

Hewlett-Packard Co.

Hardware-Accelerated iSCSI Solutions

Performance Comparison of Hardware and Software



Test
Summary

Accelerated iSCSI from ProLiant Servers to StorageWorks Storage Arrays

***Premise:** Software-based iSCSI initiators place undue burden on server resources. Hardware-accelerated iSCSI improves server performance and CPU efficiency by offloading much of the processing. Server and network administrators can establish the benefits of software versus hardware acceleration for iSCSI to make an informed decision on the type of NICs to provision in the servers.*

Hewlett-Packard Co. commissioned The Tolly Group to evaluate the performance benefits of multifunction server adapters in ProLiant servers featuring hardware-accelerated iSCSI over traditional software-based iSCSI initiators communicating to StorageWorks storage arrays.

Tolly Group engineers tested two end-to-end iSCSI storage solutions from HP. The first was a ProLiant server with an HP NC373T Gigabit Server Adapter communicating with a StorageWorks 1200 All in One (AiO) Storage System. Engineers also tested a ProLiant DL380 G5 server with an HP NC373i Gigabit Server Adapter communicating with a StorageWorks 2012i Modular Smart Array (MSA2012i).

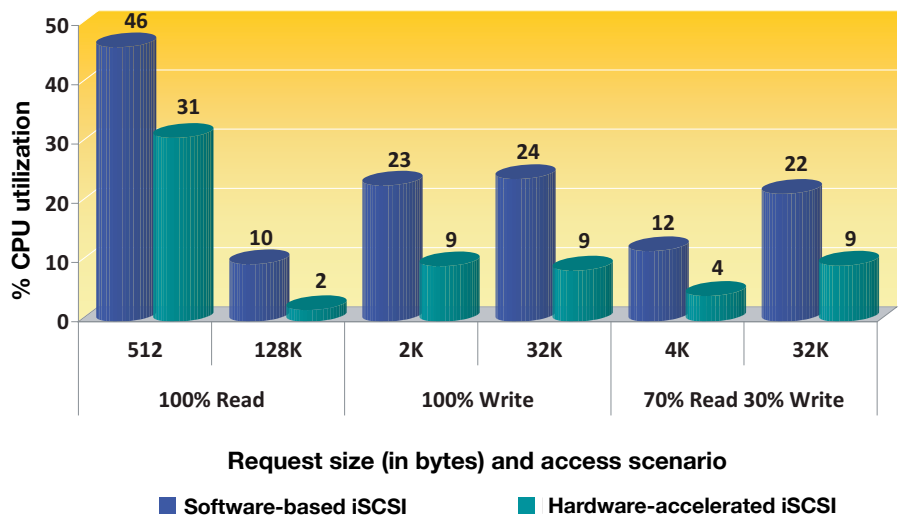
Tests compared the throughput and CPU utilization of software-based iSCSI (Microsoft iSCSI initiator) against hardware-accelerated iSCSI under various disk access scenarios, using and Iometer test tool.

Tests were conducted in April 2008.

Test Highlights

- ▶ Outperforms software-based iSCSI by reducing CPU utilization by up to 8X with hardware-accelerated iSCSI in the AiO 1200 solution
- ▶ Outperforms software-based iSCSI by up to 5X reduced CPU utilization with hardware-accelerated iSCSI in the MSA2012i solution
- ▶ Demonstrates close to theoretical maximum unidirectional iSCSI throughput over Gigabit Ethernet links, with up to 112 MBps (897 MBps) over one GbE link in the 1200 AiO solution, and up to 226 MBps (1,806 MBps) over two GbE links in the MSA2012i solution

CPU Utilization Improvement with Hardware-Accelerated iSCSI for HP StorageWorks 2012i Modular Storage Array (As Reported by Iometer 2006.07.27)



Note: StorageWorks MSA2012i was deployed with dual iSCSI controllers (across two ports) and HP ProLiant DL380 G5 Server with HP NC373i LOM Multi-function Gigabit Server Adapter.

Source: The Tolly Group, April 2008

Figure 1

Executive Summary

HP's end-to-end iSCSI solutions, equipped with NC373T and NC373i Multi-function Gigabit Server Adapters featuring hardware acceleration for iSCSI, demonstrate up to 10% improvement in throughput and up to 8X reduced CPU utilization compared to software-based iSCSI using the same server hardware.

Traditional software-based iSCSI initiators burden server resources with iSCSI processing overhead, effectively reducing the CPU efficiency and performance.

Hardware-accelerated iSCSI initiators offload most of the processing overhead from the CPU onto dedicated hardware, thereby freeing up server resources to increase the server efficiency and performance.

HP's end-to-end iSCSI solutions featuring hardware-accelerated iSCSI with either a dedicated NC373T Network Interface Card (NIC) or an HP NC373i LAN on Motherboard (LOM) hardware based on Broadcom's 5708 Network Controller. This solution was tested against software-based iSCSI using the Microsoft iSCSI Initiator built into the Windows Server 2003 operating system.

Test results in almost all test cases showed noticeable improvement in the peak

throughput and decreased CPU utilization (the lower, the better) with hardware-accelerated iSCSI versus software-based iSCSI. (See Figures 1-4 for detailed results.)

RESULTS

MSA2012i SOLUTION iSCSI PERFORMANCE

CPU utilization showed marked improvement; lower CPU utilization is better for overall performance. With a software-based solution, the CPU utilization was 10%; with hardware acceleration the CPU utilization was just 2%.

In the area of throughput, the MSA2012i solution demonstrated peak throughput of 1,675 Mbps over two GbE ports using software-based iSCSI in 100% read access mode, which

improved to 1,812 Mbps with hardware-accelerated iSCSI for the 128-KB request size.

In 100% write access mode, the MSA2012i solution demonstrated peak throughput of 1,327 Mbps over two GbE ports using software-based iSCSI, which improved to 1,481 Mbps with hardware-accelerated iSCSI for the 32 KB request size. CPU utilization also showed an improvement from 24% with a software-based solution to 9% with hardware acceleration.

In the 70% read-30% write access mode, CPU utilization again showed an improvement from 22% with a software-based solution to 9% with hardware acceleration.

In the same scenario, the MSA2012i solution demonstrated peak throughput of 1,521 Mbps

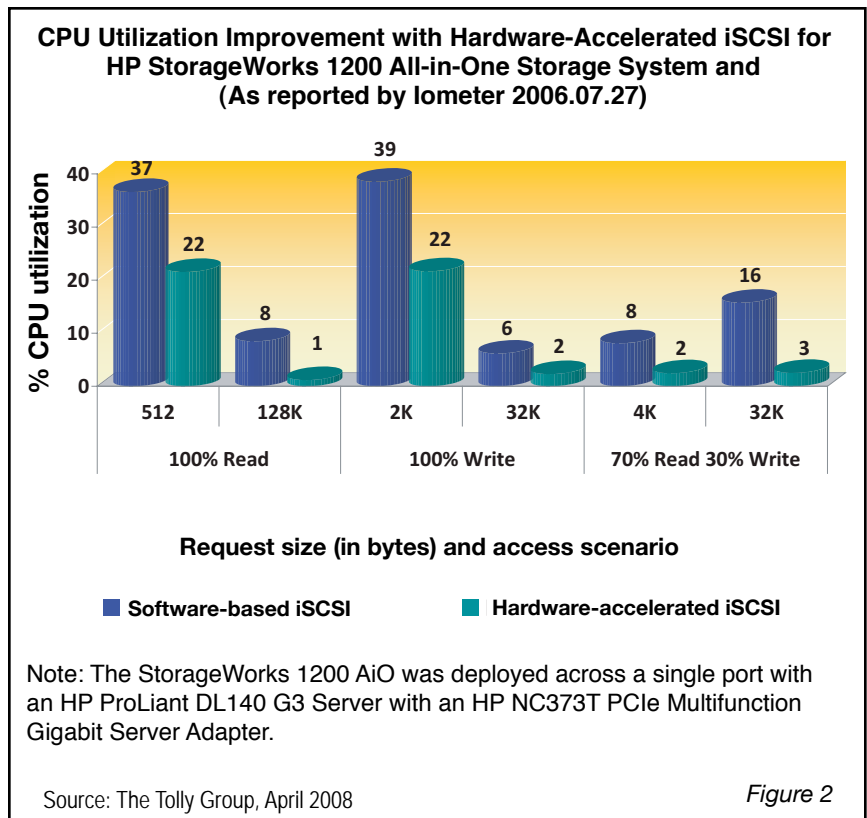


Figure 2

over two GbE ports using software-based iSCSI, which improved to 1,641 Mbps with hardware-accelerated iSCSI for the 32 KB request size. For detailed list of results for all access modes and the corresponding request sizes, see Figure 4.

AiO 1200 SOLUTION iSCSI PERFORMANCE

In 100% read access mode, CPU utilization showed a drop from 8% with software-based to 1% with hardware acceleration.

The AiO 1200 solution demonstrated peak throughput of 814 Mbps using software-based iSCSI, which improved to 897 Mbps using hardware-accelerated iSCSI for the 128

KB request size.

In 100% write access mode, CPU utilization showed a drop from 6% with a software-based solution to 2% with hardware acceleration. On the throughput front, the AiO 1200 solution demonstrated peak throughput of 825 Mbps using software-based iSCSI, which improved to 863 Mbps using hardware-accelerated iSCSI for the 32-KB request size.

In 70% read-30% write access mode, CPU utilization showed a drop from 16% with software-based acceleration to 3% with hardware acceleration. On the throughput side, the AiO 1200 solution demonstrated peak throughput of 805 Mbps using software-based iSCSI, which improved to 833 Mbps using

Hewlett-Packard Co.
Hardware-Accelerated iSCSI



Throughput and CPU Utilization Improvement Comparison

hardware-accelerated iSCSI for the 32 KB request size. For detailed results, see Figure 4.

TEST SETUP & METHODOLOGY

Tolly Group engineers tested two end-to-end iSCSI solutions from HP. See Figure 5 for hardware and software details of the solutions tested, and Figure 6 for the test bed layout.

Product Specifications

Vendor-supplied information not necessarily verified by The Tolly Group

Hardware-Accelerated iSCSI Solutions

HP NC373T

Product Specifications*

- Bus type: PCI Express (x4)
- Onboard memory: 92KB
- Connector: one RJ-45
- Cabling: Category 5, 5E, 6 or better
- Ethernet controller: Broadcom 5708
- Multifunction support: TCP/IP Offload for Windows, accelerated iSCSI
- IEEE Support: 802.3, 802.3u, 802.3x, 802.3ab, 802.1p, 802.1Q

HP ProLiant DL140 G3

Product Specifications*

Processor

- Dual-Core Intel® Xeon® Processor 5140 (2.33 GHz, 1333 MHz FSB)
 - Processor cores: 2
 - Processor cache: 8MB L2 (2x4MB)

Memory

- 1GB, PC2-5300 Fully Buffered DIMMs

Drives

- Hot plug 3.5 inch SAS 10K 36GB

Deployment

- Rack height: 1U
- Networking: Two embedded NC320i PCIe Gigabit Server Adapters
- 650-watt power supply

HP ProLiant DL380 G5

Product Specifications*

Processors

- Quad-Core Intel® Xeon® Processor E5430 (2.66 GHz, 80W, 1333 MHz FSB)
 - Processor cores: 4
 - Processor cache: 8MB L2 (2x4MB)

Memory

- 4GB, PC2-5300 DDR2 Fully Buffered DIMMs

Drives

- Hot plug 2.5 inch SAS 10K 36GB

Deployment

- Rack height: 2U
- Networking: Two embedded NC373i Multifunction Gigabit Network Adapters (with Broadcom 5708 Ethernet controllers) with TCP/IP Offload Engine for Windows and support for Accelerated iSCSI through an optional Licensing Kit
- Redundant power supplies and fans

For more information contact:

Hewlett-Packard Co.
3000 Hanover St., Palo Alto, CA 94304
Phone: (800) 474-6836, (650) 857-1501
URL: <http://www.hp.com>

The AiO 1200 solution consisted of an HP ProLiant DL140 G3 server with HP NC373T PCIe Multifunction Gigabit Server Adapter communicating with an HP StorageWorks 1200 All-in-One (AiO) Storage System.

The AiO 1200 was equipped with 12 spindles of 146 GB hard disk drives configured as eight iSCSI targets with two virtual drives of 20 GB each. Each virtual disk was mapped onto the server and offered as a target to the software or hardware iSCSI initiator.

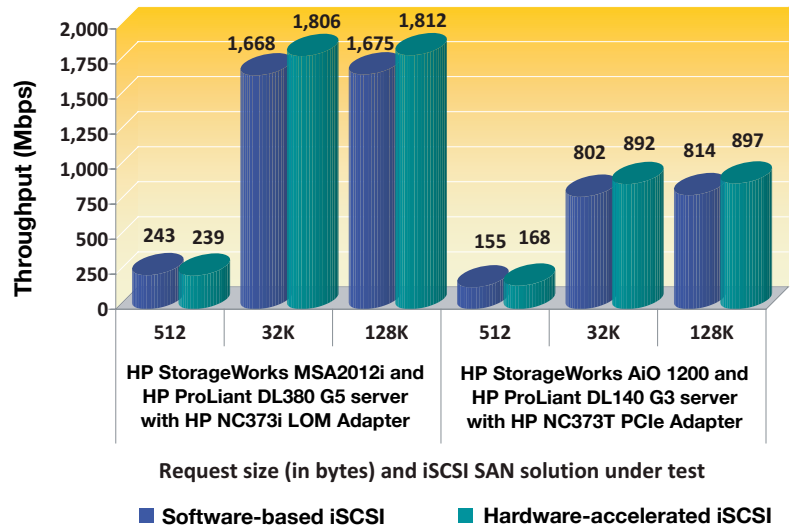
The DL140 G3 server connected to the iSCSI controller on the AiO 1200 via a Gigabit Ethernet (GbE) connection (copper cable with RJ-45 connectors) via an HP ProCurve 2824 switch.

The MSA solution consisted of an HP ProLiant DL380 G5 server with NC373i LOM Multifunction Gigabit Server Adapter communicating with an HP StorageWorks 2012i Modular Storage Array (MSA) with dual iSCSI controllers.

The MSA2012i was also equipped with 12 spindles of 146 GB hard disk drives which were divided into four virtual disks — two disks of 290 GB and two disks of 140 GB. The virtual disks were again mapped to the server and offered as targets to the software or hardware iSCSI initiators.

The DL380 G5 server was equipped with two GbE ports from two embedded NC373i LOM Multifunction Gigabit

Throughput Improvement with Hardware iSCSI Acceleration
(As Reported by Iometer 2006.07.27 with 100% Read Mode with 32 Outstanding I/O Requests)



Source: The Tolly Group, April 2008

Figure 3

Performance Improvement with Hardware-Accelerated iSCSI
(As Reported by Iometer 2006.07.27)

Test Scenario	Object Size (bytes)	Throughput (Mbps)		% CPU Utilization		
		Software-Based iSCSI	Hardware Accelerated iSCSI	Software-Based iSCSI	Hardware Accelerated iSCSI	
HP StorageWorks MSA 2012i (HP ProLiant DL380 G5 server and HP NC373i LOM adapter)	100% Read	512	243	239	46	31
		32K	1,668	1,806	14	4
		128K	1,675	1,812	10	2
	100% Write	2K	253	264	23	9
		8K	808	868	22	10
		32K	1,327	1,481	24	9
	70% Read 30% Write	4K	225	212	12	4
		8K	505	528	16	6
		32K	1,521	1,641	22	9
HP StorageWorks AiO 1200 (HP ProLiant DL140 G3 server with HP NC373T PCIe Adapter)	100% Read	512	155	168	37	22
		32K	802	892	11	2
		128K	814	897	8	1
	100% Write	2K	556	622	39	22
		8K	807	850	13	8
		32K	825	863	6	2
	70% Read 30% Write	4K	141	139	8	2
		8K	272	273	11	3
		32K	805	833	16	3

Note: 1K bytes = 1,024 bytes

Source: The Tolly Group, April 2008

Figure 4

Sever Adapters. These ports connected to the two iSCSI controllers on the MSA2012i via an HP ProCurve 2824 switch.

The Iometer test tool (version 2006.07.27) running on a ProLiant DL140 G3 server was used to measure the iSCSI performance of the two solutions.

Three different I/O access methods were tested with 100% sequential access pattern — 100% read, 100% write and 70% read 30% write. Throughput (in Mega-bytes per second, MBps) and the corresponding percentage of CPU utilization were recorded as reported by Iometer. Throughput values were reported as Megabits per second (Mbps).

For the 100% read mode, data request sizes of 512 bytes, 32 KB and 128 KB were tested. For the 100% write mode, data request sizes of 2 KB, 8 KB and 32 KB were tested. For the 70% Read 30% Write mode, data request sizes of 4 KB, 8 KB and 32 KB were tested.

For all the tests, Iometer was configured to use 32 outstanding I/O requests per target, with one worker per CPU core of the server under test. Thus, the AiO 1200 solution was tested with two workers since the ProLiant DL140 G3 server was equipped with a dual core Intel Xeon CPU. Similarly, the MSA2012i solution was tested with four workers since the ProLiant DL380 G5 server in the solution was equipped

Solutions Under Test Hardware and Software Details

The "MSA2012i Solution"	
HP ProLiant DL380 G5 server	
<ul style="list-style-type: none"> * Microsoft Windows Server 2003 Enterprise x64 Edition SP2 * Intel Xeon CPU 5150 @ 2.66 GHz, 4GB RAM, 36GB 10K RPM SAS HDD * Two HP NC373i Multifunction Gigabit Server Adapters (Driver ver 4.1.3.0) * HP NC373i iSCSI Device (Driver ver 4.1.8.0) * Microsoft iSCSI Initiator (HBA FW 1.5, Driver ver 5.2.3790.1653) * Smart Array P400 Controller (Driver ver 5.8.0.64) 	
HP StorageWorks 2012i Modular Storage Array	
<ul style="list-style-type: none"> * Dual iSCSI controllers * 12 spindles of 146GB 15K RPM SAS HDDs 	
Iometer console running on a HP ProLiant DL140 G3 server	
<ul style="list-style-type: none"> * Microsoft Windows Server 2003 Enterprise Edition SP1 * Intel Xeon CPU 5140 @ 2.33GHz, 1GB RAM, 75GB HDD * Two HP NC320i PCIe Gigabit Server Adapters (Driver ver 10.42.0.0) * Two HP NC360T PCIe Gigabit Server Adapters (Driver ver 9.9.14.0) * Iometer 2006.07.27 	
The "AiO 1200 Solution"	
HP ProLiant DL140 G3 Server	
<ul style="list-style-type: none"> * Microsoft Windows Server 2003 Enterprise x64 Edition SP2 * Intel Xeon CPU 5140 @ 2.33GHz, 1GB RAM, 75GB HDD * HP NC373T PCIe Multifunction Gigabit Server Adapter (Driver ver 4.1.3.0) * HP NC373i Virtual Bus Device (Driver ver 4.1.5.0) * HP NC373T iSCSI Device (Driver 4.1.8.0) * Microsoft iSCSI Initiator (FW ver 1.4, Driver ver 5.2.3790.1653) * Multipath Support (Driver ver 1.20.3790.2099) 	
HP StorageWorks 1200 All-in-One Storage System	
<ul style="list-style-type: none"> * Microsoft Windows Storage Server 2003 SP2 * Intel Xeon CPU 3070 @ 2.66 GHz, 2GB RAM * 12 spindles of 146GB 15K RPM SAS hard disks * Two HP NC324i PCIe Dual Port Gigabit Server Adapters (Driver ver 10.39.0.0) * HP iLO Management Channel Interface Driver (Driver ver 1.12.0.0) * Microsoft iSCSI Initiator (HBA FW ver 1.5, Driver ver 5.2.3790.1653) * Smart Array P400 Controller (RAID) (Driver ver 6.8.0.32) * WinTarget Local Mount Adapter (Driver ver 5.2.3790.3412) 	
Iometer console running on a HP ProLiant DL140 G3 server	
<ul style="list-style-type: none"> * Microsoft Windows Server 2003 Enterprise x64 Edition SP2 * Intel Xeon CPU 5140 @ 2.33GHz, 1GB RAM, 1x 75GB HDD * Two HP NC320i PCIe Multifunction Gig Server Adapters (Driver ver 10.62.1.2) * Iometer 2006.07.27 	

Source: The Tolly Group, April 2008

Figure 5

with a quad core Intel Xeon processor.

Tests initially were run with the software-based initiator enabled for the iSCSI traffic, and then repeated with the hardware acceleration enabled.

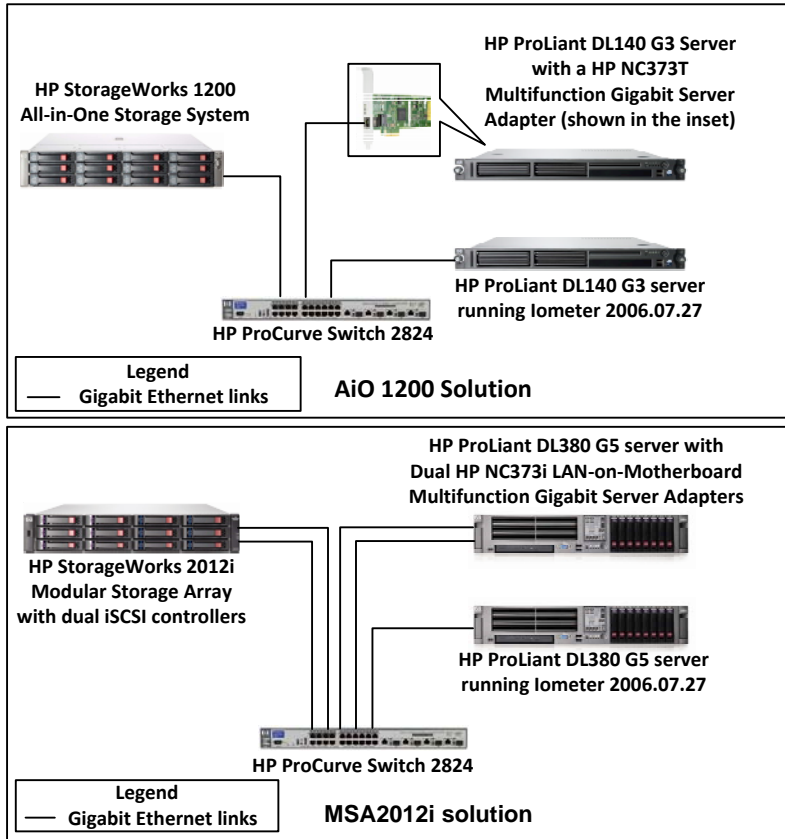
Software-based iSCSI tests used the Microsoft iSCSI Initiator built into the Windows Server 2003 Enterprise x64 Edition SP2

operating system running on the servers under test.

Hardware-accelerated iSCSI tests used the iSCSI initiator running in hardware on the NC373i LOM adapter and the NC373T NIC.

All tests were run three times, and the results were averaged to ensure consistency of results.


Test Bed Diagram for Performance Tests



Source: The Tolly Group, April 2008

Figure 6

The Tolly Group is a leading global provider of third-party validation services for vendors of IT products, components and services.



The company is based in Boca Raton, FL and can be reached by phone at (561) 391-5610, or via the Internet at:
 Web: <http://www.tolly.com>,
 E-mail: sales@tolly.com

Test Tool Summary

Vendor	Product	Web URL:
Public Domain	iometer 2006.07.27	http://www.iometer.org

Terms of Usage

USE THIS DOCUMENT ONLY IF YOU AGREE TO THE TERMS LISTED HEREIN.

This document is provided, free-of-charge, to help you understand whether a given product, technology or service merits additional investigation for your particular needs. Any decision to purchase must be based on your own assessment of suitability.

This evaluation was focused on illustrating specific features and/or performance of the product(s) and was conducted under controlled, laboratory conditions and certain tests may have been tailored to reflect performance under ideal conditions; performance may vary under real-world conditions. Users should run tests based on their own real-world scenarios to validate performance for their own networks. Commercially reasonable efforts were made to ensure the accuracy of the data contained herein but errors and/or oversights can occur. In no event shall The Tolly Group be liable for damages of any kind including direct, indirect, special, incidental and consequential damages which may result from the use of information contained in this document.

The test/audit documented herein may also rely on various test tools the accuracy of which is beyond our control. Furthermore, the document relies on certain representations by the sponsor that are beyond our control to verify. Among these is that the software/hardware tested is production or production track and is, or will be, available in equivalent or better form to commercial customers.

When foreign translations exist, the English document is considered authoritative. To assure accuracy, only use documents downloaded directly from The Tolly Group's Web site.

All trademarks are the property of their respective owners.