level performance and reliability. When overwriting data on flash memory, performance can be severely affected by having to wait for data to be erased. Each module's controller manages this issue by over provisioning storage to allow immediate writes and performing erasures when the module is idle. Another issue concerning flash technology is that after an area of the flash has experienced tens of thousands of erasures, data cannot be further erased. Using wear leveling and over provisioning, the controller

ensures that the module is able to handle data overwrites, thus enabling WarpDrive to offer a three-year warranty. As of this writing, no other PCIe flash card offers this particular security to customers.

Using the benchmarking tool IOMeter, WarpDrive has demonstrated up to 240K random 4K read IOPS and 200K random 4K Write IOPS and about 1.4GBytes/second transfer rates. As shown in section 3, multiple WarpDrive cards in servers delivers near linear scalability to meet the most demanding concurrent read/write IOPS requirements.

2. Tested Configurations and Methodology

Tests were conducted on a mid-range server configuration and a high-end server configuration. The mid-range server configuration consists of Cisco UCS C250 M2 Extended-Memory Rack-Mount Server equipped with two Intel® Xeon® Processor X5680 (3.33 GHz, 12MB L3 Cache, 130W), 192 GB of memory and 1 to 4 LSI WarpDrive cards. The high-end server configuration consists of Cisco UCS C460 M1 High-Performance Rack-Mount Server equipped with four Intel® Xeon® X7560 Processor (2.26GHz, 24MB cache, 130W), 256 GB of memory and 1 to 8 LSI WarpDrive cards as shown in Table 2. Both configurations ran Microsoft Windows® 2008 R2 operating system. In each configuration, PCIe 2.0 x8 slots were used. All the tests were conducted using default server, storage, driver and operating systems configuration settings.

IOMeter version 2008.06.18-RC [3], a tool maintained by the Open Source Development Lab (OSDL), originally developed by the Intel Corporation, and was used to simulate workloads.

IOMeter was configured with small block (4K and 8K) random read writes (70% reads and 30% writes) representative of typical transactional workloads. A series of IOMeter tests were conducted on the mid-range server configuration and high-end server configuration as depicted in Table 2 demonstrating performance and scalability. In each case processor utilization and energy consumption were captured. The power consumption was measured using a Yokogava WT210 power meter connected to the power input of the server. The measurement interval was 60 minutes with a 30 minute ramp-up and 15 minute ramp down. During the measurement interval both IOPS and power consumption were in a steady state. The reported performance and power metrics are the respective averages during the measurement interval.

Table 2. Tested Configurations

Configuration	Server	Processor	Memory	Disk Drives (OS)	Solid State Storage Configuration	Operating System
Mid-range	Cisco UCS C250 M2 Extended-Memory Rack-Mount Server	2 x Intel® Xeon® Processor X5680 (3.33 GHz, 12MB L3 Cache,130W)	24 x 8GB Dual Rank PC3L-10600 (DDR3-1333) Registered	2 x 300GB 6G SAS 10K DP disk drive	1, 2 and 4 LSI WarpDrive SLP-300	Microsoft Windows 2008 R2
High-end	Cisco UCS C460 M1 High-Performance Rack-Mount Server	4 x Intel® Xeon® X7560 Processor (2.26GHz, 24MB cache,130W)	32 x 8GB Dual Rank PC3-10600 (DDR3-1333) Registered	2 x 300GB 6G SAS 10K DP disk drive	1, 2, 4 and 4 LSI WarpDrive SLP-300	Microsoft Windows 2008 R2

3. Performance, Scaling and Energy Consumption

The results of our testing show that both Cisco UCS C250 M2 Extended-Memory Rack-Mount Server (mid-range) and Cisco UCS C460 M1 High-Performance Rack-Mount Server (high-end) configurations delivered exceptional performance, response time, energy efficiency and scalability as shown in Table 3.

Table 3. Capacity, Performance, Processor Utilization and Power Consumption

Number of LSI WarpDrive SLP-300s	Total Flash Capacity	Block size	IOPS	Average Response Time (ms)	Idle Processor Utilization (%)	Idle Power Consumption	Average Power Consumption (watts)
UCS C250 M2 Extended-Memory Rack-Mount Server							
1	300 GByte	4K	156,998	4.89	92.1	197	271
		8K	100,998	7.6	94.7	197	269
2	600 GByte	4K	357,514	2.14	84.5	209	296
		8K	222,407	3.45	89.4	209	287
4	1200 GByte	4K	696,984	2.2	73.8	239	354
		8K	431,807	3.55	76.6	239	341
Cisco UCS C460 M1 High-Performance Rack-Mount Server							
1	300 GByte	4K	155,354	3.29	92.4	822	840
		8K	101,766	5.03	94.3	822	841
2	600 GByte	4K	330,266	1.54	87.5	831	877
		8K	211,137	2.42	88.4	831	870
4	1200 GByte	4K	586,387	0.87	73.4	853	946
		8K	398,548	1.28	77.9	853	933
8	2400 GByte	4K	1,087,361	0.47	44.5	890	1,092
		8K	674,059	0.75	47	890	1,036

With a WarpDrive providing 300GB of storage capacity delivered over 155,000 IOPS for 4K block size and over 100,000 IOPS for 8K block size. High capacities and performance are achieved by installing additional WarpDrive cards. 2 and 4 WarpDrive cards were tested on UCS C250 M2 Extended-Memory Rack-Mount Server and 2, 4 and 8 WarpDrive cards were tested on Cisco UCS C460 M1 High-Performance Rack-Mount Server. The UCS C250 M2 Extended-Memory Rack-Mount Server with 2 WarpDrive cards delivered 357,514 IOPS for 4K and 222,407 IOPS for 8K block size respectively. With 4 WarpDrive cards near linear scaling was observed; 696,984 IOPS and 431,807 IOPS for 4K block size and 8K block size respectively.

The Cisco UCS C460 M1 High-Performance Rack-Mount Server with 2 WarpDrive cards delivered 330,266 IOPS for 4K block size and 211,137 IOPS for 8K block size. With 4 WarpDrive cards the throughput was 586,387 IOPS for 4K block size and 398,548 IOPS for 8K block size. The configuration with 8 WarpDrive cards delivered 1,087,361 IOPS for 4K block size and 674,059 for 8K block size. The performance and scaling of 1, 2, 4 and 8 WarpDrive cards are depicted in Figure 2.

Figure 2. IOPS Scaling with 1, 2, 4 and 8 WarpDrive Cards

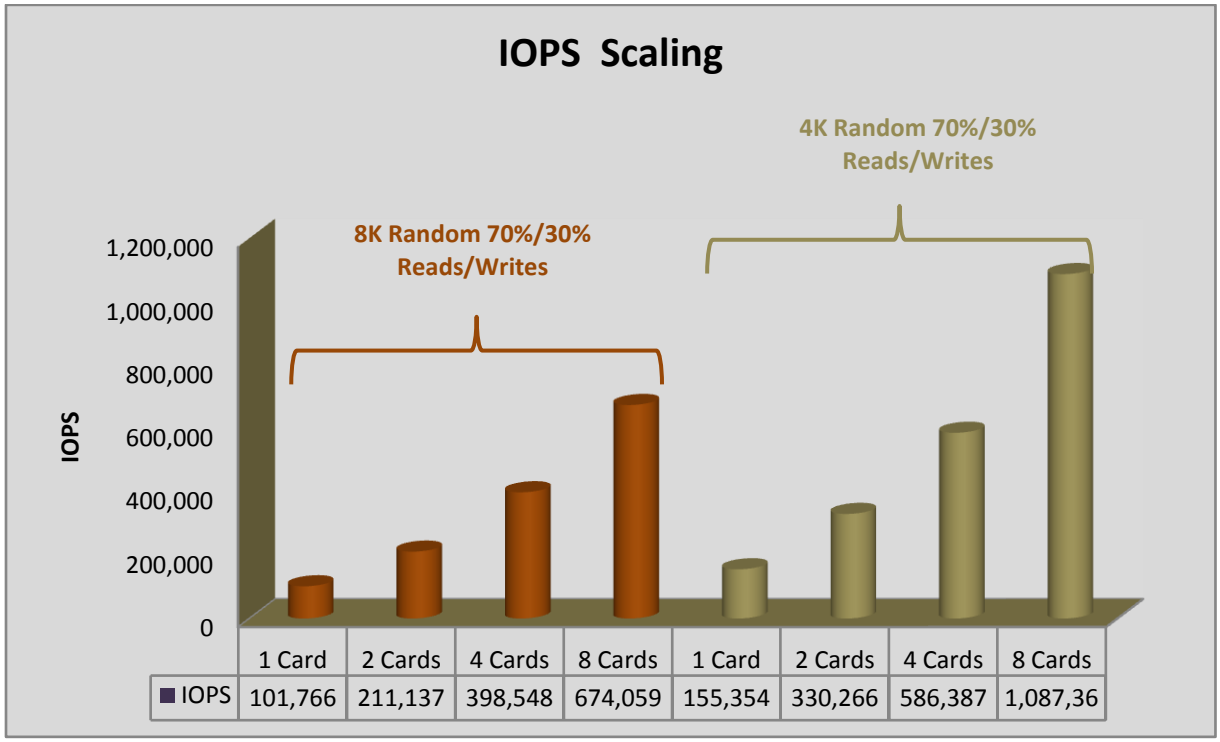

With a total capacity of 1.2TByte of flash storage capacity and 4.8 TByte of traditional hard disk drive capacity (using 8 internal 600 GByte SAS Small Form Factor hard disk drives) and up to 384 GByte of main memory, the Cisco UCS C250 M2 Extended-Memory Rack-Mount Server is a high performance and cost effective solution for mid-range transactional systems. The server has one unused PCIe slot for further scaling or other I/O options. At sustained performance of 696,984 IOPS, the average response time for IO operations was 2.2 ms, utilizing 26.2 % of processing power, and consuming 354 watts which is just 151 watts over the power consumption of the idle system with no WarpDrive Cards.

The Cisco UCS C460 M1 High-Performance Rack-Mount Server powered by Intel 7500 series of processors offers high thread density processor scalability and IO throughput. With a total capacity of 2.4 TByte of flash storage capacity, 7.2 TByte of traditional hard disk drive capacity (using 12 internal 600 GByte SAS Small Form Factor hard disk drives) and up to 1 TByte of main memory, the Cisco UCS C460 M1 High-Performance Rack-Mount Server is a high performance solution for high end transactional systems. The server has two unused PCIe slot for further scaling or other I/O options. At sustained performance of 1,087,361 IOPS, the average response time for IO operations was 0.47 ms, utilizing 55.5 % of processing power, and consuming 1092 watts which is just 292 watts over the power consumption of the idle system with no WarpDrive cards.

The traditional hard disk drives offer high capacity at much lower cost per storage capacity compared to SSDs. To achieve high IOPS, data needs to be striped across many drives so that collectively they can support the required rate. The penalty of this approach is the high cost for large number of disk drives, RAID controllers and storage enclosures to connect to the disk drives, and significantly higher power and space requirements. As shown in Figure 3, it will take over 3,600 traditional hard disk drives consuming over 24 kW power to achieve one million IOPS, assuming that 15K RPM SAS Small Form Factor hard disk drive delivers 300 IOPS and consume 5.5 watts, with an estimated 20% overheard for disk drive enclosures [5].

Figure 3. Power Estimates. WarpDrive vs. HDD

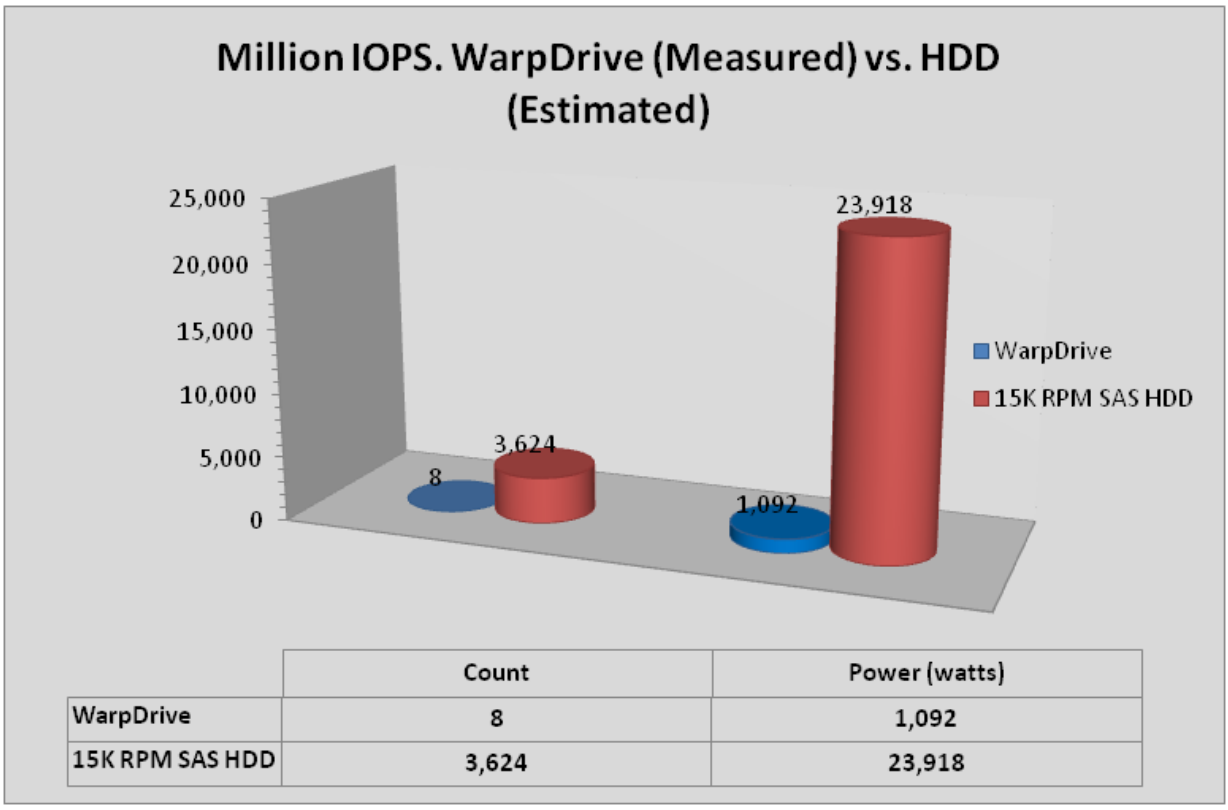

	Count	Power (watts)
WarpDrive	8	1,092
15K RPM SAS HDD	3,624	23,918

4. Conclusion

Traditional hard disk drive based storage systems offer high capacity and lower cost per capacity. With high performance at much lower power and cooling requirements, solid state storage products like the WarpDrive Acceleration card SLP-300 are becoming an increasingly viable option to replace or supplement traditional hard disk based storage systems. This study demonstrates that UCS C-Series Rack-Mounted servers and LSI Solid State Storage technology can meet the high performance and scalability demands for enterprise transactional applications at lower energy and smaller floor space footprints than traditional solutions.

5. References

1) Cisco UCS C-Series Rack-Mount Servers. http://www.cisco.com/en/US/products/ps10493/index.html
2) LSI WarpDrive PCIe solid state storage. http://www.lsi.com/storage_home/products_home/solid_state_storage/warpdrive_slp300/index.html
3) IO Meter Project. http://www.iometer.org/
4) M. Poess, R. Nambiar: Tuning servers, storage and database for energy efficient data warehouses. IEEE International Conference on Data Engineering 2010

6. Key Contributors

This document was made possible through the efforts of Raghunath Nambiar and Randy Glissmann, with thanks to Rakesh Sreekanth and Ranganathan Madhanagopal for their contributions.

Cisco Systems, Inc.

170 West Tasman Drive

San Jose, CA 95134-1706

USA

www.cisco.com

Tel: 408 526-4000

800 553-NETS (6387)

Fax: 408 527-0883

Document number: UCS-TR100030