

Iometer iSCSI performance comparison of a Gigabit Ethernet adapter and an HBA

Executive summary

Intel Corporation (Intel) commissioned Principled Technologies (PT) to compare the iSCSI performance of its Gigabit Ethernet Adapter to an iSCSI HBA (Host Bus Adapter) using Iometer. We ran transmit, receive, and bidirectional tests with small to medium I/O sizes on the following adapters:

- Intel PRO/1000 PT Dual Port Server Adapter
- QLogic QLA4052C Dual Port iSCSI TOE to HBA

To compare their overall performance, we used the average percentage performance increase of the median Iometer results of each of the controllers for the file request sizes we tested. To obtain this average percentage performance increase, we used the following formula:

$(\text{Intel adapter result} / \text{QLogic HBA result}) - 1$

We used the Megabytes per second (MB/s) measurement from the Iometer output file as the result for each adapter. We then obtained the average of the percentage win for all the file request sizes we list.

The Key Findings summarize these comparisons.

In addition, the Test results section provides complete tables of all results.

Workload

Iometer is a tool that measures the performance of a computer's I/O subsystem under a load the tester specifies. Iometer functions as both a workload generator and a measurement tool. In our testing, we used it to gauge network throughput to attached drives using two different iSCSI controllers.

The Iometer test suites we used were read, write, and RW (read/write) with the following file sizes (from smallest to largest):

- 512B
- 1KB
- 2KB
- 4KB
- 8KB
- 16KB
- 32KB
- 64KB

The test suite performed one run of each file size for each type of activity. All file size tests used 100 percent sequential distribution and 4 outstanding I/Os per target.

We report two key Iometer measurements: MB per second, and the CPU utilization that Iometer measures.

KEY FINDINGS

- The Intel PRO/1000 PT Dual Port Server Adapter achieved an average 53.6 percent increase over the QLogic QLA4052C Dual Port iSCSI TOE to HBA in 100 percent sequential read access tests.
- The Intel PRO/1000 PT Dual Port Server Adapter achieved an average 32 percent increase over the QLogic QLA4052C Dual Port iSCSI TOE to HBA in 100 percent sequential write access tests.
- The Intel PRO/1000 PT Dual Port Server Adapter achieved an average 20.3 percent increase over the QLogic QLA4052C Dual Port iSCSI TOE to HBA in 100 percent sequential read/write access tests.

Test results

Figure 1 shows the median Iometer read throughput results, in MB/s, as well as the percentage CPU utilization.

I/O size	QLogic read results		Intel read results	
	Throughput	CPU utilization	Throughput	CPU utilization
512	7.974001	5%	13.320326	28%
1KB	15.169279	6%	24.038512	39%
2KB	27.024512	5%	43.422965	26%
4KB	47.781323	4%	73.443504	36%
8KB	94.417795	7%	128.12557	27%
16KB	119.050173	5%	184.05879	34%
32KB	138.22153	3%	210.16289	22%
64KB	150.643274	2%	221.259489	17%

Figure 1: Iometer read throughput results, in MB/s, and the percentage CPU utilization.
For throughput, higher numbers are better. For CPU utilization, lower numbers are better.

Figure 1's results demonstrate that on the Iometer read tests the Intel PRO/1000 PT Dual Port Server Adapter achieved from 35 to 67 percent higher throughput than the QLogic QLA4052C Dual Port iSCSI TOE to HBA.

Figure 2 illustrates the median iometer read throughput results, in MB/s, and percentage CPU utilization, for all 8 I/O sizes we tested.

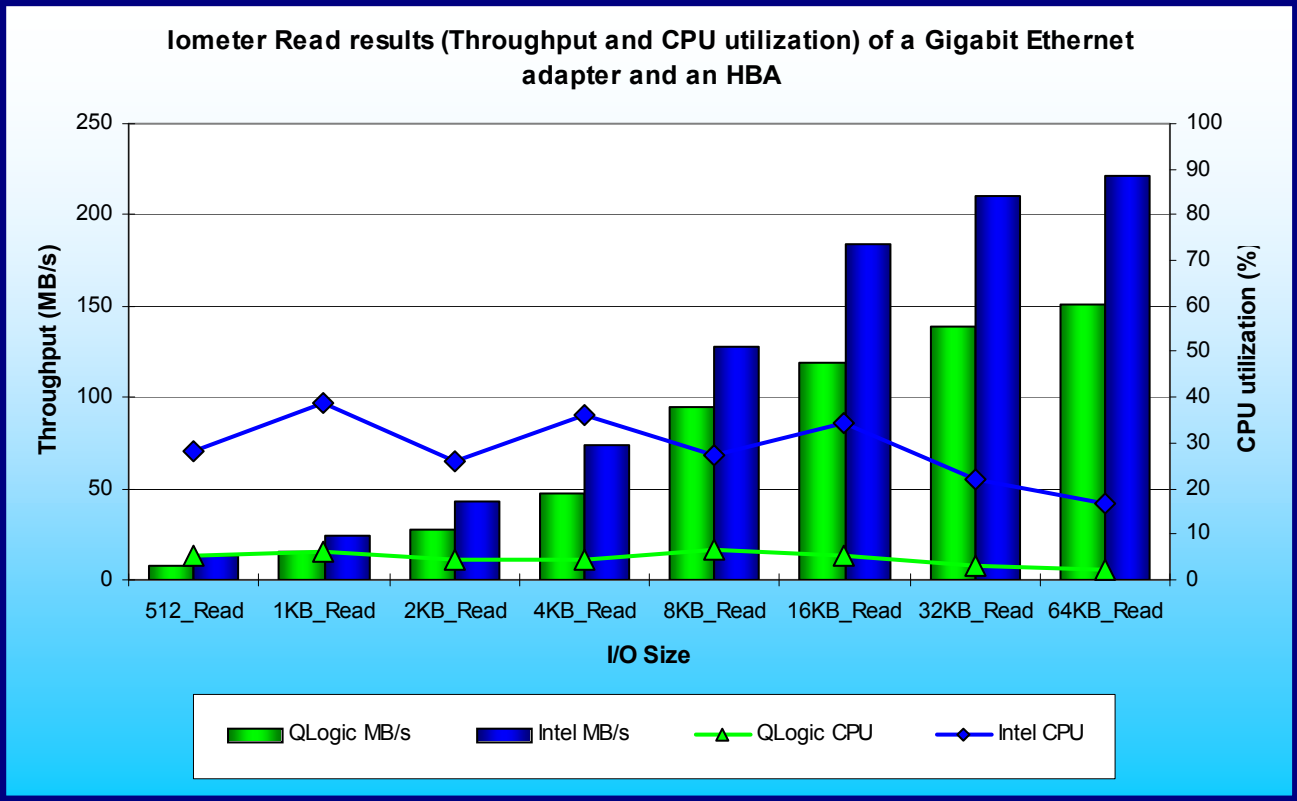


Figure 2: iometer read throughput results, in MB/s, and the percentage CPU utilization. For throughput, higher numbers are better. For CPU utilization, lower numbers are better.

Figure 3 shows the median iometer write throughput results, in MB/s, as well as the percentage CPU utilization.

I/O size	QLogic write results		Intel write results	
	Throughput	CPU utilization	Throughput	CPU utilization
512	7.516648	5%	8.403251	13%
1KB	14.15661	5%	17.105465	22%
2KB	26.511191	4%	34.165799	13%
4KB	47.921441	4%	60.83594	22%
8KB	55.868859	4%	105.541882	12%
16KB	99.839044	4%	135.306651	13%
32KB	150.189536	3%	171.809278	7%
64KB	143.130799	2%	184.484014	8%

Figure 3: iometer write throughput results, in MB/s, and the percentage CPU utilization. For throughput, higher numbers are better. For CPU utilization, lower numbers are better.

Figure 3's results demonstrate that on the iometer write tests the Intel PRO/1000 PT Dual Port Server Adapter achieved from 12 to almost 89 percent higher throughput than the QLogic QLA4052C Dual Port iSCSI TOE to HBA.

Figure 4 illustrates the median Iometer write throughput results, in MB/s, and percentage CPU utilization, for all 8 I/O sizes.

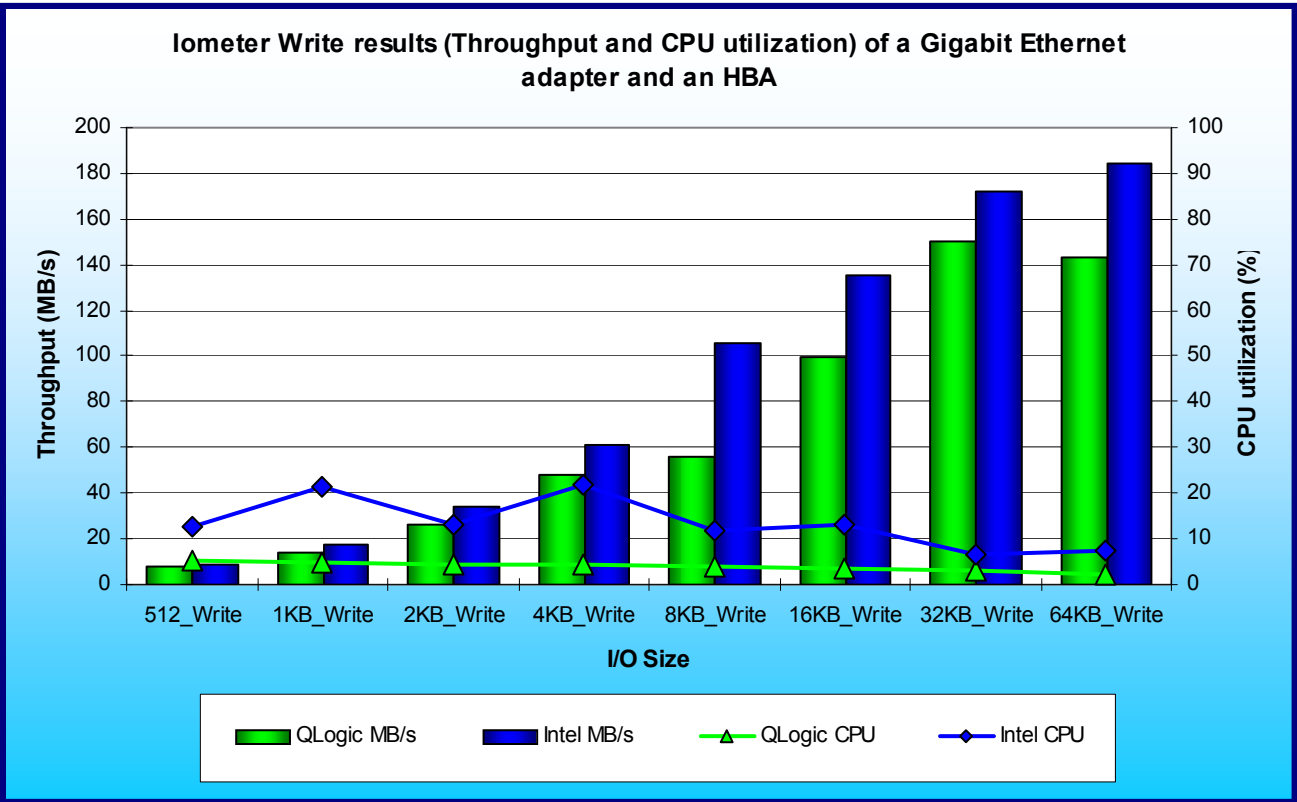


Figure 4: Iometer write throughput results, in MB/s, and the percentage CPU utilization. For throughput, higher numbers are better. For CPU utilization, lower numbers are better.

Figure 5 shows the median Iometer read/write throughput results, in MB/s, as well as the percentage CPU utilization.

I/O size	QLogic read/write results		Intel read/write results	
	Throughput	CPU utilization	Throughput	CPU utilization
512	7.464417	5%	9.604424	17%
1KB	11.884849	4%	14.245055	21%
2KB	19.089945	3%	24.053575	12%
4KB	26.464326	2%	34.618611	15%
8KB	47.453299	3%	56.997068	10%
16KB	77.009767	3%	84.451919	12%
32KB	114.513548	2%	128.53355	10%
64KB	143.666406	2%	165.098311	13%

Figure 5: Iometer read/write throughput results, in MB/s, and the percentage CPU utilization. For throughput, higher numbers are better. For CPU utilization, lower numbers are better.

Figure 5's results demonstrate that on the Iometer read/write tests the Intel PRO/1000 PT Dual Port Server Adapter achieved from almost 10 to over 30 percent higher throughput than the QLogic QLA4052C Dual Port iSCSI TOE to HBA.

Figure 6 illustrates the median Iometer read/write throughput results, in MB/s, and percentage CPU utilization, for all 8 I/O sizes.

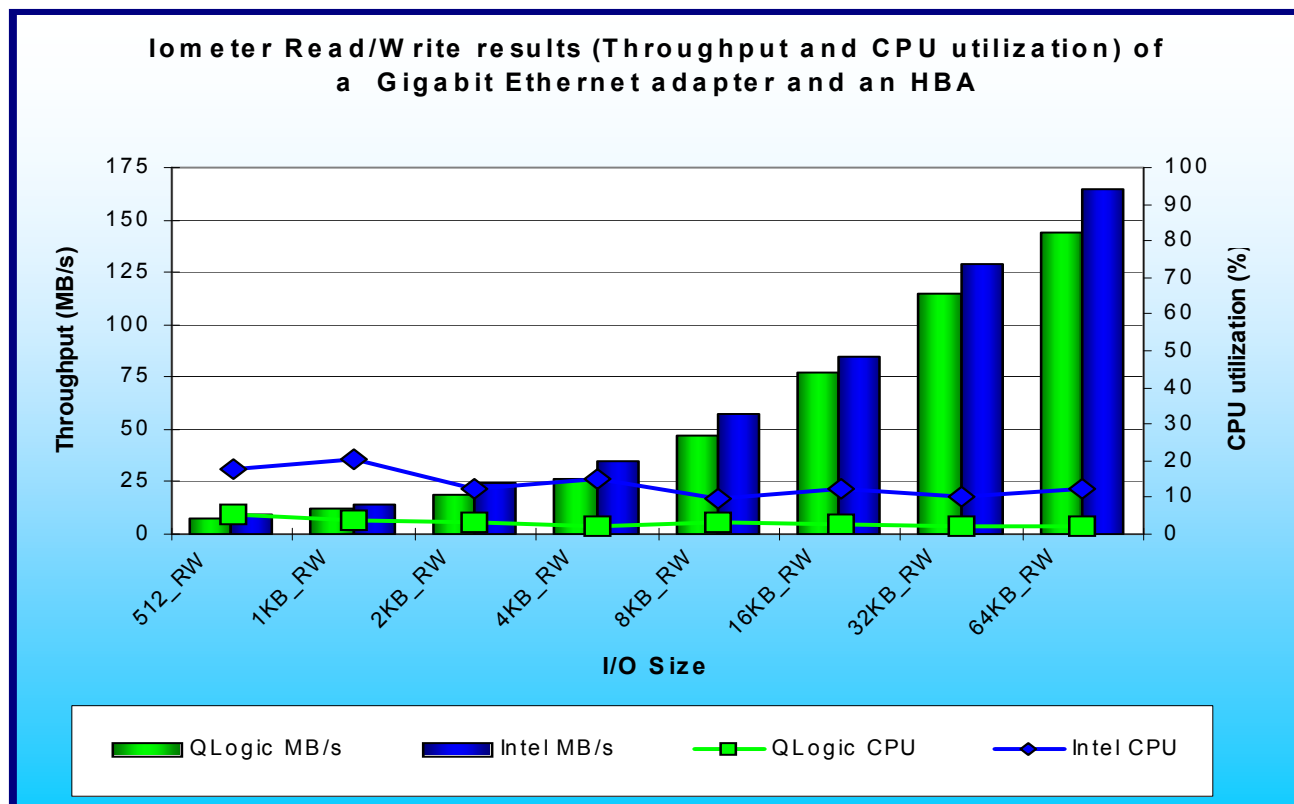


Figure 6: Iometer read/write throughput results, in MB/s, and the percentage CPU utilization. For throughput, higher numbers are better. For CPU utilization, lower numbers are better.

Test methodology

Intel configured and provided all the components in our test bed. Intel shipped the server with a fresh copy of Microsoft Windows 2003 Server, x64 Enterprise Edition Service Pack 1 installed and all the current updates from the Microsoft Windows Update site applied. The following tools were also on the server:

- Iometer v2004.07.30
- Microsoft iSCSI Initiator v2.03
- QLogic SANsurfer iSCSI HBA Manager v4.03.18

We used Microsoft iSCSI Initiator to configure the storage system for testing with the Intel PRO/1000 PT Dual Port Server Adapter. For the QLogic QLA4052C Dual Port iSCSI TOE to HBA, we used the QLogic SANsurfer iSCSI HBA Manager to control the storage system.

Setting up Microsoft iSCSI Initiator

We performed the following steps to configure and set up Microsoft iSCSI Initiator:

1. Launch the iSCSI Initiator Properties control panel by clicking the desktop shortcut.
2. Select the Discovery tab.
3. Click Add.
4. At the Add Target Portal screen, enter the IP address or DNS name and Port in the appropriate spaces.
 - For the IP address or DNS name, type 11.1.1.240
 - For the Port, type 3260
5. Click OK.
6. Click Add again, and enter the second IP address, 12.1.1.240, and the same Port as above.
7. Click OK.
8. Make sure that the Target Portals properties displayed on the Discovery tab are correct for both IP addresses.
9. Select the Targets tab. This will display a list of available targets that you can log on to. Prior to logon, the Status for each target should be Inactive.
10. Click the target you want to log on to, and click Log On.
11. At the Log On to Target dialog, check the box to Automatically restore this connection when the system boots, and uncheck the box to Enable multi-path.
12. Click OK.
13. Repeat steps 10 through 12 for the second target.
14. At the iSCSI Initiator Properties dialog, click OK to exit.

Setting up QLogic SANsurfer iSCSI HBA Manager

We performed the following steps to configure and set up QLogic SANsurfer iSCSI HBA Manager:

1. Launch the SANsurfer application by clicking the desktop shortcut.
2. Click connect to connect to the host system
3. Expand the HBA menu to show both adapters.
4. On the Target Settings tab for both adapters, click the plus symbol, and enter the IP address in the space provided.
 - For the IP address, type 11.1.1.240 for the first adapter.
5. Click OK.
6. Click Save Target Settings to save the changes.
7. Click Add again, and enter the second IP address, 12.1.1.240.
8. Click OK.
9. Click Save Target Settings to save the changes.
10. Select File→Exit to close the application.

Running the Iometer workloads

To run the Iometer tests, we did the following:

1. Reboot the server.
2. Login.
3. Wait 10 minutes after logging in to ensure a consistent starting state.
4. Open the Iometer.test mix file.
5. Under Topology, select server.
6. Under Disk Targets, select PHYSICALDRIVE:1 and 2.
7. Click the green flag to start.
8. Save the log file.

Test bed configuration

Figure 7 shows the test bed configuration we used in our tests.

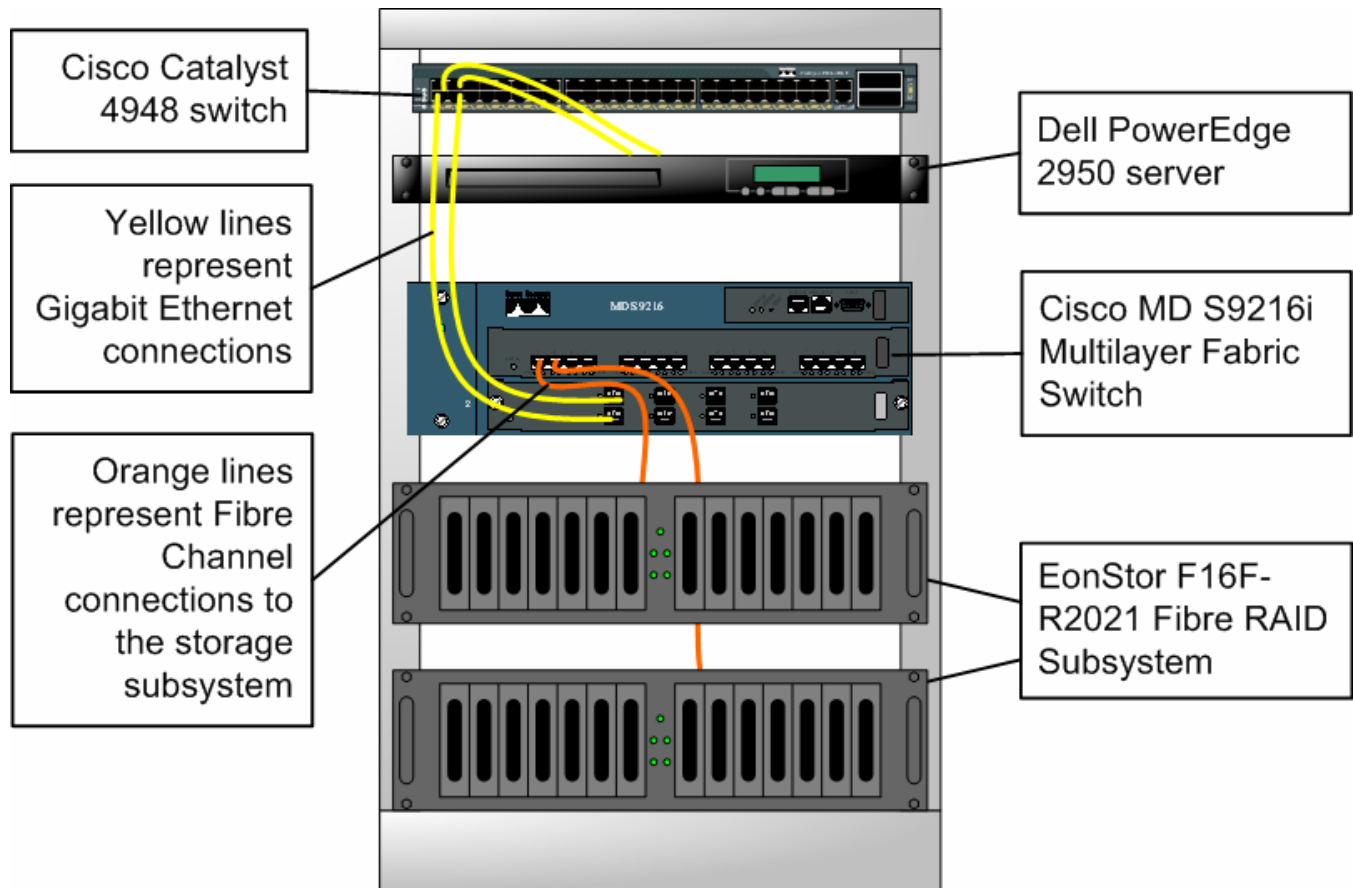


Figure 7: The test bed we used to run the Iometer workloads.

Appendix A – Test system configuration information

This appendix provides detailed configuration information about the test server system and the two iSCSI controllers. (We used the same server for both adapters during testing.)

Server	Dell PowerEdge 2950
System configuration information	
General	
Number of processor packages	2
Number of cores per processor package	2
Number of hardware threads per core	1
System Power Management Policy	Always on
CPU	
Vendor	Intel
Name	Dual-Core Intel Xeon processor 5160
Stepping	6
Socket type	LGA 771
Core frequency (GHz)	3.0 GHz
System bus	1333 MHz
L1 Cache	32 KB + 32 KB (per core)
L2 Cache	4 MB (shared by 2 cores)
Platform	
Vendor and model number	Dell PowerEdge 2950
Motherboard model number	0NH278
Motherboard chipset	Intel 5000X Chipset
Motherboard revision number	A03
Motherboard serial number	CN-0NH-278-13740-671-0107
BIOS name and version	Dell, Inc. 1.2.0, 10/18/2006
BIOS settings	Default
Chipset INF driver	Intel 1.2.66.0
Memory module(s)	
Vendor and model number	Infineon HYS72T128020HFN-3-7-A
Type	PC2-4200 FBDIMM
Speed (MHz)	533
Speed in the system currently running @ (MHz)	533
Timing/Latency (tCL-tRCD-iRP-tRASmin)	4-4-4-12
Size	4096 MB
Number of RAM modules	4
Chip organization	Double-sided
Hard disk	
Vendor and model number	Maxtor Atlas 15k II
Number of disks in system	1
Size	73 GB
Buffer Size	8 MB
RPM	15,000
Type	SAS
Controller	Dell SAS 5/i Integrated Controller
Controller driver	Dell 1.21.8.0
Operating system	
Name	Microsoft Windows 2003 Server, x64 Enterprise Edition

Build number	3790
Service Pack	Service Pack 1
Microsoft Windows update date	02/09/2007
File system	NTFS
Kernel	ACPI Multiprocessor x64-based PC
Language	English
Microsoft DirectX version	9.0c
Graphics	
Vendor and model number	ATI ES1000
Chipset	ATI ES1000 PCI
BIOS version	01.00
Type	Integrated
Memory size	16 MB
Resolution	1024 x 768
Driver	Microsoft 5.2.3790.1830
Optical drive	
Vendor and model number	TSSTcorp CDRW/DVD
Type	DVD-ROM
Interface	IDE
Dual/Single layer	Single
USB ports	
Number	4
Type	USB 2.0

Figure 8: Detailed system configuration information for the test system.

iSCSI controller	Intel PRO/1000 PT Dual Port Server Adapter	QLogic QLA4052C Dual Port iSCSI TOE to HBA
Network card/subsystem		
Vendor and model number	Intel PRO/1000 PT Dual Port Server Adapter	QLogic QLA4052C Dual Port iSCSI TOE to HBA
Type	PCI Express	PCI-X
Driver	Intel 9.6.31.0	QLogic 2.0.1.8

Figure 9: Detailed system configuration information for the test iSCSI controllers.



Principled Technologies, Inc.
1007 Slater Road, Suite 250
Durham, NC 27703
www.principledtechnologies.com
info@principledtechnologies.com

Principled Technologies is a registered trademark of Principled Technologies, Inc.
All other product names are the trademarks of their respective owners.

These results were obtained using Iometer version 2003.12.16, Copyright 2001 by Intel Corporation. Intel does not endorse any Iometer results.

Disclaimer of Warranties; Limitation of Liability:

PRINCIPLED TECHNOLOGIES, INC. HAS MADE REASONABLE EFFORTS TO ENSURE THE ACCURACY AND VALIDITY OF ITS TESTING, HOWEVER, PRINCIPLED TECHNOLOGIES, INC. SPECIFICALLY DISCLAIMS ANY WARRANTY, EXPRESSED OR IMPLIED, RELATING TO THE TEST RESULTS AND ANALYSIS, THEIR ACCURACY, COMPLETENESS OR QUALITY, INCLUDING ANY IMPLIED WARRANTY OF FITNESS FOR ANY PARTICULAR PURPOSE. ALL PERSONS OR ENTITIES RELYING ON THE RESULTS OF ANY TESTING DO SO AT THEIR OWN RISK, AND AGREE THAT PRINCIPLED TECHNOLOGIES, INC., ITS EMPLOYEES AND ITS SUBCONTRACTORS SHALL HAVE NO LIABILITY WHATSOEVER FROM ANY CLAIM OF LOSS OR DAMAGE ON ACCOUNT OF ANY ALLEGED ERROR OR DEFECT IN ANY TESTING PROCEDURE OR RESULT.

IN NO EVENT SHALL PRINCIPLED TECHNOLOGIES, INC. BE LIABLE FOR INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH ITS TESTING, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL PRINCIPLED TECHNOLOGIES, INC.'S LIABILITY, INCLUDING FOR DIRECT DAMAGES, EXCEED THE AMOUNTS PAID IN CONNECTION WITH PRINCIPLED TECHNOLOGIES, INC.'S TESTING. CUSTOMER'S SOLE AND EXCLUSIVE REMEDIES ARE AS SET FORTH HEREIN.