



Document Number	Revision	Page
PC4034501	A	1 of 26

---

---

# QLA1040 SCSI Host Adapter Performance Benchmark Report

QLogic Engineering



Document Number	Revision	Page
PC4034501	A	2 of 26

---

---

## TABLE OF CONTENTS

<b>1. Benchmark Goals.....</b>	<b>3</b>
<b>2. Competitive Comparison List .....</b>	<b>5</b>
<b>3. Test Environment .....</b>	<b>6</b>
<b>4. Performance Metrics .....</b>	<b>8</b>
<b>5. Test Description .....</b>	<b>10</b>
<b>6. Results .....</b>	<b>11</b>
6.1 Dual Processor Results .....	11
6.1.1 Sequential Read Operations .....	11
6.1.2 Sequential Write Operations .....	13
6.1.3 Random Read Operations .....	14
6.1.4 Random Write Operations .....	15
6.1.5 Throughput.....	16
6.2 Quad Processor Results .....	18
6.2.1 Sequential Reads (Quad).....	18
6.2.2 IOPS Related to CPU Utilization.....	20
6.2.3 Sequential Writes (Quad).....	21
6.2.4 Random Reads (Quad).....	22
6.2.5 Random Writes (Quad).....	23
6.2.6 Throughput.....	24
<b>Summary .....</b>	<b>26</b>



Document Number	Revision	Page
PC4034501	A	3 of 26

---

---

## 1. Benchmark Goals

In an effort to establish an accurate comparison performance analysis, we elected to run a proprietary program with the following characterizations and goals:

- 1) Evaluate the overall performance of the QLogic QLA1040 SCSI adapter in a controlled pre-determined environment as described in section 3 of this document.
- 2) Compare the performance of the Qlogic QLA1040 adapter against the current major competitors in today's market. These adapters were chosen to represent the leading competition as determined by performance claims, similar architecture and popularity.
- 3) Analyze the results to determine areas of strength or weakness. Additionally, we'll attempt to identify limitations and/or hidden value such as system overhead or affects of queue management.
- 4) Present meaningful information to the QLogic Sales and Marketing staff. Our intention is to run the performance tests and visually graph the results for immediate representation. In addition, we will provide an analysis as to our conclusions.

### Futures:

- 1). Investigation as to the validity of various performance tests that are now available. The marketplace constantly changes in respect to these programs that become available, each making claims to be leader in performance analysis.
- 2). Additional test criteria to determine saturation points of the various components as it relates to system response rates and throughput.



Document Number	Revision	Page
PC4034501	A	4 of 26

- 
- 3). Characterization of various disk manufactures and how they respond to SCSI bus utilization, drive cache partitioning and other external factors.
  - 4). Participate in an independent organization that is comprised of many participants that represent this industry segment. This is a non-profit group founded to standardize and promote I/O benchmarking, know as the Storage Performance Council.



Document Number	Revision	Page
PC4034501	A	5 of 26

---

---

## 2. Competitive Comparison List

The benchmarks were performed on the following controllers:

- QLogic QLA1040  
PCI to Ultra Wide Single Ended SCSI Host Bus  
Adapter  
Top Assembly: PC2010403-09 Rev. E  
Firmware: 6.04 BIOS
- Adaptec 2940UW  
PCI to Ultra Wide Single Ended SCSI Host Bus  
Adapter  
Top Assembly: 945300-01 Rev. D  
Firmware: 589217 Rev. D
- Symbios Logic SYM8751SP  
PCI to Ultra Wide Single Ended SCSI Host Bus  
Adapter  
Top Assembly: 348-0031366 Rev. D  
Firmware: V4.03 BIOS



Document Number	Revision	Page
PC4034501	A	6 of 26

---

---

### 3. Test Environment

---

---

#### Test Systems

Since single processor systems are not capable of generating the required number of I/Os to saturate our controller, only Dual and Quad processor systems were used. Details of the systems used are shown below.

	Motherboard	PCI Chipset	Processors	Ram
System #1	SuperMicro P6DNF	440 FX	Dual 200MHZ Pentium-Pro	32 MB
System #2	Intel 647195-317	450 GX	Quad 200MHZ Pentium-Pro	256MB

---

---

#### Operating Systems

All tests were run under the following operating system:

NT 4.0 Build 1381 Service Pack 2

---

---

#### Peripherals

For this pass, we elected to use the following hard disk drives for the tests.

Seagate	Interface	Speed	Firmware
ST31055W	Wide Ultra-SCSI	5400 RPM	0276 & 0528



Document Number	Revision	Page
PC4034501	A	7 of 26

---

---

## Drivers

The latest **released** drivers were used for all the products tested. Drivers for each of the products were obtained from bulletin boards or web sites. Each was the latest version that could be readily obtained from these sources. Driver versions for each product are listed below.

Product Under Test	Driver Version
QLogic QLA1040	QL10WNT.SYS Version 2.24
Adaptec 2940UW	AIC78XX.SYS Version 2.00
Symbios Logic 8751	SYMC8XXX.SYS version 4.01

---

---

## Configuration

The QLogic Host Adapter parameters used for the test cycle were set to the following values.

Fifo Threshold	3
Command Complete	1
Fast Memory Timing	1
Completion Posting Mode	1
Execution Throttle	32



Document Number	Revision	Page
PC4034501	A	8 of 26

---

---

## 4. Performance Metrics

Performance of a SCSI host adapter consists of several different characteristics; each of which are best measured with a specific configuration of test parameters. The characteristics that are important for performance are:

**Throughput** - The maximum sustained data transfer rate achievable. This test is performed by maximizing the transfer size per drive, and then increasing the number of disks until we reach the combined aggregate. We stack the commands so the system has a constant queue depth and request the largest block size. The large I/O size maximizes bursting on the busses involved and accomplishes two things. The first is to effectively mask the mechanical latencies of the disk drives to the point that they do not affect the test. The second is to eliminate, as much as possible, the overheads on the SCSI bus associated with switching context between commands frequently. The throughput result is a measurement of amount of raw data handled by the Host Adapter in its simplest form.

**I/O's per Second (IOPS)** - Number of transfers achievable per second. This measurement is a result of the command overhead contributed to each I/O by the system and host adapter. IOPS are measured by executing small block transfers (Typically 512 bytes to 8K bytes) to many disk drives simultaneously. In essence, we time stamp the ability of the system to generate the I/O request, issue the read command to the SCSI port, and wait until a completion status is posted, then start all over again. By queuing several commands to each disk, the mechanical latencies are again masked, but the overhead in the SCSI bus and the host adapter are now maximized. CPU availability is critical here as it requires substantial involvement to build each request. This will exemplify the ability of the controller to respond to system requests (by the number of I/O's processed) and also demonstrates the impact of the adapter as it relates to system overhead (CPU utilization). The resulting number equates to how many commands per second can be completed. Lower the command overhead results in a higher IOPS number.



Document Number	Revision	Page
PC4034501	A	9 of 26

---

**CPU Utilization-** The amount of CPU time required for processing each I/O. This is a measure of the CPU time required to execute each command to the drive. It represents the efficiency of the Software Driver/Host Adapter interface, and the intelligence on the host adapter. The environment in which this is measured is identical to that of the IOPS test. This measurement is normally expressed as a percentage of the total CPU power consumed during the test that was required to process I/O's. Lower percentages are an indication of the Host Adapter's ability to off-load the CPU. The advantage of having additional CPU resources will be readily apparent in the system's response to other processes waiting to be serviced.



Document Number	Revision	Page
PC4034501	A	10 of 26

---

---

## 5. Test Description

---

---

### General

The tests that were run fall into four specific types:

- 1) Sequential Reads
- 2) Sequential Writes
- 3) Random Reads
- 4) Random Writes

During each of these tests there are three parameters that can be changed.

- 1) Number of disk drives used.
- 2) Block size of each I/O (Amount of data transferred).
- 3) Queue Depth (Number of commands stacked to a single drive).

This creates a multitude of combinations that can be run and graphed. We have chosen to fix the Number of Drives and the Queue Depth parameters at values that produce the highest performance numbers, and graph Block Size against either IOPS, Throughput, or CPU Utilization. The number of drives for this operation was optimally configured for eight and we limited the application queue depth to four outstanding commands per drive.

Throughput can be derived by simply multiplying the IOPS by the block size. The throughput numbers will not be charted for all the IOPS tests shown in the following section. We elected to show a single throughput chart for reads and one for writes on sequential operations only.

---

---

### Test Program

The test program used for the benchmarks is a QLogic internal application called Bench32'. This program provides the controlling parameters required for each of the tests. The version of this program is V1.3



Document Number	Revision	Page
PC4034501	A	11 of 26

## 6. Results

### 6.1 Dual Processor Results

#### 6.1.1 Sequential Read Operations

In sequential read operations there is some lower performance noted for the Adaptec controller, but all three eventually merge with 16K block sizes. QLogic and Symbios maintain relative equal performance in IOPS. Adaptec caps at 7919 IOPS in this Dual CPU system environment. QLogic hits a maximum of 8901 and Symbios mirrors this result. As block size increases, those relative positions are held while the absolute numbers converge to near equality as the block size increases to 64K.

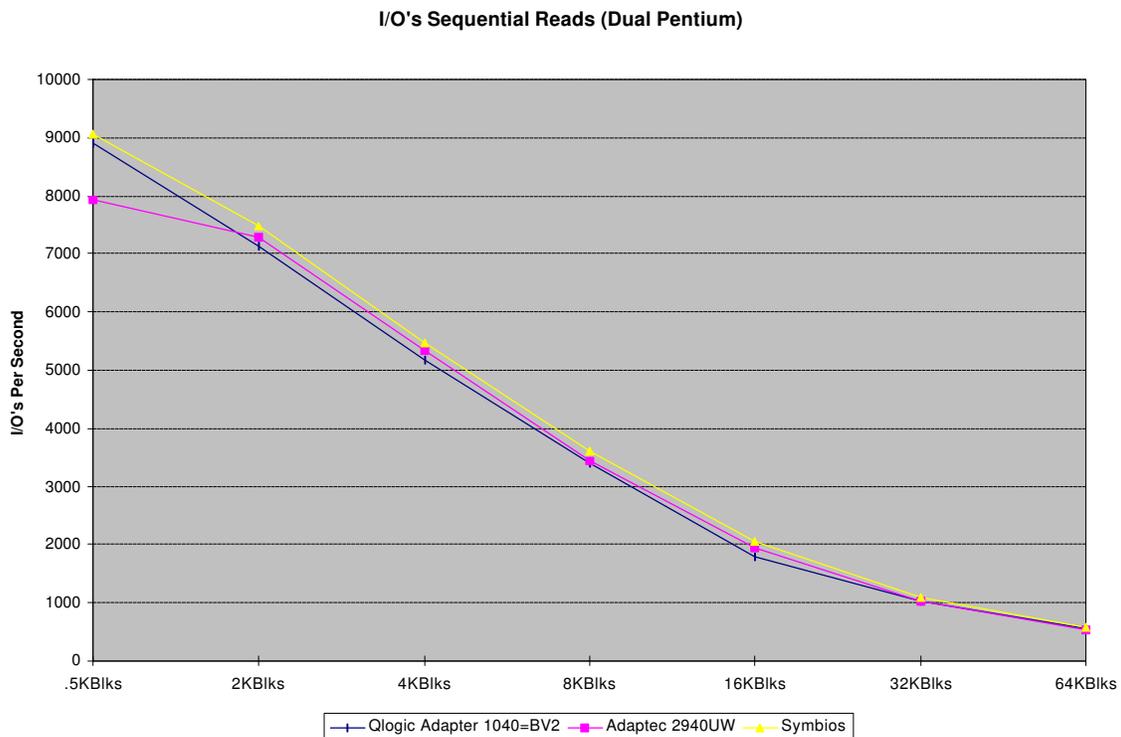


Figure 1



Document Number	Revision	Page
PC4034501	A	12 of 26

CPU Utilization Sequential Reads (Dual Pentium)

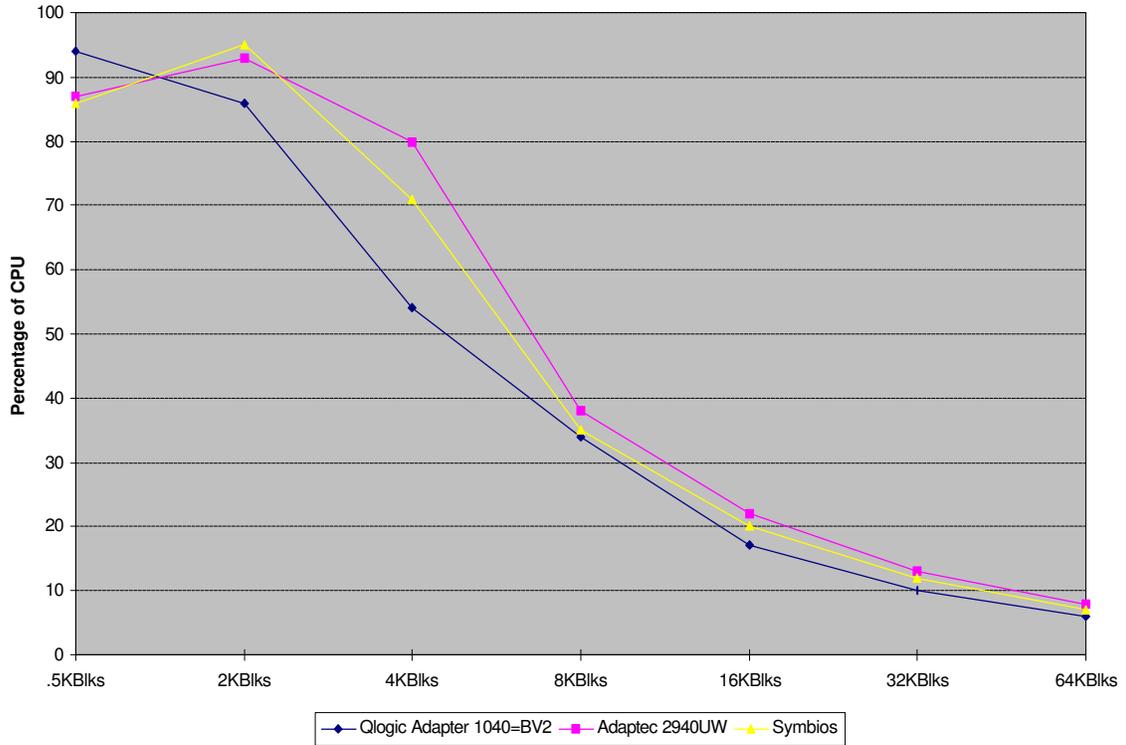


Figure 2



Document Number	Revision	Page
PC4034501	A	13 of 26

## 6.1.2 Sequential Write Operations

The Write tests showed QLogic with a clear consistent advantage in sequential writes across the range of block sizes. Differences in sequential write operations are more dramatic because of the number of disk drive commands the Host Adapter is able to queue for each drive. The Adaptec Driver only allowed a total of 23 commands spread across all drives, whereas the QLogic driver is only limited by the Nvram parameter. The value used on our tests was 4 commands per drive or 32 total on 8 drives. QLogic maintains a significant performance advantage in throughput at the smaller block sizes with 2721 I/O's compared to 2425 and 1765 respectively. This equates to a performance advantage of approximately 13% over Symbios and 55% over Adaptec.

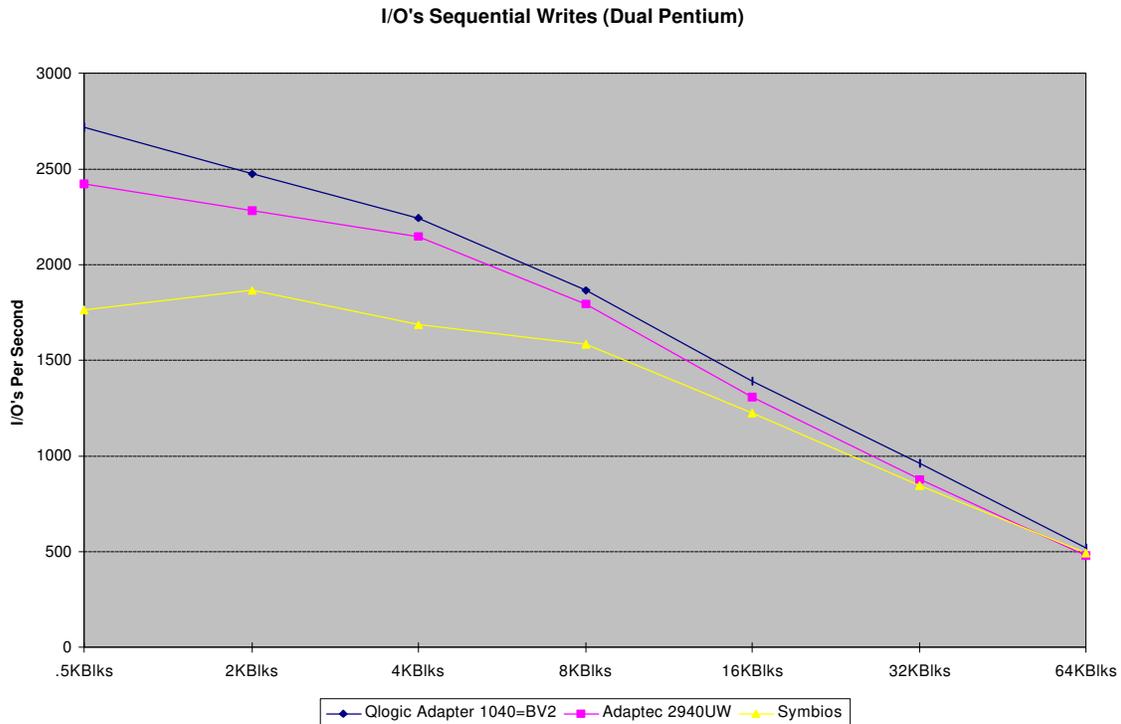


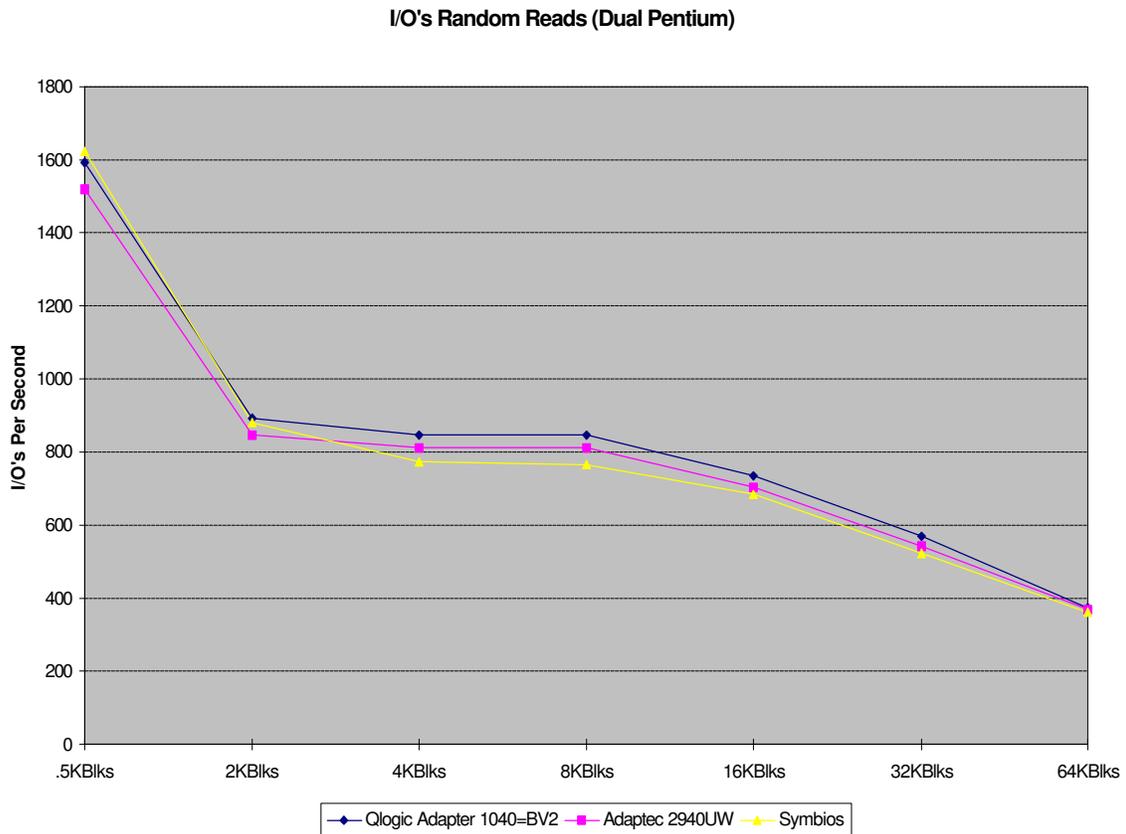
Figure 3



Document Number	Revision	Page
PC4034501	A	14 of 26

### 6.1.3 Random Read Operations

Because of mechanical latencies in the disk drives, this will limit random IOPS to relatively low numbers. The differences are really inconsequential across the block sizes. They become more important at the higher block sizes where throughput instead of I/O's per second becomes the issue. All three adapters mimic each other's pattern in the random read tests.



**Figure 4**



Document Number	Revision	Page
PC4034501	A	15 of 26

### 6.1.4 Random Write Operations

The pattern of the random writes are nearly identical to the random reads in that all the controllers are slaved to the drives and their capability. This is to be expected to a certain degree, as the limiting factor would be the mechanical latency of the drive. However, the Random Write tests showed QLogic with a clear consistent advantage in writes across the range of block sizes. The reason for this performance gain is the ability of our controller to queue the requested amount of commands as determined in our Nvram parameters. The competition's driver will place an internal driver limit on the number of requests queued.

I/O's Random Writes (Dual Pentium)

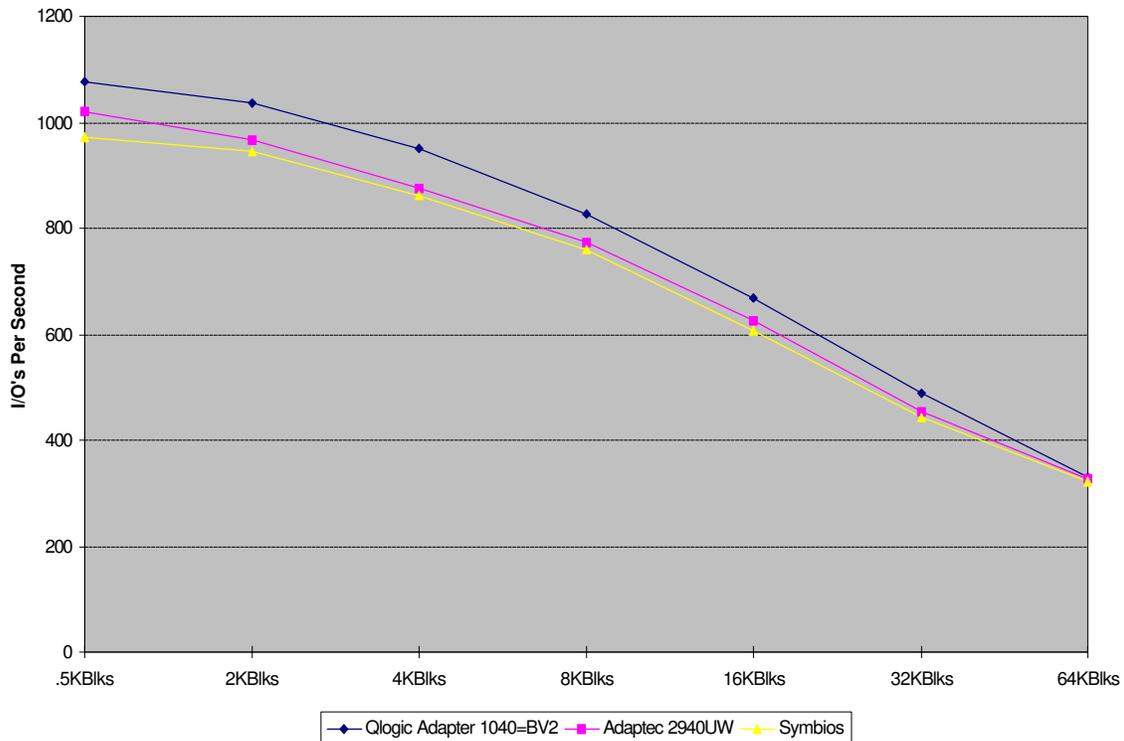


Figure 5



Document Number	Revision	Page
PC4034501	A	16 of 26

### 6.1.5 Throughput

Transfer rates are determined by the maximum amount of data the disk subsystem can deliver to the host system. This is imperative for applications that require large chunks of data to be moved. The objective is for the subsystem to transfer large blocks of data (64K) while limiting the impact of CPU overhead. The result is a host system that is more productive.

Transfer Rate Sequential Reads (Dual System)

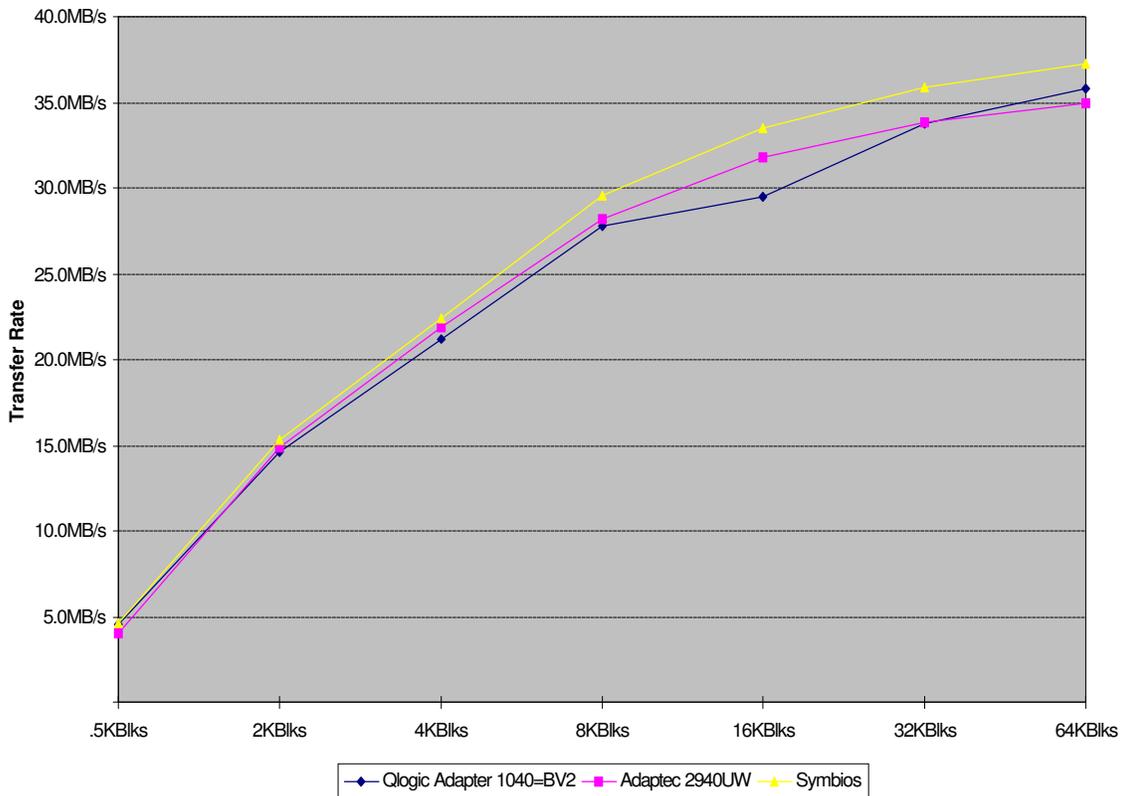


Figure 6



Document Number	Revision	Page
PC4034501	A	17 of 26

Transfer Rate Sequential Writes (Dual System)

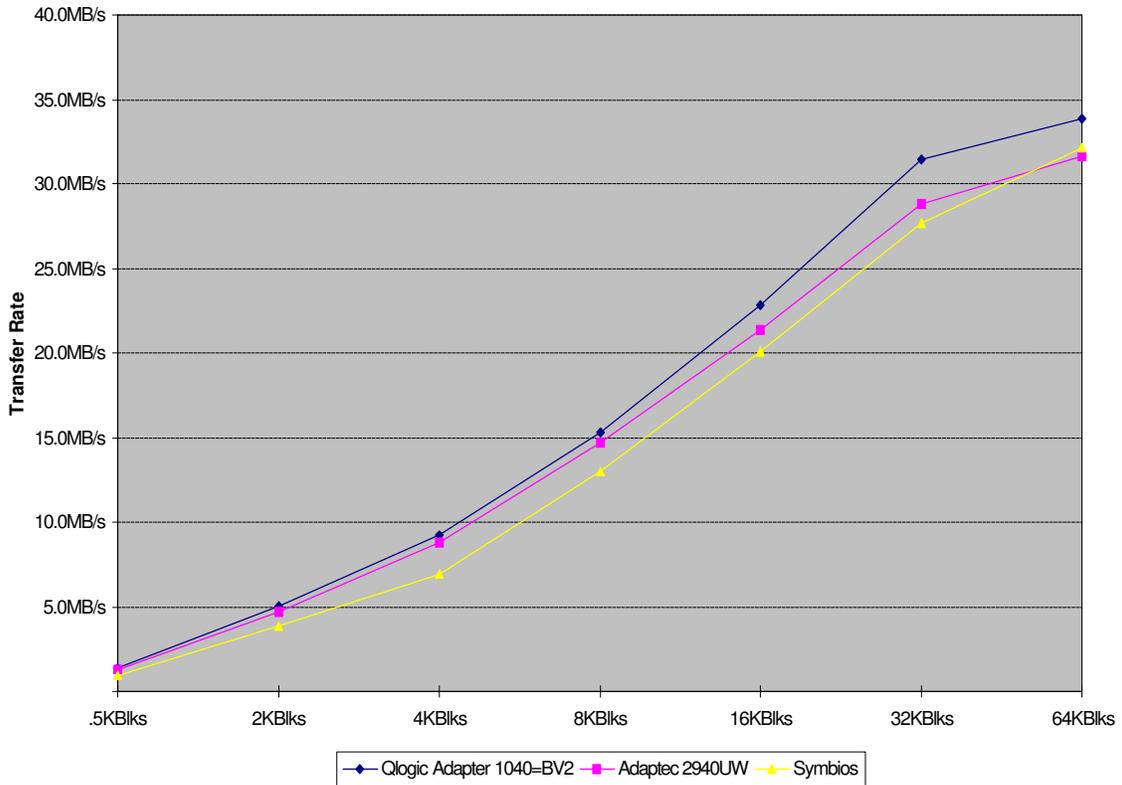


Figure 7



Document Number	Revision	Page
PC4034501	A	18 of 26

## 6.2 Quad Processor Results

### 6.2.1 Sequential Reads (Quad)

The increased performance is now demonstrated on the Adaptec card due to the higher CPU availability. This advantage is quickly eliminated about the 8K transfer size. The key here however, is CPU utilization. At all points on the curve, the QLogic adapter is using significantly less CPU bandwidth to generate the I/O's. For example, at the highest IOPS numbers, Adaptec is using 47% more of the CPU's time and generating only 14% more IOPS. The Symbios board is generating less IOPS with 36% more CPU utilization. This is an important reference point since it shows significantly higher CPU overhead to create a minimal increase in IOPS.

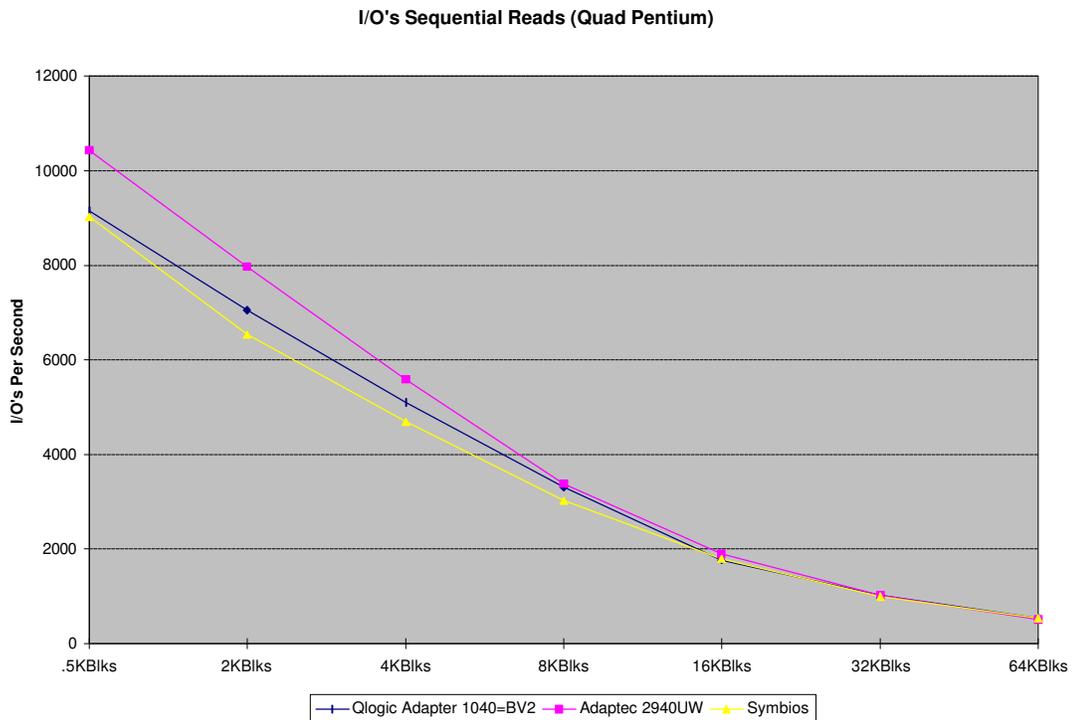


Figure 8



Document Number	Revision	Page
PC4034501	A	19 of 26

CPU Utilization Sequential Reads (Quad Pentium)

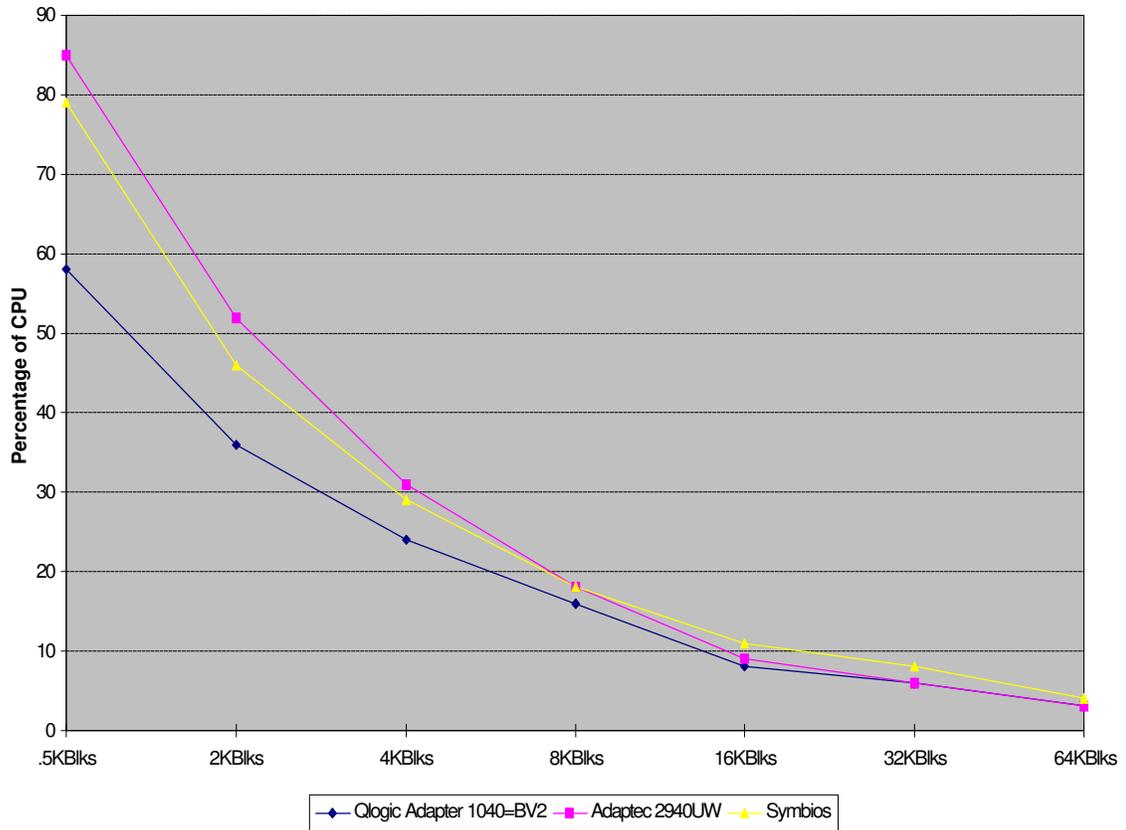


Figure 9



Document Number	Revision	Page
PC4034501	A	20 of 26

## 6.2.2 IOPS Related to CPU Utilization

The following chart will depict the number of I/O's per second generated as rated against the amount of CPU overhead. This is an overall summary of sequential reads and how it impacts the system. The QLogic adapter is using substantially less CPU time when compared to the competition. The example shows the amount of CPU time when calculated against the maximum number of I/O's each controller is capable of performing. This margin is most important when the system is transferring smaller block sizes. This additional CPU availability will be reflected in overall system performance since the processor is available to perform other tasks such as background computations or additional users. Notice the Adaptec controller using approximately 46% more of the CPU and only generates 14% more IOPS.

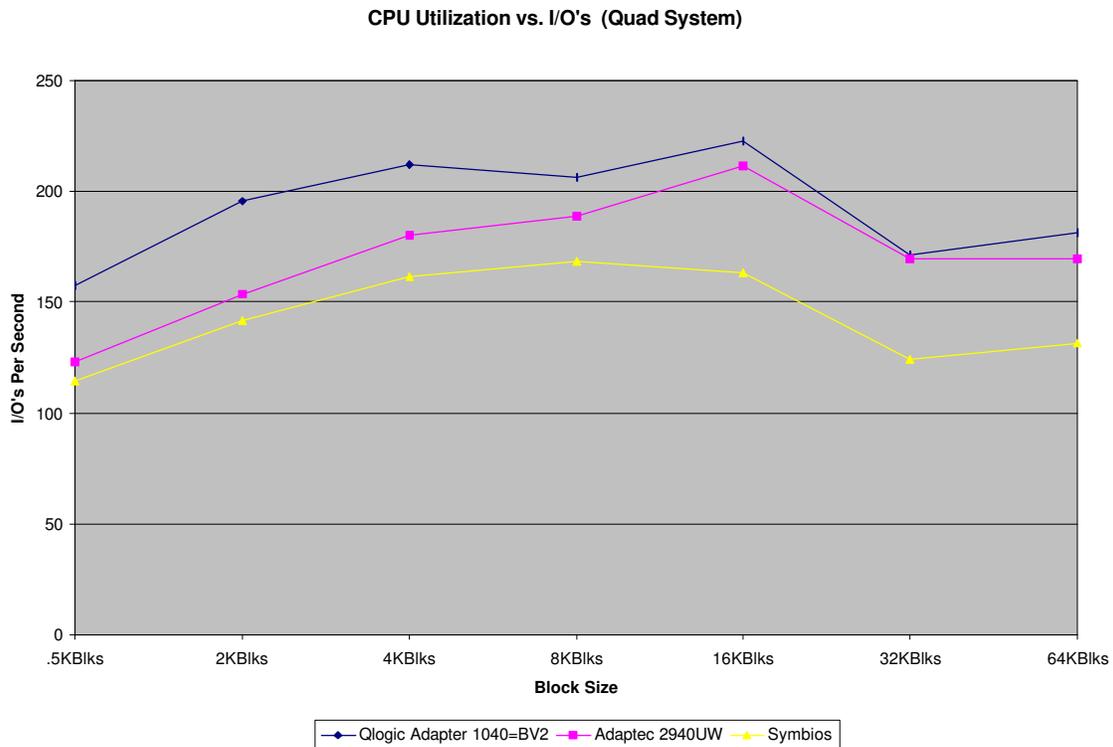


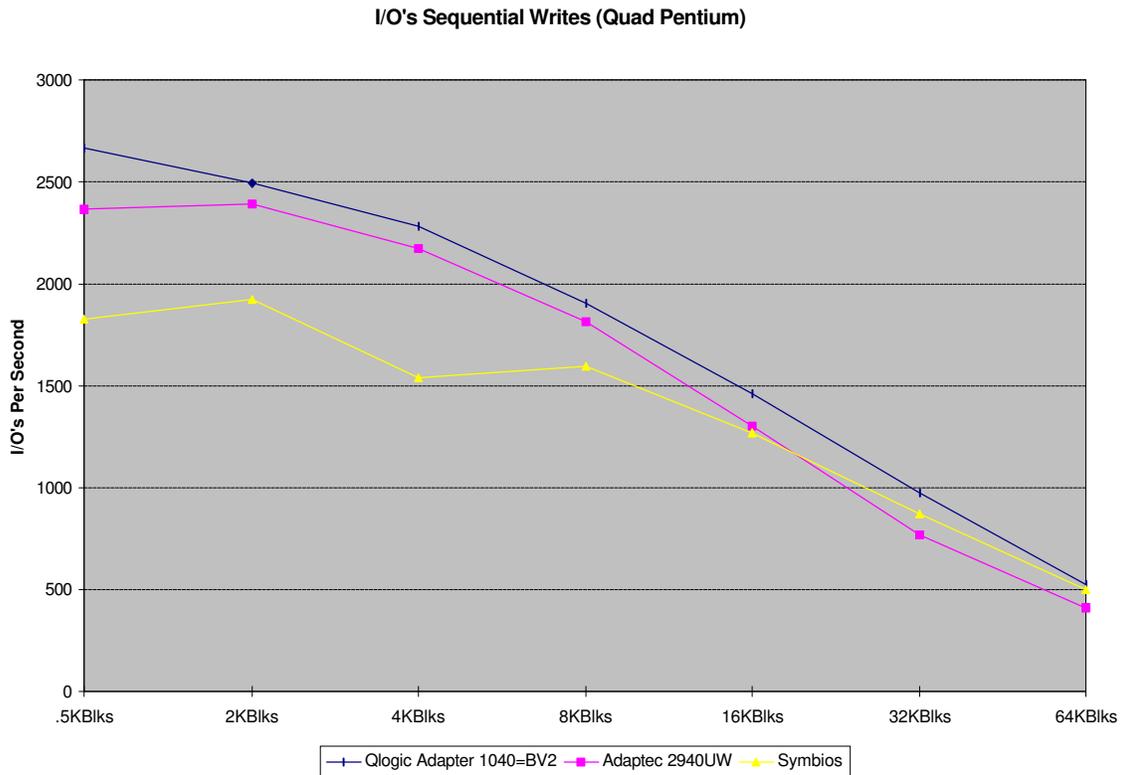
Figure 10



Document Number	Revision	Page
PC4034501	A	21 of 26

### 6.2.3 Sequential Writes (Quad)

The Write tests showed QLogic with a clear consistent advantage in sequential writes across the range of block sizes. Differences in sequential write operations are more dramatic because of the number of disk drive commands the Host Adapter is able to queue for each drive. The Adaptec Driver only allowed a total of 23 commands spread across all drives, whereas the QLogic driver is only limited by the Nvram parameter. The value used on our tests was 4 commands per drive or 32 total on 8 drives. QLogic maintains a significant performance advantage in throughput at the smaller block sizes with 2666 I/O's compared to 2367 and 1826 respectively. This equates to a performance advantage of approximately 13% over Adaptec and 49% over Symbios.



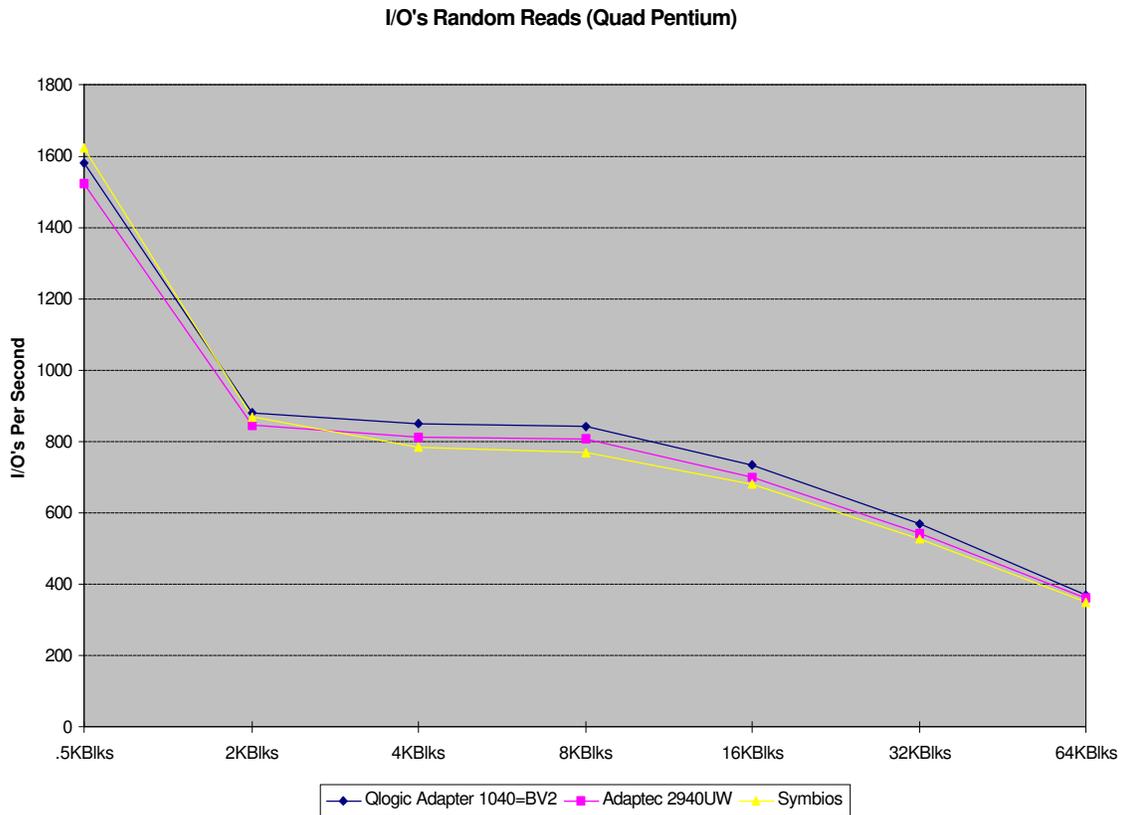
**Figure 11**



Document Number	Revision	Page
PC4034501	A	22 of 26

### 6.2.4 Random Reads (Quad)

Random Reads indicate the QLogic controller at a slight advantage starting at the edge of the 2K boundary. Because of mechanical latencies in the disk drives, this will limit random IOPS to relatively low numbers. The differences displayed are really inconsequential across the block sizes. All three adapters mimic each other's pattern in the random read tests.



**Figure 12**



Document Number	Revision	Page
PC4034501	A	23 of 26

### 6.2.5 Random Writes (Quad)

It's evident in the graph that Qlogic is the consistent leader in performance throughout the different block sizes. At the larger block sizes where throughput is the key factor Qlogic is approximately 28% faster than Symbios and 10% over Adaptec.

The pattern of the random writes are nearly identical to the random reads in that all the controllers are slaved to the drives and their capability. This is to be expected to a certain degree, as the limiting factor would be the mechanical latency of the drive. However, the Random Write tests showed QLogic with a clear consistent advantage in writes across the range of block sizes. The reason for this performance gain is the ability of our controller to queue the requested amount of commands as determined in our Nvram parameters. The competition's driver will place an internal driver limit on the number of requests queued

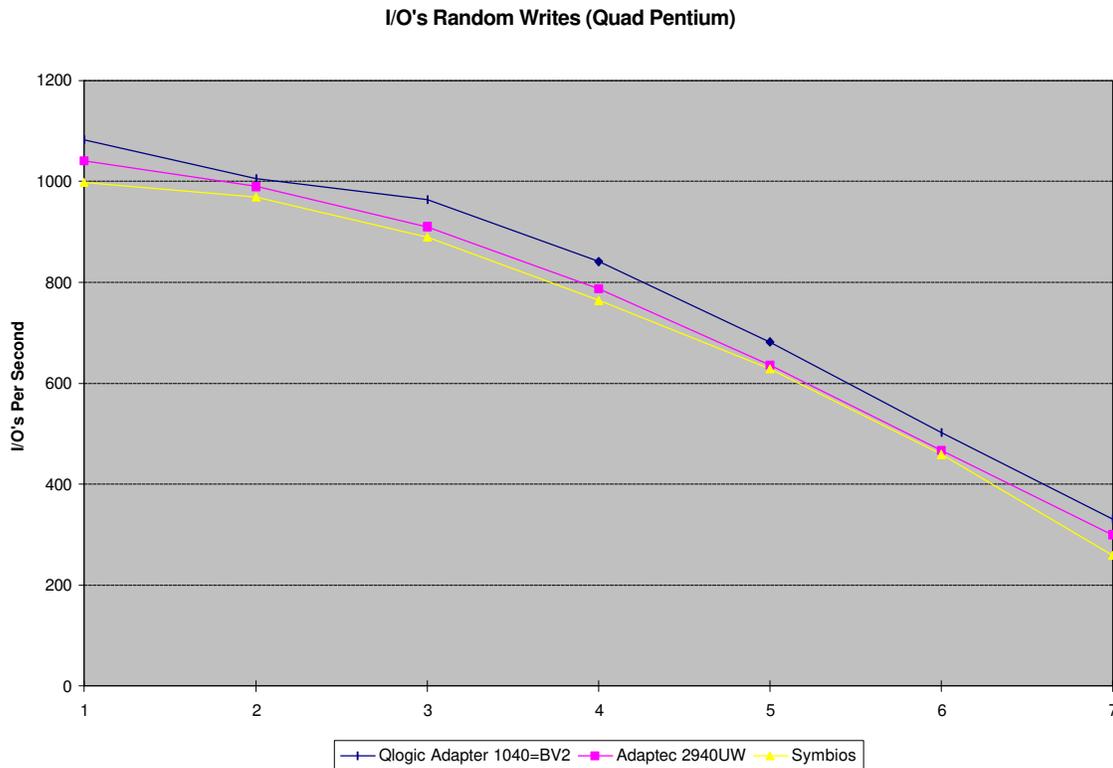


Figure 13



Document Number	Revision	Page
PC4034501	A	24 of 26

## 6.2.6 Throughput

Transfer rates are determined by the maximum amount of data the disk subsystem can deliver to the host system. This is imperative for applications that require large chunks of data to be moved. The objective is for the subsystem to transfer large blocks of data (64K) while limiting the impact of CPU overhead. The result is a host system that is more productive.

Transfer Rate Sequential Reads (Quad System)

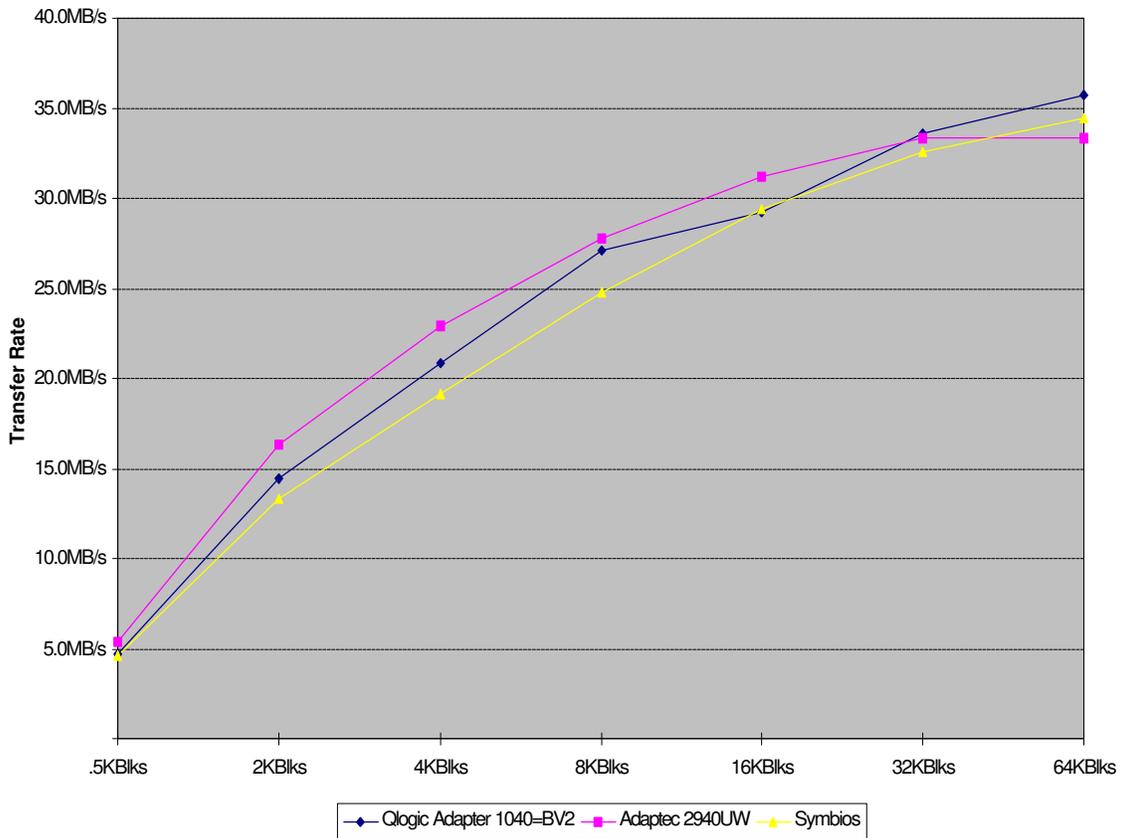


Figure 14



Document Number	Revision	Page
PC4034501	A	25 of 26

Transfer Rate Sequential Writes (Quad System)

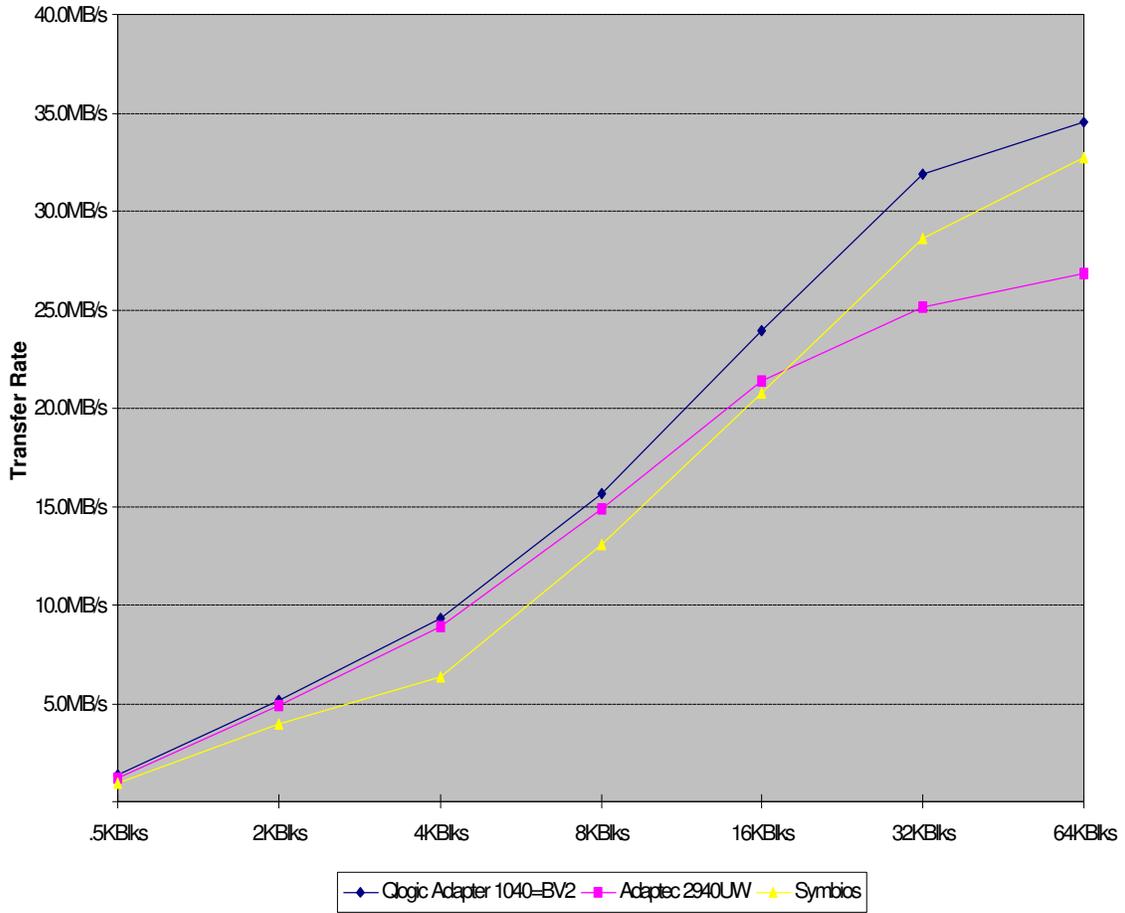


Figure 15



Document Number	Revision	Page
PC4034501	A	26 of 26

---

---

## Summary

Please note that we are still in the process of testing, compiling, and analyzing additional data. This report will be updated periodically to reflect the new data as it becomes available. It is also important to remember that all of the manufactures are constantly updating drivers and BIOS. We have run this series of tests many times due to driver changes posted on WEB pages. The results are much different every time a change is incorporated. We have found performance shifts as much as 50% in some cases. It is important to note that these tests were executed with the main objective to compare several different adapters in the same environment. This was a controlled test with nearly identical parameters set for each device (as close as the physical limitations would allow). These results are not representative of an actual operating system or customer specific environment. Typical operating systems under normal user loads (referring to servers, not desktop systems) will be performing random reads and writes throughout the disk. Many systems will favor certain areas of the disk due to directory structures and file placement. Performance is relative to numerous external factors both controllable and non-controllable. Other comparisons that need to be considered are items such as driver interoperability and user features. Our next series will attempt to incorporate a reference chart of features. Again, this is difficult since the BIOS changes so rapidly. At this time, all three vendors offer the same basic support for items such as CD-ROM boot, selectable drive boot and other user options. The QLogic board offers increased flexibility in configuring specific system and drive variables.

The results show that all 3 adapters that we examined were relatively close in performance on the sequential read and random read tests regardless of the block sizes. The distinguishing factor for these tests was the CPU utilization. It demonstrates dramatically less CPU overhead with the Qlogic adapter at the smaller block sizes. The sequential write and random write tests showed a clear advantage with the Qlogic adapter over the Adaptec and Symbios cards. This was attributed to the QLogic driver's ability to queue the optimum number of outstanding commands. If we increase the queue depth, the difference is even more dramatic. Further studies are planned to chart this out.

Lastly, we have found specific drive settings can greatly impact the test results. For instance, by changing the cache partition table in the drive or setting various mode select parameters, it can result in substantial gains or degradation. In the future, we plan to incorporate a series of tests to help us identify these specific settings so we can optimize the drive performance.