



Comparison of the Performance of the DS4300 (FAStT600) and FAStT200 Storage Servers

Charles Stephan

IBM eServer xSeries Performance Development & Analysis
Research Triangle Park, North Carolina, USA

Abstract

This white paper has been updated to reflect the renaming of the FAStT family of products as of September 7, 2004. The only changes to this paper are the affected product names (e.g., FAStT becomes DS4000, FAStT600 becomes DS4300, FAStT EXP700 becomes DS4000 EXP700). The FAStT200 Storage Server was withdrawn before the introduction of the new product names.

This paper provides the baseline performance of the IBM® TotalStorage® DS4300 (FAStT600) Storage Server, which employs IBM DS4000 storage technology (formerly Fibre Array Storage Technology, i.e., FAStT).

The purpose of this paper is to present the results obtained using the Iometer tool to measure the raw performance of the DS4300 Storage Server's RAID subsystem. The DS4300's performance is compared to that of its predecessor, the FAStT200.¹

The paper is organized in four sections. The first section briefly describes the tool used to measure the performance of the DS4300 and FAStT200, and defines the workloads used in the measurements. The second section describes the hardware and software measurement environment. The third section presents the results of the measurements and explains how the results should be interpreted. Finally, the fourth section summarizes the performance gains demonstrated by the DS4300.

Important lessons learned from this performance study are highlighted in boxes at appropriate points throughout the paper.

Questions about the information presented should be directed to the author at stephanc@us.ibm.com or Charles T Stephan/Raleigh/IBM@IBMU.S.

¹ The measurement results in this paper represent data that was written to disks or read from disks. The results do not represent data that was read strictly from RAID controller cache or written strictly to RAID controller cache. While both methods produce valid data, the "out-of-cache" or "to-cache" measurements do not fit within the scope of this document.

Measurement Tool and Workloads

Iometer is a workload generator and a measurement tool originally developed by the Intel Corporation. It is now maintained under an Intel® Open Source License, and it is available at <http://sourceforge.net>.

Iometer is designed to generate workloads and record measurement results for server disk and network subsystems—*not* desktop disk and network subsystems. In this context, the use of the words “server” and “desktop” is not a trivial matter. Consider the following example.

The single-threaded utility *copy* is routinely used to test server disk subsystems. The *copy* utility is a fine benchmark for a laptop or desktop machine, but not for a server. Why is it used so often for measuring server disk subsystem performance? It is probably used for two reasons. First, *copy* is easy to execute, and does not require large amounts of resources. The second reason is that the differences between server architecture and desktop architecture may not have been understood by the people implementing the benchmark.

Desktop machines are designed to manage one task at a time, and they do this very well. In fact, when *copy* is executed, a desktop machine with a single hard drive will usually perform better than a server with an array of multiple drives. The reason for the performance disparity is based on the design differences of the two machines. Servers are designed to handle multiple tasks in parallel. So, when *copy* is executed on a server, some server operating systems will bounce the *copy* process from CPU to CPU, because it is designed to keep all of the CPUs busy (Microsoft® Windows® 2003 will no longer do this). This is very costly with regard to performance. Furthermore, since *copy* is single-threaded, each I/O request must be satisfied before another I/O request can be generated. Therefore, the multiple-drive array is not being utilized efficiently, because only one drive is required to satisfy each I/O request.

One way to measure the performance of a server disk subsystem is to use Iometer. Iometer, by default, provides “workers” for each CPU in the system. This satisfies the need to keep all CPUs busy, and thus, multiple I/O requests can be issued in parallel so that all of the drives in an array can be kept busy just as it is done by a high-performance SMP server application. Iometer also provides a configurable parameter, called “outstanding I/Os,” which can be used to increase the load on a server disk subsystem. The measurement results contained in this paper were generated by increasing the number of outstanding I/Os queued at the drives up to and beyond what would be typical in a production environment.

Do not use desktop-oriented tools or single-threaded utilities, such as *copy*, to measure a server’s disk subsystem performance. Iometer is specifically designed to generate workloads on servers that utilize all of the CPUs in parallel, which ensures that I/O requests are issued in parallel to the disk subsystem.

The measurement results in this paper were obtained using Iometer version 2003.02.15, Copyright 1996-1999 Intel Corporation.

The workloads used to yield the results in this document were the On-Line Transaction Processing workload, Streaming Reads workload, Streaming Writes workload, File Server workload, Web Server workload, Random Reads workload, and the Random Writes workload. The characteristics for each workload are described in the following sections.

On-Line Transaction Processing Workload

The On-Line Transaction Processing (OLTP) workload is designed to emulate a transactional database workload. It is defined as 100% random accesses, 67% reads, and 33% writes. This workload is measured using transfer request sizes of 4K, 8K, 16K, 32K, and 64K.

Streaming Reads Workload

The Streaming Reads workload is designed to emulate a read-intensive multimedia streaming application. It is defined as 100% sequential accesses and 100% reads. This workload is measured using transfer request sizes of 512 bytes, 1K, 2K, 4K, 8K, 16K, 32K, 64K, 128K, 256K, and 512K.

Streaming Writes Workload

The Streaming Writes workload is designed to emulate a write-intensive multimedia streaming application. It is defined as 100% sequential accesses and 100% writes. This workload is measured using transfer request sizes of 512 byte, 1K, 2K, 4K, 8K, 16K, 32K, 64K, 128K, 256K, and 512K.

File Server Workload

The File Server workload consists of a mixture of various transfer request sizes. It is defined as 100% random accesses, 80% reads, and 20% writes. The mixture of transfer request sizes is defined as:

- 10% 512 Byte
- 5% 1K
- 5% 2K
- 60% 4K
- 2% 8K
- 4% 16K
- 4% 32K
- 10% 64K

Web Server Workload

The Web Server workload is designed to emulate a Web server delivering static content. It is defined as 100% random accesses and 100% reads. This workload consists of a mixture of transfer request sizes that is defined as:

- 22% 512 Byte
- 15% 1K
- 8% 2K
- 23% 4K
- 15% 8K
- 2% 16K
- 6% 32K
- 7% 64K
- 1% 128K
- 1% 512K

Comparison of the Performance of the DS4300 (FAStT600) and FAStT200 Storage Servers

Random Reads Workload

The Random Reads workload is defined as 100% random accesses and 100% reads. This workload is measured using transfer request sizes of 4K and 8K.

Random Writes Workload

The Random Writes workload is defined as 100% random accesses and 100% writes. This workload is measured using transfer request sizes of 4K and 8K.

Measurement Environment

The measurements were conducted using the IBM® eServer® xSeries® 345 with two Intel Pentium® 4 3.0GHz Xeon™ processors and 512MB of system memory. The operating system installed was Microsoft® Windows® 2000 Advanced Server 5.0.2195 (Build 2195) with Service Pack 4.

The x345 contained two IBM TotalStorage DS4000 FC-2-133 Host Bus Adapters using driver version 8.1.5.62.

The Fibre Channel management software installed on the x345 was Redundant Disk Array Controller (RDAC) version 08.30.95.03 and Storage Manager Client version 08.33.G5.03 (unless otherwise noted).

The Fibre Channel RAID controller software used on the DS4300 RAID controllers was Appware version 05.33.07.00, Bootware version 05.33.07.00, and NVSRAM version N1722F600R833V02 (unless noted otherwise on a table or chart). Each DS4300 RAID controller contained 128MB of cache.

The storage backend consisted of six IBM TotalStorage DS4000 EXP700 Storage Expansion Enclosures with ESM firmware version 9319. The DS4000 EXP700 enclosures contained 18.2GB 15K rpm drives with firmware version B947.



Comparison of the Performance of the DS4300 (FAStT600) and FAStT200 Storage Servers

Configuration Diagram for the Measured Hardware

xSeries 345 loaded with Windows 2000 Advanced Server



Dual DS4300s (128MB of cache per controller)



DS4000 EXP700s
containing 18GB15 rpm drives

Measurement Results and Analysis

The performance information contained in this section was derived under specific operating and environmental conditions. The results obtained in your operating environments may vary significantly.

The measurement results in this section represent the maximum sustainable performance for a configuration using two DS4300 RAID controllers with the number of hard disk drives (HDDs) utilized that correspond to an average response time of approximately 20 milliseconds (ms). The results in this section may not correspond to what is commonly referred to as the "peak" performance. Peak performance typically refers to a measurement result with the highest number of I/Os or MBps regardless of the average response time associated with that result.

RAID-5 OLTP Workload Results

Table 1 contains RAID-5 measurement results for the OLTP workload for various transfer request sizes. The data in Table 1 corresponds to the results achieved at an average response time of 20 ms. The drives were configured in arrays of 10, and only 8% of the total capacity of the drives was used. This is true for all of the measurements unless otherwise noted. Finally, a 64K segment was configured for the arrays and the cache block size was 4K. Cache mirroring was disabled.

Table 1. RAID-5 Dual DS4300s OLTP IOps*


Workload		OLTP 4K		OLTP 8K		OLTP 16K		OLTP 32K		OLTP 64K	
RAID-5		IOps	MBps	IOps	MBps	IOps	MBps	IOps	MBps	IOps	MBps
20 HDDs	All Caches Disabled	2750	11	2650	21	2420	38.5	1895	60.5	1320	84.5
	Write Cache Enabled	3480	14	3450	27.5	3050	49	2225	71	1600	102
40 HDDs	All Caches Disabled	5395	21.5	5100	41	4380	70	3250	104	2000	128
	Write Cache Enabled	6990	28	6740	54	4820	77	3410	109	2160	138
60 HDDs	All Caches Disabled	7460	30	6400	51	4820	77	3400	109	2070	132
	Write Cache Enabled	8080	32	6500	52	4775	76.5	3390	108	2160	138
80 HDDs	All Caches Disabled	8050	32	6650	53	4920	78.5	3375	108	2100	134
	Write Cache Enabled	8170	32.5	6460	51.5	4875	78	3350	107	2125	136

*Results correspond to an average response time of 20 ms.

Comparison of the Performance of the DS4300 (FAST600) and FAST200 Storage Servers

A drive stroke of 8% does not reflect typical capacity usage in a production environment. Think of the “disk stroke percentage” as the percentage of disk capacity required for a user’s working data set. The phrases “working data set” and “disk stroke percentage” are used interchangeably. As the working data set grows, so does the time for seek operations. One can expect production workloads to have longer seek times because of both capacity utilization and disk fragmentation. Chart 1 illustrates the effect that the drive stroke percentage plays on performance with respect to the number of I/Os an array of 10 drives can sustain for the RAID-5 8K OLTP workload. The results in Chart 1 correspond to a constant queue depth of 15 outstanding I/Os at the array.

Chart 1. RAID-5 OLTP 8K I/Os vs. Drive Stroke Percentage



Results in Chart 1 are meant only to show that the performance of the drives decreases as more of the storage capacity is used in workloads that are characterized by a large percentage of random accesses. For example, if a user has 40 drives attached to dual DS4300 controllers, and the 40 drives are dedicated to transaction processing with an average transfer request size of 8K, then the performance the user could expect would depend on the size of the working data set and the capacity utilization. The maximum performance the user could expect would be 5100 I/Os according to Table 1, assuming all caches were disabled. But that performance is based on a disk stroke percentage of 8%. Assuming the user’s working data set spans approximately 45% of the capacity of the drives, a more realistic estimate of the performance expected would be approximately 3775 I/Os.

This result was obtained by calculating the performance gained when the disk stroke percentage is reduced from 45% to 8% of the total data capacity. The gain is determined using the equation $[(1050 \text{ I/Os} - 765 \text{ I/Os}) / 765 \text{ I/Os}] = 35\%$. Next, the actual performance the user could expect, X, is determined using the equation $[5100 \text{ I/Os} = X(0.35) + X]$. Solving for X yields approximately 3775 I/Os. Remember, the goal is to obtain approximate performance and not precise performance carried out to numerous significant figures.

Comparison of the Performance of the DS4300 (FAStT600) and FAStT200 Storage Servers

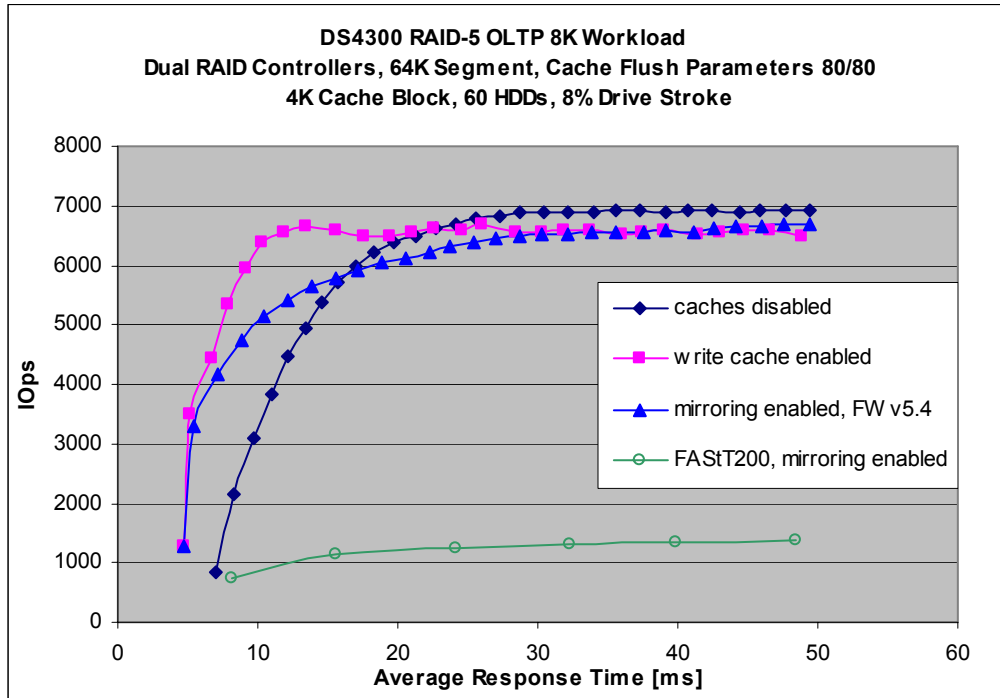
The effect on performance illustrated by Chart 1 does not hold true for workloads that access data sequentially. For sequential workloads, the performance of the drives does not decrease significantly, if at all, as a result of utilizing more of the capacity of the drives.

Hard drive performance decreases as the working data set spans more of the capacity of the drives while executing a workload characterized by a large percentage of random data accesses.

Hard drive performance does not decrease significantly, if at all, as the working data set spans more of the capacity of the drives while executing a workload characterized by a large percentage of sequential data accesses.

Chart 2 illustrates the RAID-5 OLTP 8K workload performance of dual DS4300s and dual FAStT200s.

Chart 2. DS4300 RAID-5 OLTP 8K vs. FAStT200 RAID-5 OLTP 8K



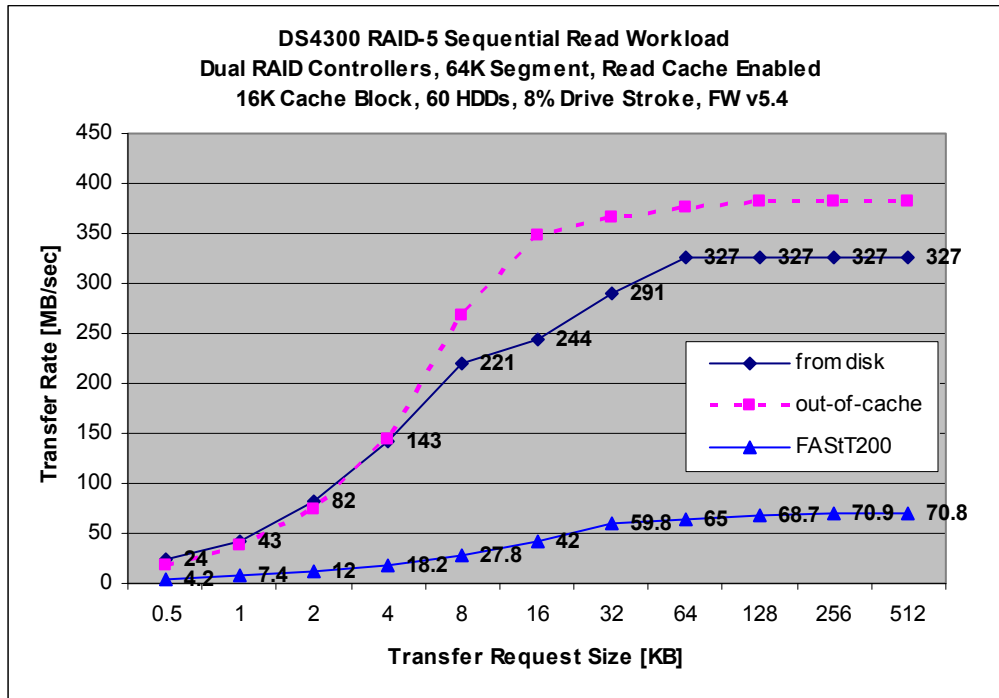
Comparing the “mirroring enabled” and “FAStT200” curves in Chart 2, and sampling the data that corresponds to a 20 ms average response time reveals that the DS4300 performs approximately 400% better than the FAStT200. However, the FAStT200 measurements were conducted using 9GB 10K rpm drives. Would 18.2GB 15K rpm drives improve the performance? The performance might improve by about 5% to 10% because the FAStT200 controllers were the bottleneck. The slight gain would come from the reduced latencies of the new drive technology. So, correcting for the slight gain, the DS4300 performs approximately 355% better than the FAStT200 for the RAID-5 OLTP 8K workload.

RAID-5 Streaming Reads and Streaming Writes Workload Results

Chart 3 illustrates the RAID-5 Streaming Reads performance of dual DS4300 controllers. Be advised that these results may not correspond to “peak” performance. All of the transfer rates correspond to average response times no greater than 20 ms.

Note that due to the large number of measurements conducted, the segment size used for both streaming workloads was 64K, which is the segment size used for the vast majority of the measurements. Using a larger segment size for both streaming reads and writes, and adjusting the read-ahead multiplier for streaming reads may result in better performance.

Chart 3. RAID-5 Dual DS4300s Streaming Reads Transfer Rate

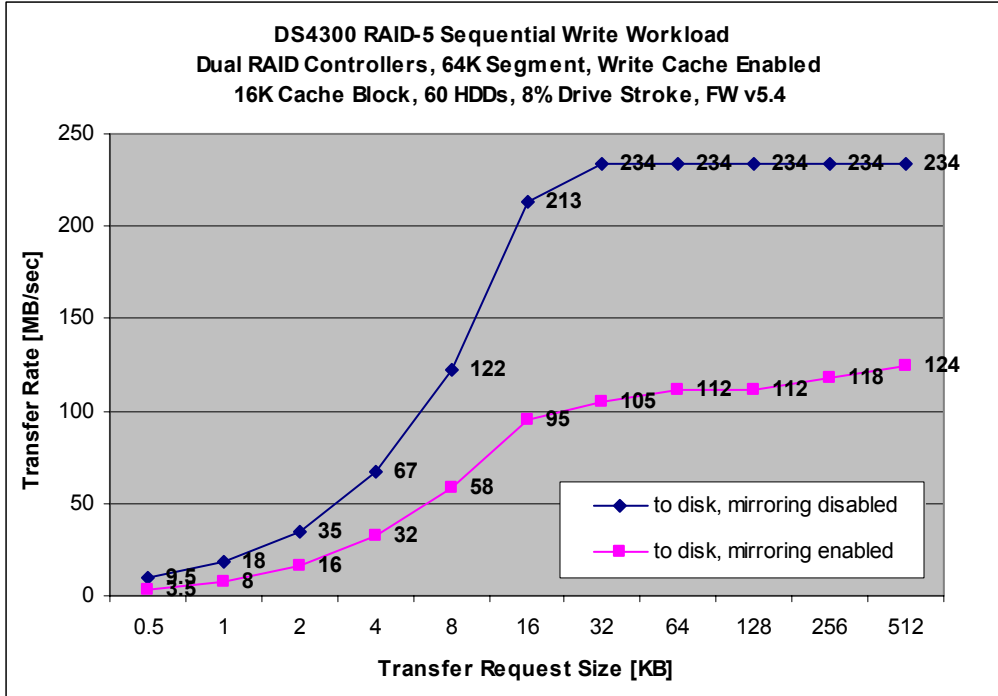


The transfer rate for large transfer blocks, greater than or equal to 64K, is approximately 70MBps for the FAST200. After compensating for older drive technology by boosting the FAST200 transfer rate by 10%, the DS4300 performs approximately 300% better than the FAST200 for large block transfers for the Streaming Reads workload. It must be noted that the FAST200 would not have been able to achieve greater than approximately 190MBps because it would have been limited by the two 1Gbit Fibre links it utilized. In fact, the FAST200 could not even saturate the 1Gbit links.

Comparison of the Performance of the DS4300 (FAST600) and FAST200 Storage Servers

Chart 4 illustrates the Streaming Writes performance of dual DS4300 controllers. Be advised that these results may not correspond to “peak” performance. All of the transfer rates correspond to an average response time no greater than 20 ms.

Chart 4. RAID-5 Dual DS4300s Streaming Writes Transfer Rate



RAID-5 File Server and Static Web Server Workload Results

Table 2 contains the results of the File Server and the Static Web Server workload measurements. The results should not be interpreted as, for example, “This is what a user would see if the user attached dual DS4300 controllers to a machine that is designated to be a file server or static Web server.” There are too many other factors involved with that line of thinking that might influence performance as well, such as the number of active users accessing files, the type of machine being used as the file server or Web server, and the general characteristics of the production workload, to name a few.

The File Server and Static Web Server workloads are used primarily to track product performance, serve as a comparison between similar products, and exercise a product with a workload that consists of various transfer request sizes as opposed to a uniform transfer request size.

Table 2. RAID-5 Dual DS4300s File Server and Static Web Server Transfer Rate*

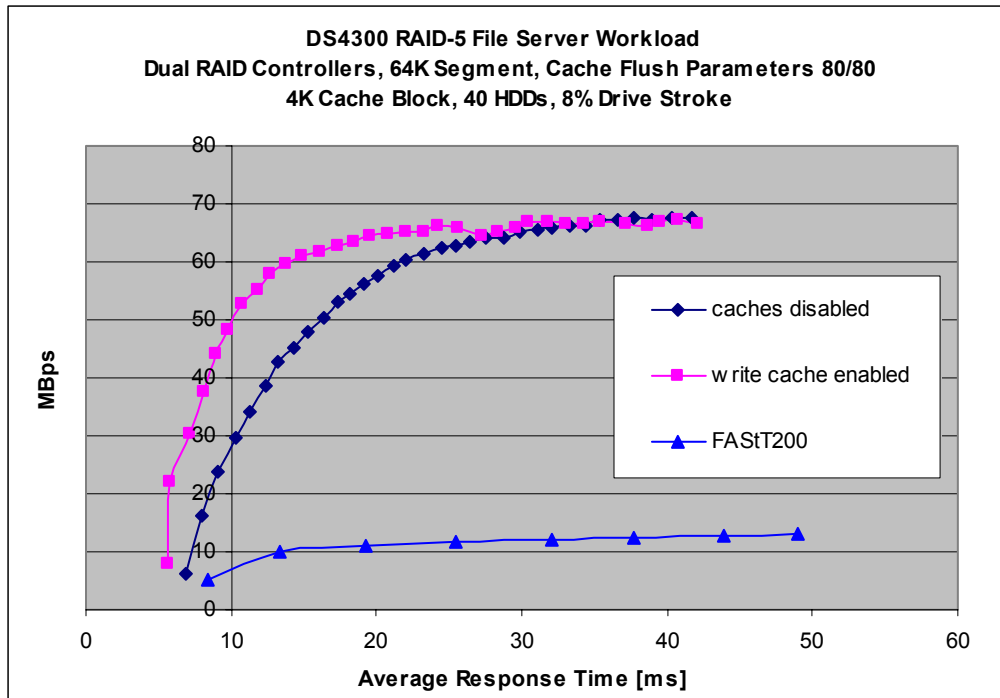
Workload		File Server [MBps]	Static Web Server [MBps]
RAID-5			
20 HDDs	All Caches Disabled	31	76
	Write Cache Enabled	36	N/A
40 HDDs	All Caches Disabled	57	137
	Write Cache Enabled	64	N/A
60 HDDs	All Caches Disabled	65	142
	Write Cache Enabled	66	N/A
80 HDDs	All Caches Disabled	66	142
	Write Cache Enabled	65	N/A

*Results correspond to an average response time of 20 ms.

Comparison of the Performance of the DS4300 (FAStT600) and FAStT200 Storage Servers

Chart 5 illustrates a comparison of the DS4300 and FAStT200 executing the File Server workload. Sampling the data at an average response time of 20 ms for the curves labeled "FAStT200" and "write cache enabled" yields a transfer rate of 64MBps for the DS4300 and approximately 11MBps for the FAStT200. Remember, however, that the FAStT200 measurements were conducted using 9GB 10K rpm drives. The performance of the FAStT200 is boosted by 10% to compensate for the drive technology discrepancy. Running the numbers reveals that the DS4300 performs approximately 430% better than the FAStT200 for the File Server workload. It should be noted that the result for the FAStT200 was achieved using 48 hard drives – eight more than the DS4300.

Chart 5. RAID-5 Dual DS4300s File Server vs. Dual FAStT200s



RAID-5 Random Reads and Random Writes Workload Results

Table 3 contains the results for the Random Reads 4K and 8K workloads, and the results for the Random Writes 4K and 8K workloads. Like the File Server workload and Web Server workload, these random read and write workloads are used to track product performance and serve as a comparison between similar products. In addition, these workloads are found in some production environments.

Table 3. RAID-5 Dual DS4300s Random Reads and Random Writes IOps*

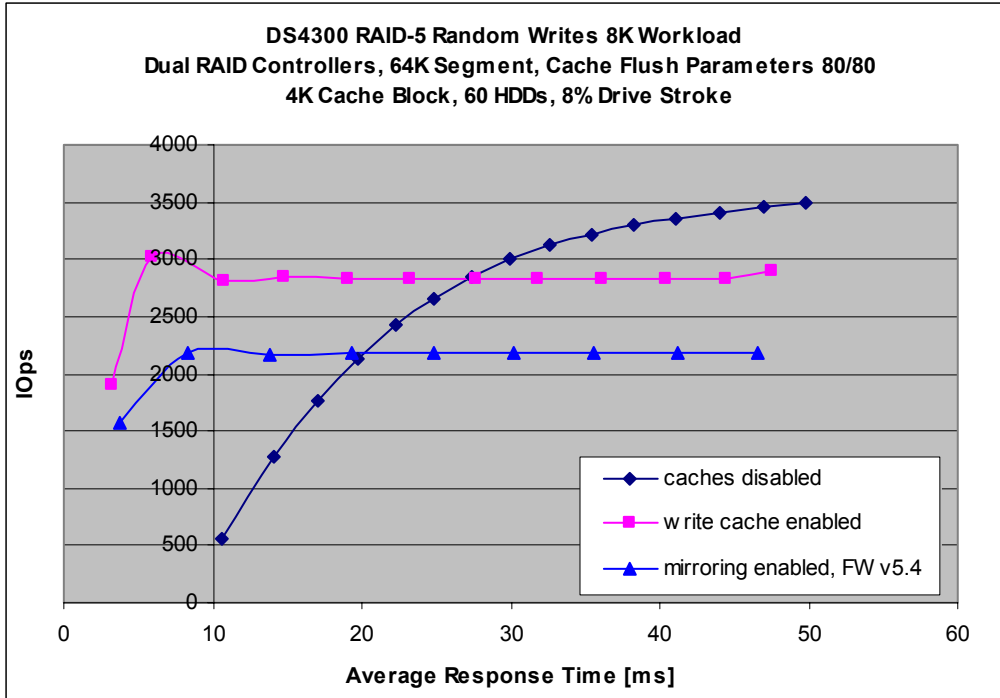
Workload		Random Reads 4K		Random Reads 8K		Random Writes 4K		Random Writes 8K	
		IOps	MBps	IOps	MBps	IOps	MBps	IOps	MBps
RAID-5									
20 HDDs	All Caches Disabled	7520	30	7220	57.5	800	3	775	6
	Write Cache Enabled	N/A	N/A	N/A	N/A	1775	7	1820	14.5
40 HDDs	All Caches Disabled	14,830	59	13,600	108	1600	6.5	1500	12
	Write Cache Enabled	N/A	N/A	N/A	N/A	4000	16	3020	24
60 HDDs	All Caches Disabled	16,750	67	14,130	113	2350	9.5	2140	17
	Write Cache Enabled	N/A	N/A	N/A	N/A	3600	14.5	2830	22.5
80 HDDs	All Caches Disabled	16,850	67.5	14,050	112	3000	12	2600	21
	Write Cache Enabled	N/A	N/A	N/A	N/A	3550	14	2815	22.5

*Results correspond to an average response time of 20 ms.

It is important to note the “all caches disabled” results for the Random Writes 4K and 8K workload. The performance is significantly lower than the “write cache enabled” results. The reason is that when the write cache is disabled for this workload in a RAID-5 environment, the performance is severely degraded due to the read modify write disk operation required for each write request. With the write-back cache enabled, the entire stripe set can be retrieved at once, and all updates can be done within the cache, eliminating the need for separate disk-read-modify-write operations to each individual stripe unit. Once the entire RAID-5 stripe set has been modified in cache, the stripe set can be written directly to disk. See the IBM Redbook, “Tuning IBM eServer xSeries Servers for Performance,” for a review of the RAID levels. In production, a user would be well-advised to enable write caching in order to reap the benefits provided by the DS4000 implementation of the RAID-5 algorithm. The benefits provided by enabling write cache in a RAID-5 environment for a workload that consists of a large percentage of data modifications is illustrated in Chart 6.

Comparison of the Performance of the DS4300 (FAStT600) and FAStT200 Storage Servers

Chart 6. RAID-5 Dual DS4300s Random Writes 8K Workload



The Random Reads workload can be used to compare the performance of the DS4300 to the FAStT200. With 40 drives attached, the FAStT200 was capable of achieving approximately 3780 8K IOPS. Table 3 shows that the DS4300 is capable of achieving 13,600 8K IOPS. Again, compensating by 10% for the older drive technology used with the FAStT200, the DS4300 performs approximately 225% better than the FAStT200 for the RAID-5 Random Reads workload.

RAID-10 OLTP Workload Results

The RAID-10 results will contain nearly the same tables and charts as the RAID-5 results. The examples, interpretations, “boxed” statements, and so forth used with the RAID-5 results will not be repeated with the RAID-10 results. However, there will be comparisons of DS4300 performance and FAStT200 performance.

Table 4 contains RAID-10 measurement results for the OLTP workload for various transfer request sizes. The data in Table 4 corresponds to the results achieved at an average response time of 20 ms.

Table 4. RAID-10 Dual DS4300s OLTP IOps*

Workload		OLTP 4K		OLTP 8K		OLTP 16K		OLTP 32K		OLTP 64K	
RAID-10		IOps	MBps	IOps	MBps	IOps	MBps	IOps	MBps	IOps	MBps
20 HDDs	All Caches Disabled	4880	19.5	4660	37	4280	68.5	3180	101	2285	146
	Write Cache Enabled	5080	20	4915	39	4510	72	3275	104	2380	152
40 HDDs	All Caches Disabled	9320	37	8815	70.5	7360	117	5360	171	3340	213
	Write Cache Enabled	10,800	43	9375	75	7250	116	5050	161	3250	208
60 HDDs	All Caches Disabled	11,240	45	9800	78	7650	122	5580	178	3380	216
	Write Cache Enabled	12,030	48	9335	74.5	7080	113	5220	167	3270	209
80 HDDs	All Caches Disabled	11,400	45.5	9910	79	7690	123	5590	178	3375	216
	Write Cache Enabled	11,650	46.5	9175	73	7000	112	5190	166	3245	207

*Results correspond to an average response time of 20 ms.

Chart 7 illustrates the affect that the drive stroke percentage plays on performance with respect to the number of IOps an array of 10 drives can be sustained for the RAID-10 8K OLTP workload. The results in Chart 7 correspond to a constant queue depth of 15 outstanding I/Os at the array. The same trend observed for the RAID-5 measurements (i.e., drive performance decreases as the drive stroke percentage increases) is clearly demonstrated in Chart 7.

Comparison of the Performance of the DS4300 (FAStT600) and FAStT200 Storage Servers

Chart 7. RAID-10 OLTP 8K IOps vs. Drive Stroke Percentage

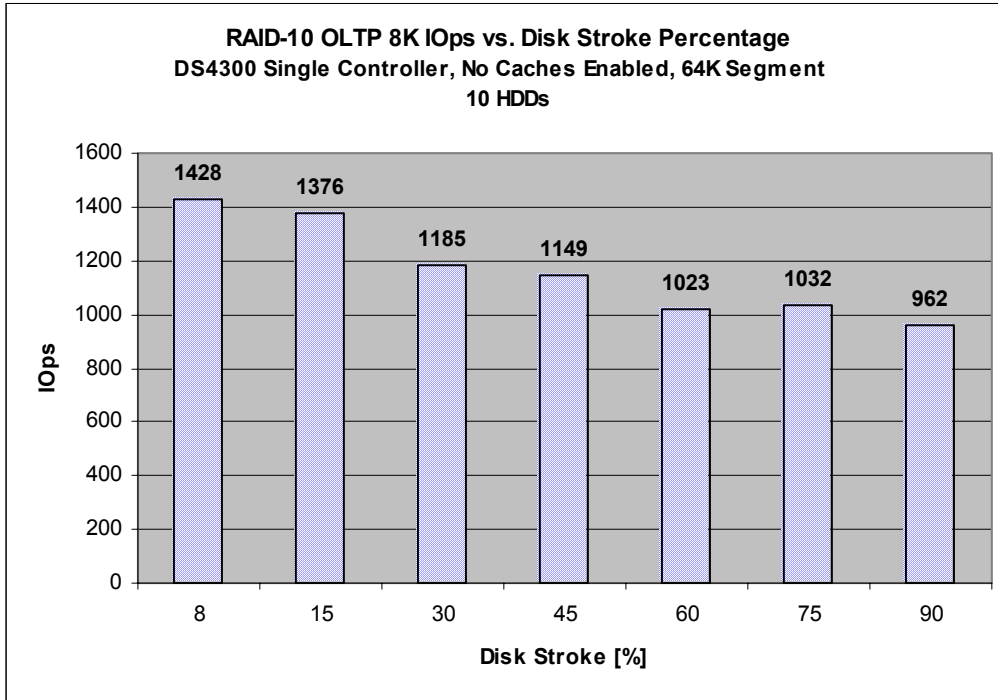
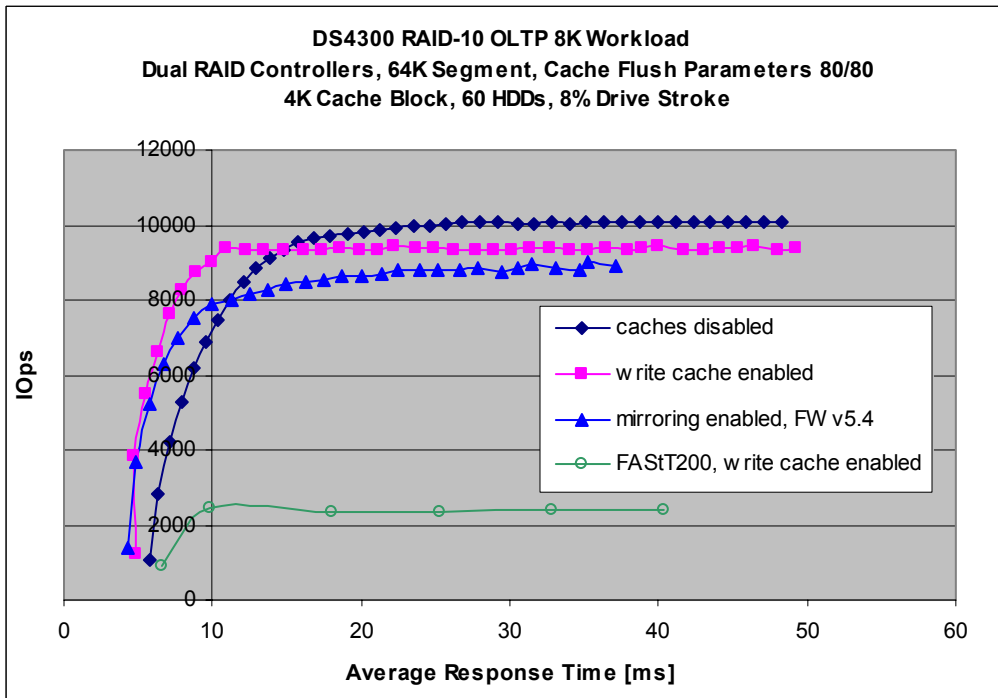


Chart 8 compares the performance of dual DS4300s with dual FAStT200s for the RAID-10 OLTP 8K workload.

Chart 8. DS4300 RAID-10 OLTP 8K vs. FAStT200



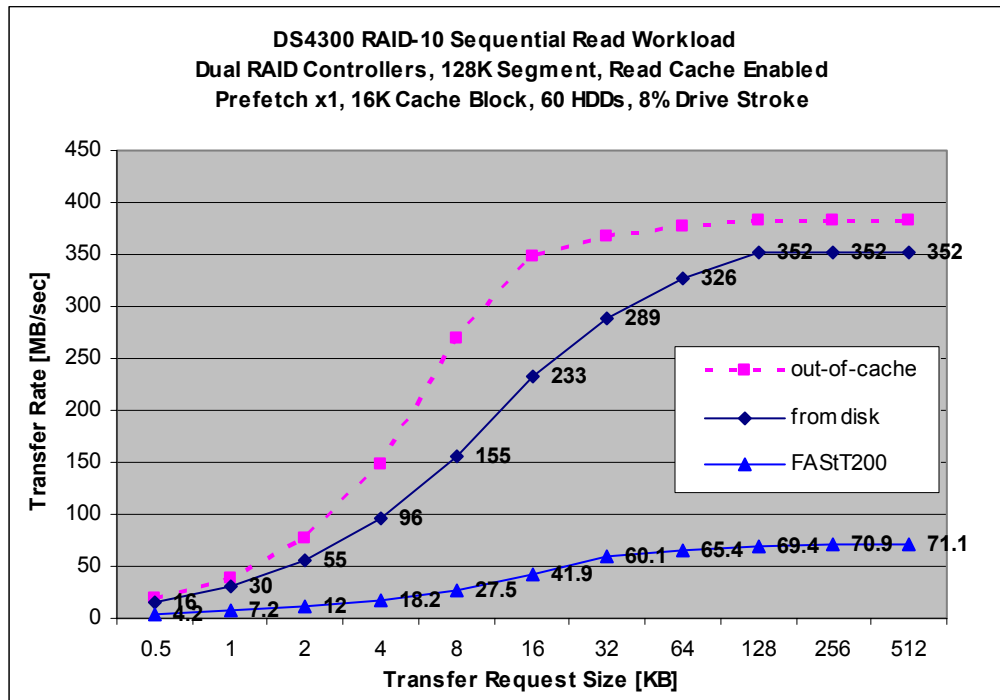
Comparison of the Performance of the DS4300 (FAST600) and FAST200 Storage Servers

Comparing the “write cache enabled” and the “FAST200” curves in Chart 8, and sampling the data that corresponds to a 20 ms average response time reveal that the DS4300 performs approximately 290% better than the FAST200. After correcting for the drive technology discrepancies, the DS4300 performs approximately 250% better than the FAST200 for the RAID-10 OLTP 8K workload.

RAID-10 Streaming Reads and Streaming Writes Workload Results

Chart 9 illustrates the RAID-10 Streaming Reads performance of dual DS4300 controllers. Be advised that these results may not correspond to “peak” performance. All of the transfer rates correspond to average response times no greater than 20 ms.

Chart 9. RAID-10 Dual DS4300s Streaming Reads Transfer Rate

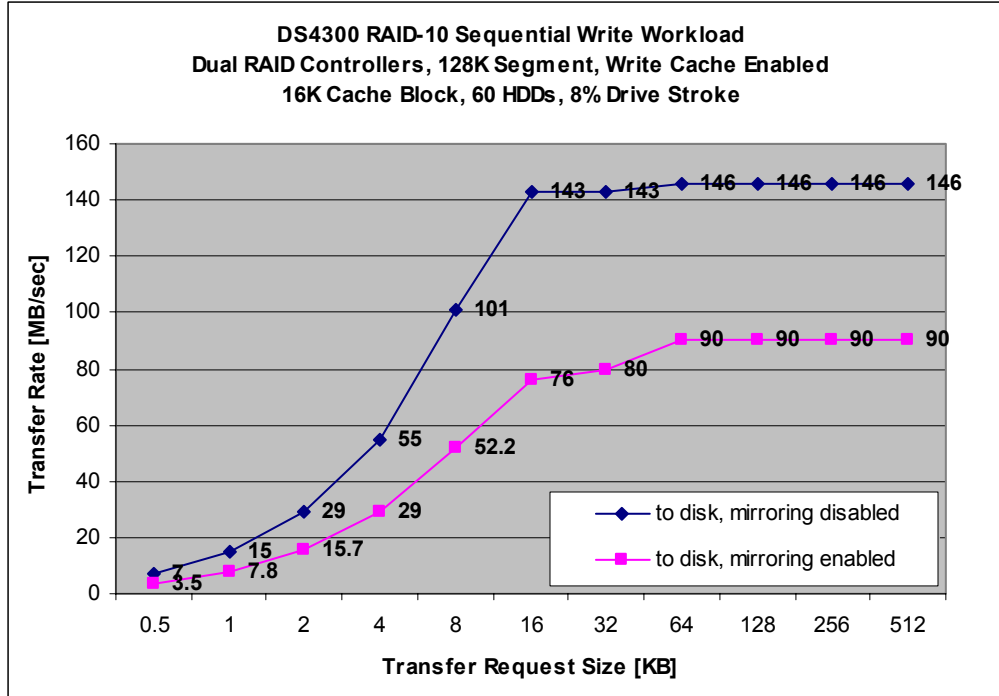


The transfer rate for large transfer blocks, greater than or equal to 64K, is approximately 70MBps for the FAST200 and 350MBps for the DS4300. After compensating for older drive technology, the DS4300 performs approximately 335% better than the FAST200 for the RAID-10 Streaming Reads workload.

Comparison of the Performance of the DS4300 (FASTt600) and FASTt200 Storage Servers

Chart 10 illustrates the Streaming Writes performance of dual DS4300 controllers. Be advised that these results may not correspond to “peak” performance. All of the transfer rates correspond to an average response time no greater than 20 ms.

Chart 10. RAID-10 Dual DS4300s Streaming Writes Transfer Rate



RAID-10 File Server and Static Web Server Workload Results

Table 5 contains the results of the File Server and the Static Web Server workload measurements. The same warning stated in the RAID-5 File Server / Web Server section for interpreting these results applies in this section, as well.

The File Server and Static Web Server workloads are used primarily to track product performance, serve as a comparison between similar products, and exercise a product with a workload that consists of various transfer request sizes.

Comparison of the Performance of the DS4300 (FASTt600) and FASTt200 Storage Servers

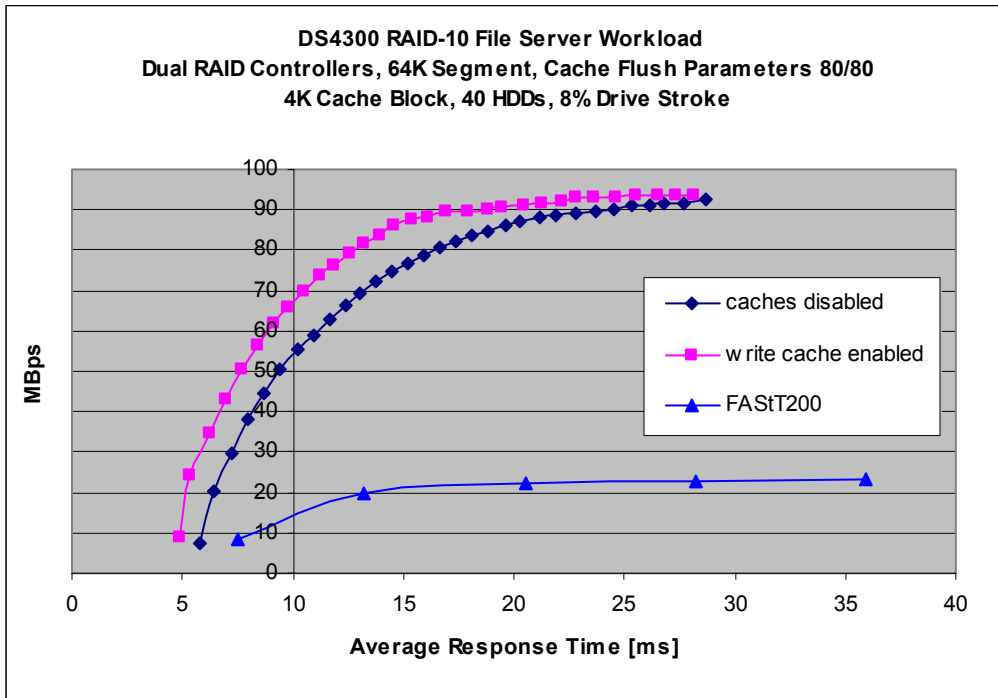
Table 5. RAID-10 Dual DS4300 File Server and Static Web Server Transfer Rate*

Workload		File Server [MBps]	Static Web Server [MBps]
RAID-10			
20 HDDs	All Caches Disabled	47	77
	Write Cache Enabled	48	N/A
40 HDDs	All Caches Disabled	87	137
	Write Cache Enabled	90	N/A
60 HDDs	All Caches Disabled	92	144
	Write Cache Enabled	94	N/A
80 HDDs	All Caches Disabled	93	144
	Write Cache Enabled	93	N/A

*Results correspond to an average response time of 20 ms.

Chart 11 illustrates a comparison of the DS4300 and FASTt200 executing the File Server workload. Sampling the data at an average response time of 20 ms for the curves labeled “FASTt200” and “write cache enabled” yields a transfer rate of 90MBps for the DS4300 and approximately 22MBps for the FASTt200. Correcting again for the drive technology discrepancies between those used for the FASTt200 measurements and the DS4300 measurements reveals that the DS4300 performs approximately 270% better than the FASTt200 for the File Server workload. It should be noted that the result for the FASTt200 was achieved using 48 hard drives – eight more than the DS4300.

Chart 11. RAID-10 Dual DS4300s File Server vs. Dual FASTt200s



RAID-10 Random Reads and Random Writes Workload Results

Table 6 contains the results for the Random Reads 4K and 8K workloads, and the results for the Random Writes 4K and 8K workloads. Like the File Server workload and Web Server workload, these random read and write workloads are used to track product performance and serve as a comparison between similar products. In addition, these workloads are found in some production environments.

Table 6. RAID-10 Dual DS4300s Random Reads and Random Writes IOps*

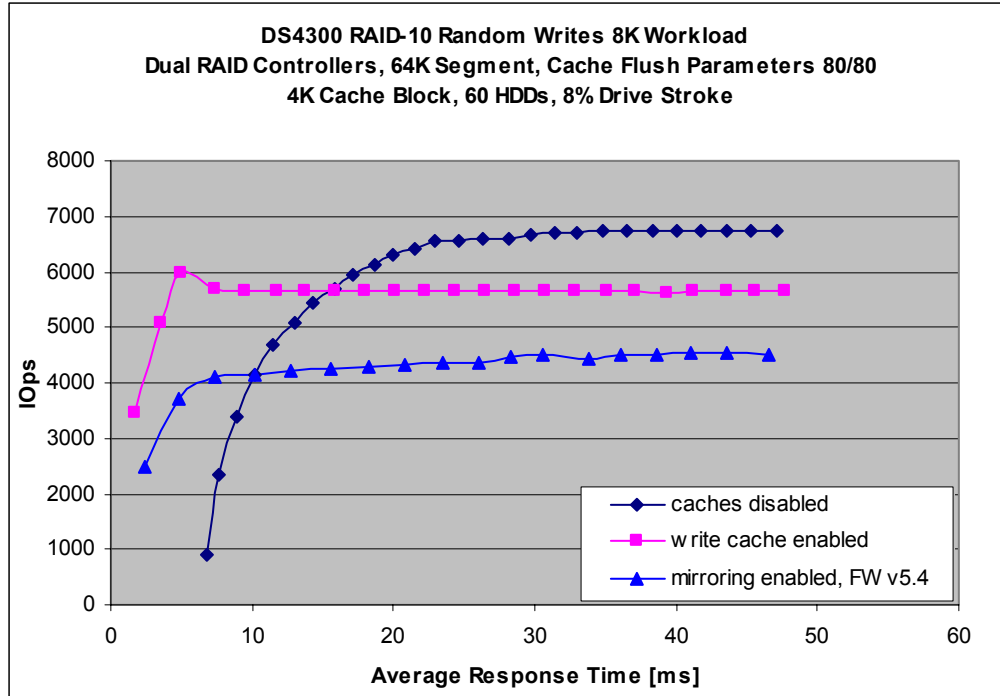
Workload		Random Reads 4K		Random Reads 8K		Random Writes 4K		Random Writes 8K	
		IOps	MBps	IOps	MBps	IOps	MBps	IOps	MBps
RAID-10									
20 HDDs	All Caches Disabled	7555	30	7285	58	2750	11	2620	21
	Write Cache Enabled	N/A	N/A	N/A	N/A	3785	15	3640	29
40 HDDs	All Caches Disabled	14,700	58.5	13,600	108	5215	20.5	4975	39.5
	Write Cache Enabled	N/A	N/A	N/A	N/A	6725	27	5850	46.5
60 HDDs	All Caches Disabled	16,850	67	14,100	113	6885	27.5	6305	50
	Write Cache Enabled	N/A	N/A	N/A	N/A	6625	26.5	5665	45
80 HDDs	All Caches Disabled	17,000	68	14,180	113	7320	29	6525	52
	Write Cache Enabled	N/A	N/A	N/A	N/A	6510	26	5665	45

*Results correspond to an average response time of 20 ms.

Chart 12 illustrates the performance of dual DS4300 RAID controllers for the RAID-10 Random Writes 8K workload.

Comparison of the Performance of the DS4300 (FAStT600) and FAStT200 Storage Servers

Chart 12. RAID-10 Dual DS4300 Random Writes 8K Workload



The Random Reads workload can be used to compare the performance of the DS4300 to the FAStT200. With 40 drives attached, the FAStT200 was capable of achieving approximately 3175 8K IOPs. Table 3 shows that the DS4300 is capable of achieving 13,600 8K IOPs. Finally, compensating for older drive technology used with the FAStT200, the DS4300 performs approximately 290% better than the FAStT200 for the Random Reads workload.

Summary

The DS4300 offers substantial performance improvement over its predecessor, the FAStT200. Some examples of the approximate performance gains are:

- 350% for the RAID-5 OLTP 8K workload
- 300% for the RAID-5 Streaming Reads workload
- 430% for the RAID-5 File Server workload
- 225% for the RAID-5 Random Reads workload

- 250% for the RAID-10 OLTP 8K workload
- 335% for the RAID-10 Streaming Reads workload
- 270% for the RAID-10 File Server workload
- 290% for the RAID-10 Random Reads workload

Comparison of the Performance of the DS4300 (FAStT600) and FAStT200 Storage Servers



© IBM Corporation 2004

IBM Systems and Technology Group

Department MX5

Research Triangle Park NC 27709

Produced in the USA

9-04

All rights reserved.

IBM, the IBM and eServer logos, eServer, TotalStorage, ServeRAID and xSeries are trademarks or registered trademarks of IBM Corporation in the United States and/or other countries.

Pentium and Xeon are trademarks or registered trademarks of Intel Corporation.

Microsoft and Windows are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Other company, product, and service names may be trademarks or service marks of others.

The performance data contained herein was obtained in a controlled, isolated environment. Actual results that may be obtained in other operating environments may vary significantly. While IBM has reviewed each item for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere.

The responsibility for use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's or user's ability to evaluate and integrate them into their operating environment. Customers or users attempting to adapt these techniques to their own environments do so at their own risk. In no event shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.

IBM reserves the right to change specifications or other product information without notice. References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. IBM PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

IBM xSeries servers are assembled in the U.S., Great Britain, Japan, Australia and Brazil and are composed of U.S. and non-U.S. parts.

This publication may contain links to third party sites that are not under the control of or maintained by IBM. Access to any such third party site is at the user's own risk and IBM is not responsible for the accuracy or reliability of any information, data, opinions, advice or statements made on these sites. IBM provides these links merely as a convenience and the inclusion of such links does not imply an endorsement.