Comparison of the Performance of the FAStT900 and FAStT700 Storage Servers

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

Abstract

This paper provides the baseline performance of the IBM® TotalStorage® FAStT900 Storage Server, which employs IBM Fibre Array Storage Technology (FAStT).

The purpose of this paper is to present the results obtained using the Iometer tool to measure the raw performance of the FAStT900's RAID subsystem. The FAStT900's performance is compared to that of its predecessor, the IBM TotalStorage FAStT700 Storage Server.[1]

The paper is organized in four sections. The first section briefly describes the tool used to measure the performance of the FAStT900 and FAStT700, 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 FAStT900.

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@IBMUS.

---

[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, a discussion of "out-of-cache" or "to-cache" measurements does not fit within the scope of this document.

Measurement Tool and Workloads

Iometer is a workload generator and a measurement tool originally developed by 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. Intel does not endorse any Iometer results.

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

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 and the IBM eServer xSeries® 235 with two Intel Pentium 4 2.4GHz Xeon processors and 1024MB of system memory. The operating system installed was Microsoft Windows 2000 Advanced Server 5.0.2195 (Build 2195) with Service Pack 4.

Both systems contained two IBM TotalStorage FAStT FC-2-133 Host Bus Adapters using driver version 8.2.3.63.

The Fibre Channel management software used was Redundant Disk Array Controller (RDAC) version 08.41.95.02 and Storage Manager Client version 08.41.G5.01.

The Fibre Channel RAID controller software used on the FAStT900 RAID controllers was Appware version 05.40.06.00, Bootware version 05.40.02.00, and NVSRAM version N1742F900R840V02. Each FAStT900 RAID controller contained 1024MB of cache.

The storage backend consisted of 14 IBM TotalStorage FAStT EXP700 Storage Expansion Enclosures with ESM firmware version 9325. The EXP700 enclosures contained 18.2GB, 15K rpm drives with firmware version B947.

Configuration Diagrams for the Measured Hardware

Random Access Measurements Hardware Configuration Diagram

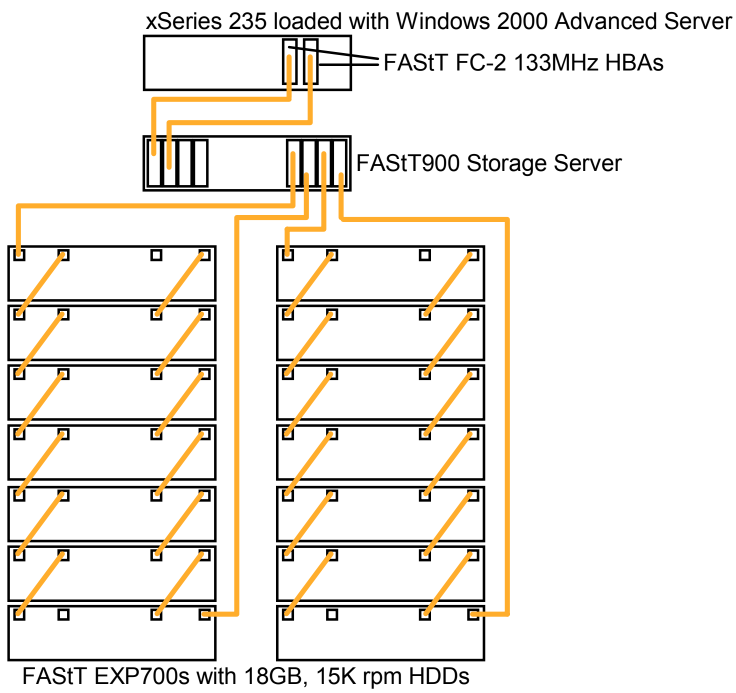

Sequential Access Measurements Hardware Configuration Diagram

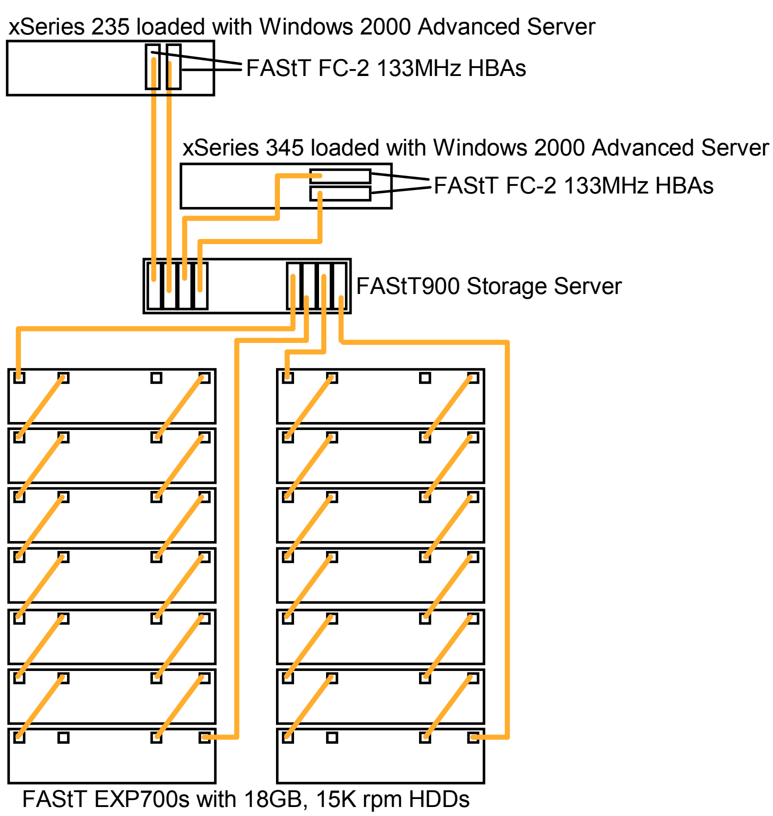

Measurement Results and Analysis

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 dual FAStT900 RAID controllers with a specified number of hard disk drives (HDDs) utilized that corresponds 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 IOps or MBps regardless of the average response time associated with that result. However, since most server applications will not wait forever for disk I/O to complete, the 20ms threshold was used because it represents a reasonable amount of time for an application to wait for completion of disk I/O before overall performance begins to decrease.

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, because this produced optimal disk subsystem throughput for these workloads. Cache mirroring was disabled.

Table 1. FAStT900 RAID-5 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	3095	12	3000	23	2790	43	2185	68	1660	104
	Write Cache Enabled	3860	15	3800	29	3620	56	2620	82	1880	117
60 HDDs	All Caches Disabled	9280	36	8880	69	8080	126	5120	160	2780	174
	Write Cache Enabled	11,360	44	10,960	85	9885	154	5525	172	3150	196
100 HDDs	All Caches Disabled	15,340	60	14,320	111	9055	141	5155	161	2755	172
	Write Cache Enabled	24,340	95	15,880	124	10,315	161	5765	180	3215	200
140 HDDs	All Caches Disabled	20,320	79	15,025	117	9060	141	5170	161	2770	173
	Write Cache Enabled	20,500	80	16,525	129	10,340	161	5780	180	3200	200
180 HDDs	All Caches Disabled	18,965	74	15,040	117	9130	142	5200	162	2780	173
	Write Cache Enabled	20,540	80	16,450	128	10,345	161	5800	180	3200	200

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

It is worth observing two things concerning Table 1. First, it can be observed that the IOps tend to decrease from left to right across the table as the transfer request size increases. However, notice that the transfer rate (MBps) increases from left to right. Transferring larger amounts of data for each request utilizes the fiber bandwidth more efficiently.

Second, the IOps and transfer rate tend to increase from top to bottom as more drives are utilized. Furthermore, notice that the IOps and transfer rate tend to level off or drop slightly in the bottom half of the table. For example, following the 8K column down, it can be observed that the IOps with 140 HDDs utilized and write cache enabled are 16,525. The IOps with 180 HDDs utilized and write cache enabled are 16,450. This means that 16,525 IOps is probably the maximum number of RAID-5 OLTP 8K IOps to disk, that dual FAStT900 RAID controllers can achieve.

Do not conclude that adding more than 140 HDDs will not increase the number of 8K IOps the FAStT900 Storage Server can achieve in production environments. In production environments where the working data set spans a greater percentage of the disk capacity, a user could expect to see an increase in the number of IOps as HDDs are added beyond 140 HDDs. However, the IOps will be less than the results in Table 1.

Remember, the data in Table 1 was achieved using a disk stroke of 8% and does not reflect typical storage 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. Table 2 contains information that can be used to approximate the performance of the FAStT900 Storage Server in a RAID-5 OLTP environment assuming different working data set sizes.

Table 2. RAID-5 OLTP Drive Stroke Percentage Multiplier

Transfer Request Size	**To approximate the RAID-5 OLTP IOps for the following disk stroke percentages, multiply the 8% disk stroke IOps from Table 1 by the corresponding multiplier…**					
	15%	30%	45%	60%	75%	90%
RAID-5 OLTP 4K, 8K, and 16K	0.88	0.80	0.74	0.72	0.66	0.65

For example, if a user has 100 drives attached to dual FAStT900 RAID controllers, and the 100 drives are dedicated to transaction processing with an average transfer request size of 8K, the performance the user could expect from a hardware perspective would greatly depend on the size of the working data set and the capacity utilization. The maximum performance the user could expect would be 14,320 IOps 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 maximum performance expected can be determined by multiplying 14,320 IOps by 0.74 (from Table 2), which equals 10,595 IOps.

Looking at the multipliers in Table 2 from left to right, it can be observed that the estimated performance will decrease as the disk stroke percentage increases. This is due to the mechanical positioning latencies of the HDDs as a result of seeking for random data accesses. Use Table 2 to obtain approximate performance and not precise performance. This estimation method does not consider the effect of users on performance.

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. This is due to the mechanical latencies of disk drives.

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 1 illustrates the performance of dual FAStT900 RAID controllers and dual FAStT700 RAID controllers for the RAID-5 OLTP 8K workload.

Chart 1. FAStT900 RAID-5 OLTP 8K vs. FAStT700 RAID-5 OLTP 8K

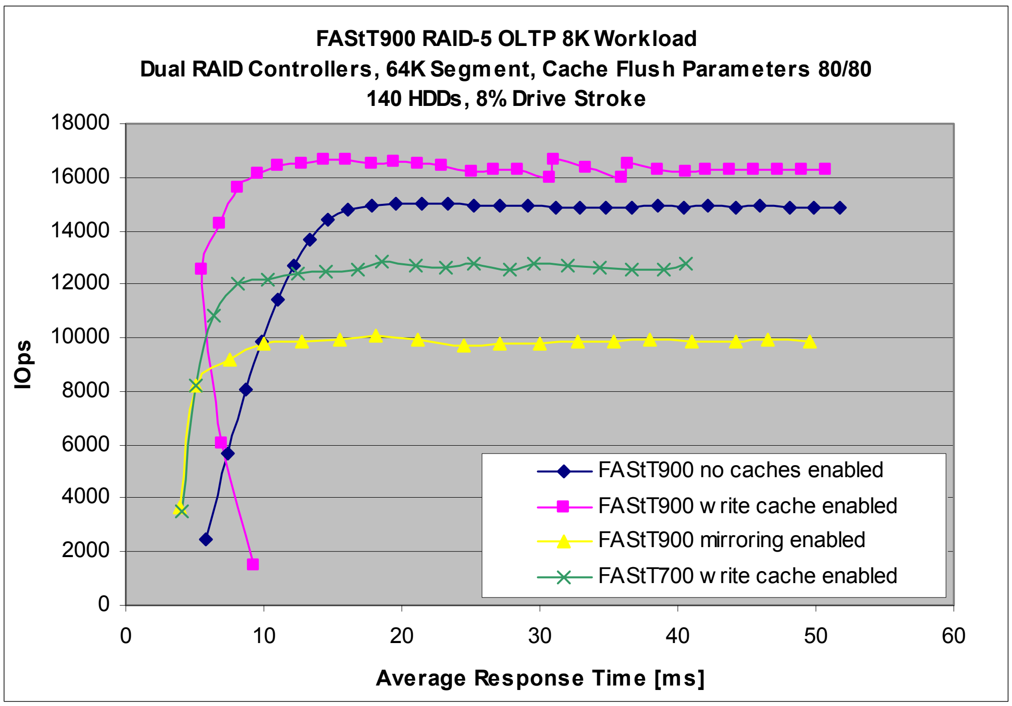

Comparing the "FAStT900 write cache enabled" and "FAStT700 write cache enabled" curves in Chart 1, and sampling the data that corresponds to a 20 ms average response time reveals that the FAStT900 performs approximately 30% better than the FAStT700 for the RAID-5 OLTP 8K workload utilizing 140 HDDs. The FAStT700 was measured using the same hardware and software as the FAStT900 for this comparison.

The FAStT900 Storage Server performs approximately 5% to 30% better than the FAStT700 Storage Server for the RAID-5 OLTP 8K workload depending on the number of HDDs utilized.

RAID-5 Streaming Reads and Streaming Writes Workload Results

Chart 2 illustrates the RAID-5 Streaming Reads performance of dual FAStT900 RAID controllers versus dual FAStT700 RAID controllers. The FAStT700 was measured using 108 10K rpm HDDs, a 64K segment, a read-ahead multiplier of 8 (prefetch x8), and version 5.0 firmware.

Chart 2. FAStT900 RAID-5 Streaming Reads vs. FAStT700 RAID-5 Streaming Reads

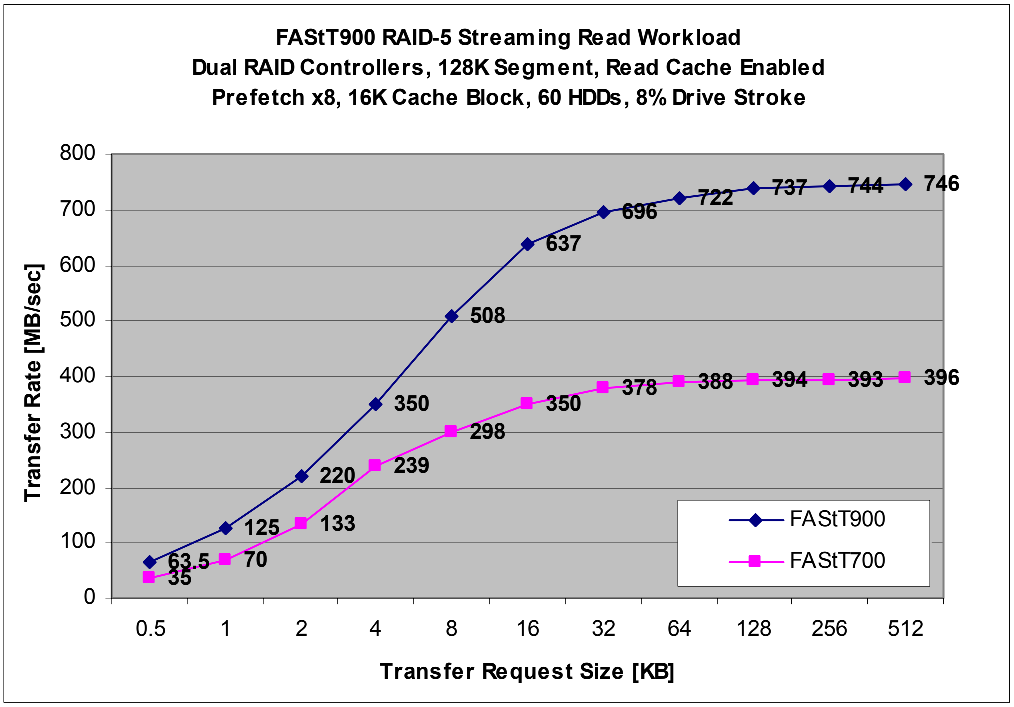

The FAStT700 transfer rate for 64K transfer requests is 388MBps, and the FAStT900 transfer rate for 64K transfer requests is 722MBps. After compensating for older drive technology by boosting the FAStT700 transfer rate by 5%, the FAStT900 performs approximately 75% better than the FAStT700 for 64K transfers for the RAID-5 Streaming Reads workload. The performance difference between the 10K rpm HDDs and the 15K rpm HDDs is probably more than 5%. However, the FAStT700 is not capable of achieving more than about 400MBps, so it is assumed that the controllers would bottleneck before the performance increase attributed to the 15K rpm HDDs could be realized.

The FAStT900 Storage Server performs approximately 50% to 75% better than the FAStT700 Storage Server for the RAID-5 Streaming Reads workload over a transfer request range of 512 bytes to 512KB utilizing 60 HDDs.

Chart 3 illustrates the RAID-5 Streaming Writes performance of dual FAStT900 RAID controllers versus dual FAStT700 RAID controllers. The FAStT700 was measured using 58 10K rpm HDDs, a 64K segment, write cache enabled, write cache mirroring disabled, and version 5.0 firmware.

Chart 3. FAStT900 RAID-5 Streaming Writes vs. FAStT700 RAID-5 Streaming Writes

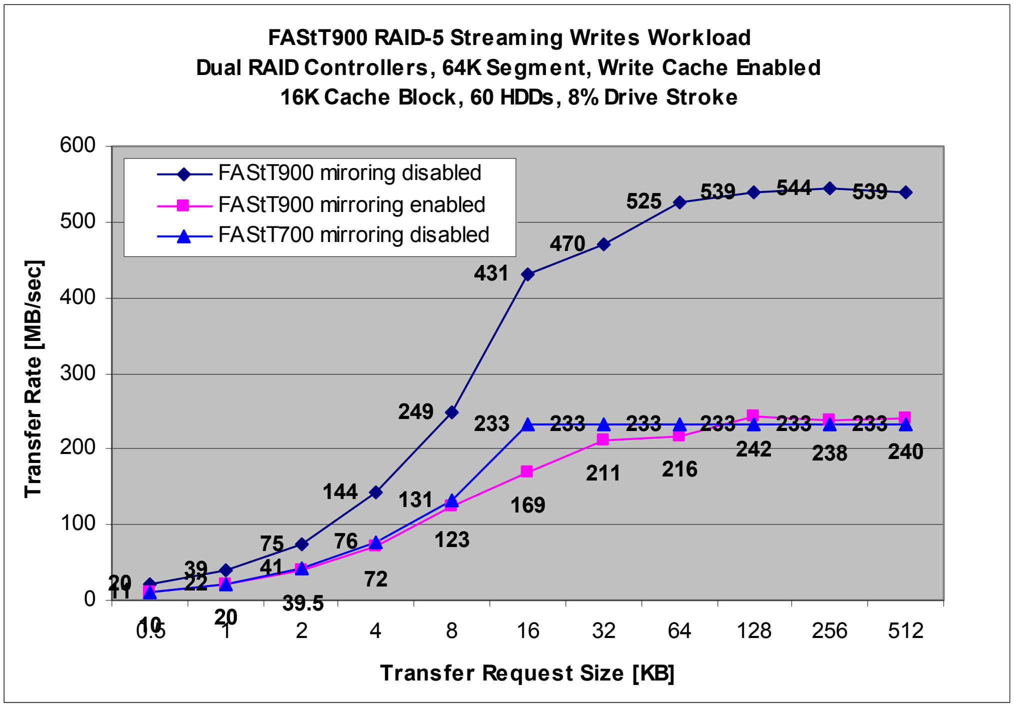

The FAStT700 transfer rate for 64K transfer requests is 233MBps, and the FAStT900 transfer rate for 64K transfer requests is 525MBps. The FAStT700 maximum transfer rate for streaming writes is approximately 240MBps. Therefore, the FAStT900 performs approximately 118% better than the FAStT700 for 64K transfers for the RAID-5 Streaming Writes workload.

The FAStT900 Storage Server performs approximately 70% to 120% better than the FAStT700 Storage Server for the RAID-5 Streaming Writes workload over a transfer request range of 512 bytes to 512KB utilizing 60 HDDs.

RAID-5 File Server and Static Web Server Workload Results

Table 3 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 a FAStT900 Storage Server 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 3. FAStT900 RAID-5 File Server and Static Web Server Transfer Rates*

Workload		File Server [MBps]	Static Web Server [MBps]
RAID-5			
20 HDDs	All Caches Disabled	35	90
	Write Cache Enabled	42	N/A
60 HDDs	All Caches Disabled	105	238
	Write Cache Enabled	117	N/A
100 HDDs	All Caches Disabled	133	245
	Write Cache Enabled	131	N/A
140 HDDs	All Caches Disabled	133	246
	Write Cache Enabled	132	N/A
180 HDDs	All Caches Disabled	133	242
	Write Cache Enabled	132	N/A

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

Notice that the IOps and transfer rate tend to level off in the bottom half of the table as did the OLTP results in Table 1. For example, following the File Server column down, it can be observed that the transfer rate does not increase from 100 HDDs to 180 HDDs. This means that 133MBps is the maximum transfer rate for the Iometer File Server workload that dual FAStT900 RAID controllers can achieve.

Chart 4 illustrates a comparison of the FAStT900 and FAStT700 executing the File Server workload. Sampling the data at an average response time of 20 ms yields a transfer rate of 133MBps for the FAStT900 and 114MBps for the FAStT700. The FAStT900 performs approximately 16% better than the FAStT700 for the RAID-5 File Server workload with 140 HDDs. The FAStT700 was measured using the same hardware and software as the FAStT900 for this comparison.

The FAStT900 Storage Server performs approximately 5% to 15% better than the FAStT700 Storage Server for the RAID-5 File Server workload depending on the number of HDDs utilized.

Chart 4. RAID-5 FAStT900 File Server Workload vs. FAStT700 File Server Workload

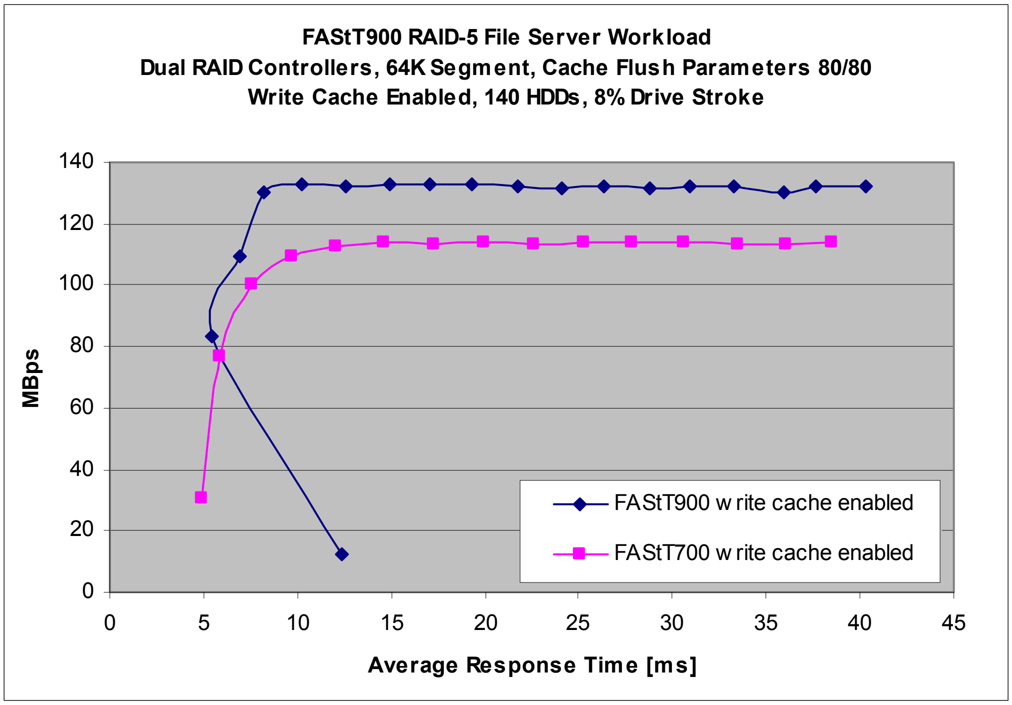

RAID-5 Random Reads and Random Writes Workload Results

Table 4 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 4. FAStT900 RAID-5 Random Reads and Random Writes IOps*

Workload		Random Reads 4K		Random Reads 8K		Random Writes 4K		Random Writes 8K	
RAID-5		IOps	MBps	IOps	MBps	IOps	MBps	IOps	MBps
20 HDDs	All Caches Disabled	8670	33.5	8365	65	898	3.5	890	6.5
	Write Cache Enabled	N/A	N/A	N/A	N/A	1100	4	460	3.5
60 HDDs	All Caches Disabled	25,080	97	24,395	190	2710	10.5	2640	20.5
	Write Cache Enabled	N/A	N/A	N/A	N/A	5015	19	5050	39
100 HDDs	All Caches Disabled	36,650	143	30,180	235	4480	17.5	4275	33
	Write Cache Enabled	N/A	N/A	N/A	N/A	8835	34	7675	60
140 HDDs	All Caches Disabled	42,555	166	30,020	234	6100	23	5700	44.5
	Write Cache Enabled	N/A	N/A	N/A	N/A	11,000	43	8325	65
180 HDDs	All Caches Disabled	40,900	159	29,480	230	7000	27.5	6285	49
	Write Cache Enabled	N/A	N/A	N/A	N/A	10,000	39	8400	65.5

*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 FAStT implementation of the RAID-5 algorithm. Chart 5 illustrates the benefits of enabling write cache for a workload that consists of a large percentage of data modifications.

Chart 5. FAStT900 RAID-5 Random Writes 8K Workload

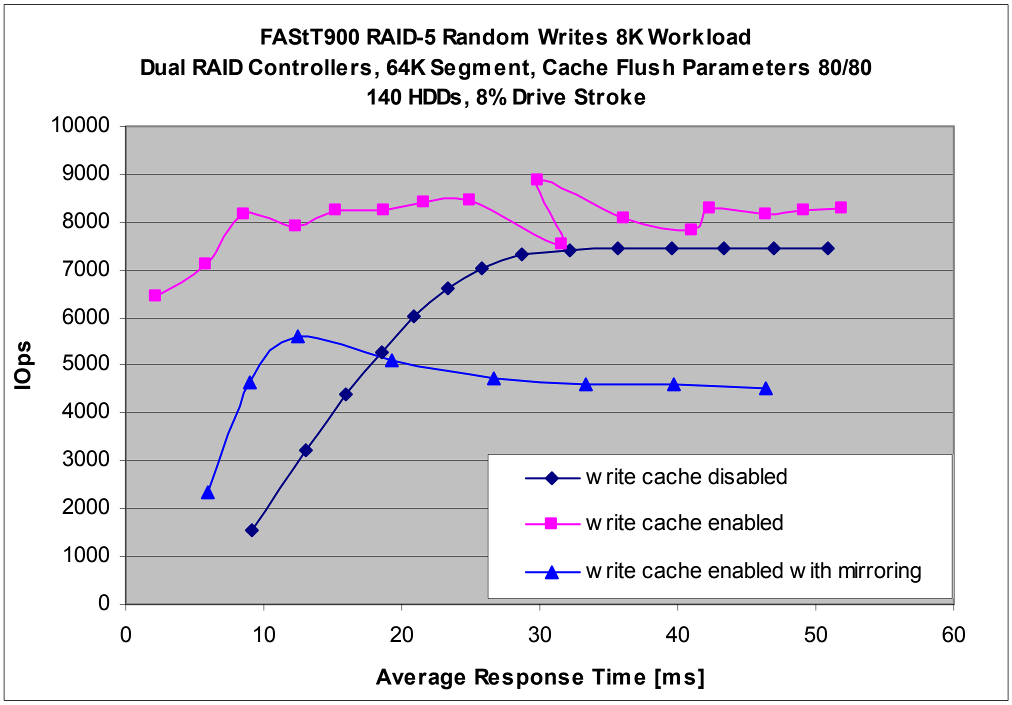

The Random Reads workload can be used to compare the performance of the FAStT900 to the FAStT700. Chart 6 illustrates this comparison.

Chart 6. RAID-5 FAStT900 8K Random Reads vs. FAStT700 8K Random Reads

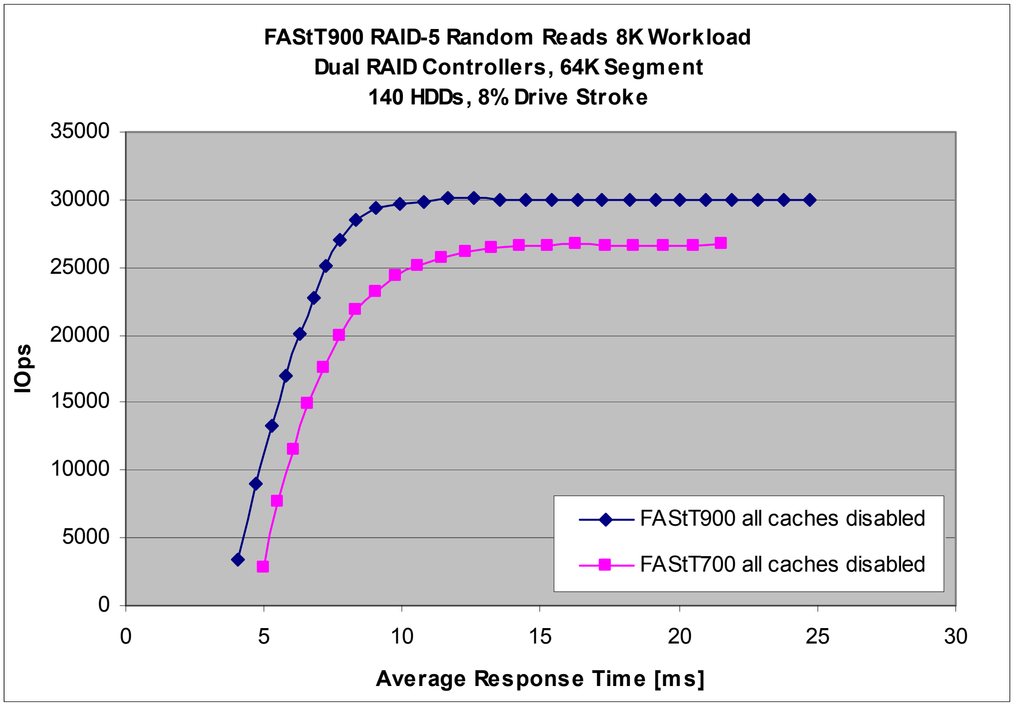

With 140 drives attached, the FAStT900 is capable of achieving approximately 30,020 8K IOps. The FAStT700 is capable of achieving approximately 26,640 8K IOps. Therefore, the FAStT900 performs approximately 13% better than the FAStT700 for the RAID-5 Random Reads workload with 140 HDDs.

The FAStT900 Storage Server performs approximately 5% to 15% better than the FAStT700 Storage Server for the RAID-5 Random Reads 8K workload depending on the number of HDDs utilized.

RAID-10 OLTP Workload Results

The RAID-10 results will contain the same tables and charts as the RAID-5 results. The examples, interpretations, and so forth used with the RAID-5 results will not be repeated with the RAID-10 results. However, there will be comparisons of FAStT900 performance and FAStT700 performance.

Table 5 contains RAID-10 measurement results for the OLTP workload for various transfer request sizes. The data in Table 5 corresponds to the results achieved at an average response time of 20 ms.

Table 5. FAStT900 RAID-10 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	5040	19	4840	37.5	4475	70	3440	107	2780	173
	Write Cache Enabled	5895	23	5660	44	5185	81	3975	124	2800	174
60 HDDs	All Caches Disabled	15,425	60	15,000	117	13,455	210	8420	262	4585	286
	Write Cache Enabled	15,435	60	15,000	117	14,400	225	8700	272	5065	316
100 HDDs	All Caches Disabled	25,125	98	23,425	183	14,575	227	8450	264	4580	286
	Write Cache Enabled	26,100	102	23,800	186	15,500	242	9370	292	5070	316
140 HDDs	All Caches Disabled	30,480	118	23,450	183	14,485	226	8425	263	4570	286
	Write Cache Enabled	34,330	134	25,075	195	16,240	253	9430	294	5080	317
180 HDDs	All Caches Disabled	25,450	99	21,735	169	14,600	228	8535	266	4630	289
	Write Cache Enabled	25,225	98.5	25,330	198	16,350	255	9380	293	5060	316

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

Table 6 contains information that can be used to approximate the performance of the FAStT900 Storage Server in a RAID-10 OLTP environment assuming different working data set sizes.

Table 6. RAID-10 OLTP Drive Stroke Percentage Multiplier

Transfer Request Size	**To approximate the RAID-10 OLTP IOps for the following disk stroke percentages, multiply the 8% disk stroke IOps from Table 1 by the corresponding multiplier…**					
	15%	30%	45%	60%	75%	90%
RAID-10 OLTP 4K, 8K, and 16K	0.95	0.82	0.80	0.72	0.72	0.67

The same trend observed for the RAID-5 measurements (i.e., drive performance decreases as the drive stroke percentage increases) is clearly observed in Table 6.

Chart 7 illustrates the performance of dual FAStT900 RAID controllers and dual FAStT700 RAID controllers for the RAID-10 OLTP 8K workload.

.

Chart 7. FAStT900 RAID-10 OLTP 8K vs. FAStT700

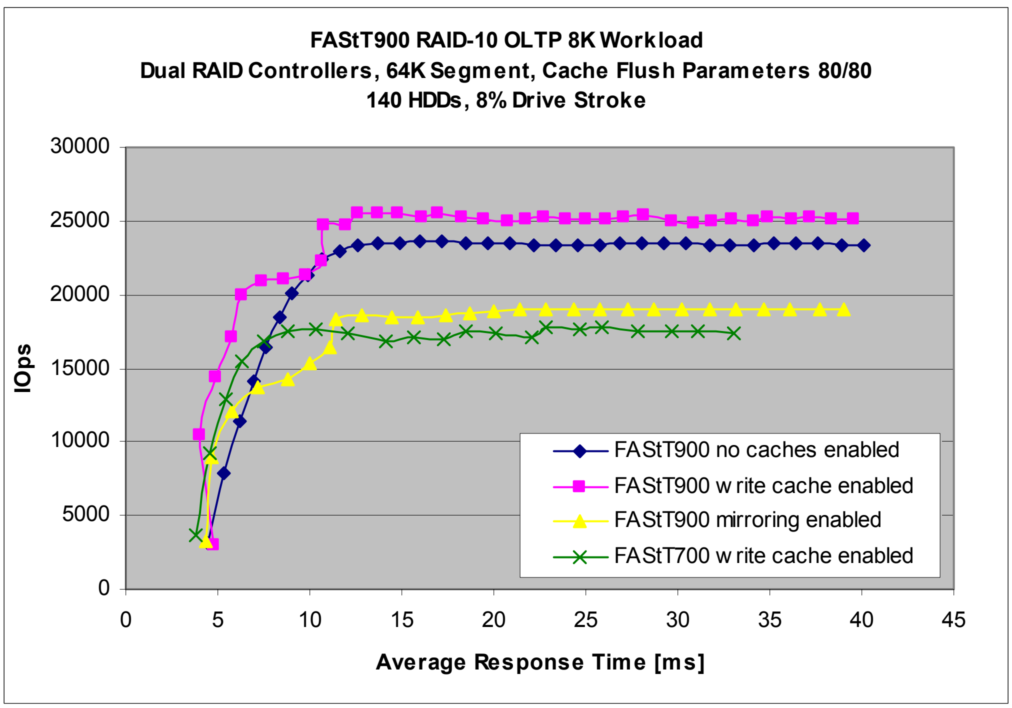

Comparing the "FAStT900 write cache enabled" and "FAStT700 write cache enabled" curves in Chart 7, and sampling the data that corresponds to a 20 ms average response time reveals that the FAStT900 performs approximately 45% better than the FAStT700 for the RAID-10 OLTP 8K workload utilizing 140 HDDs. The FAStT700 was measured using the same hardware and software as the FAStT900 for this comparison.

The FAStT900 Storage Server performs approximately 5% to 45% better than the FAStT700 Storage Server for the RAID-10 OLTP 8K workload depending on the number of HDDs utilized.

RAID-10 Streaming Reads and Streaming Writes Workload Results

Chart 8 illustrates the RAID-10 Streaming Reads performance of dual FAStT900 RAID controllers versus dual FAStT700 RAID controllers. The FAStT700 was measured with 96 10K rpm HDDs, a 64K segment, a read-ahead multiplier of 8 (prefetch x8), and version 5.0 firmware.

Chart 8. FAStT900 RAID-10 Streaming Reads Transfer Rate

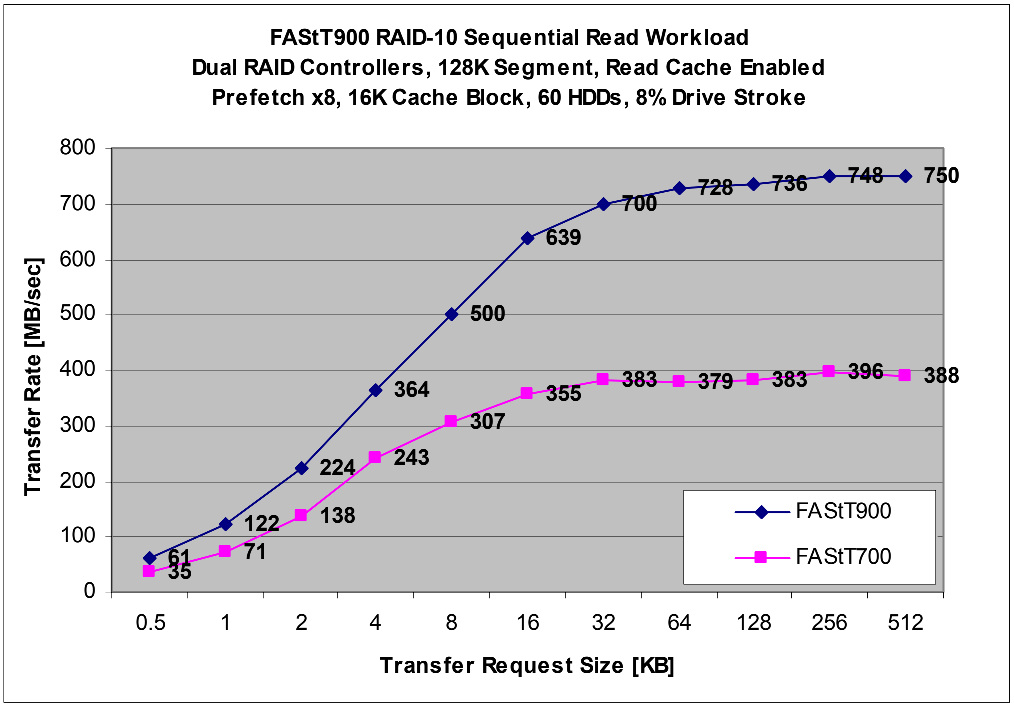

The FAStT700 transfer rate for 64K transfer requests is 379MBps, and the FAStT900 transfer rate for 64K transfer requests is 728MBps. After compensating for older drive technology by boosting the FAStT700 transfer rate by 5% for the same reason explained for the RAID-5 Streaming Reads workload results, the FAStT900 performs approximately 83% better than the FAStT700 for 64K transfers for the RAID-10 Streaming Reads workload.

The FAStT900 Storage Server performs approximately 50% to 85% better than the FAStT700 Storage Server for the RAID-10 Streaming Reads workload over a transfer request range of 512 bytes to 512KB utilizing 60 HDDs.

Chart 9 illustrates the RAID-10 Streaming Writes performance of dual FAStT900 RAID controllers versus dual FAStT700 RAID controllers. The FAStT700 was measured using 58 10K rpm HDDs, a 64K segment, write cache enabled, write cache mirroring disabled, and version 5.0 firmware.

Chart 9. FAStT900 RAID-10 Streaming Writes Transfer Rate

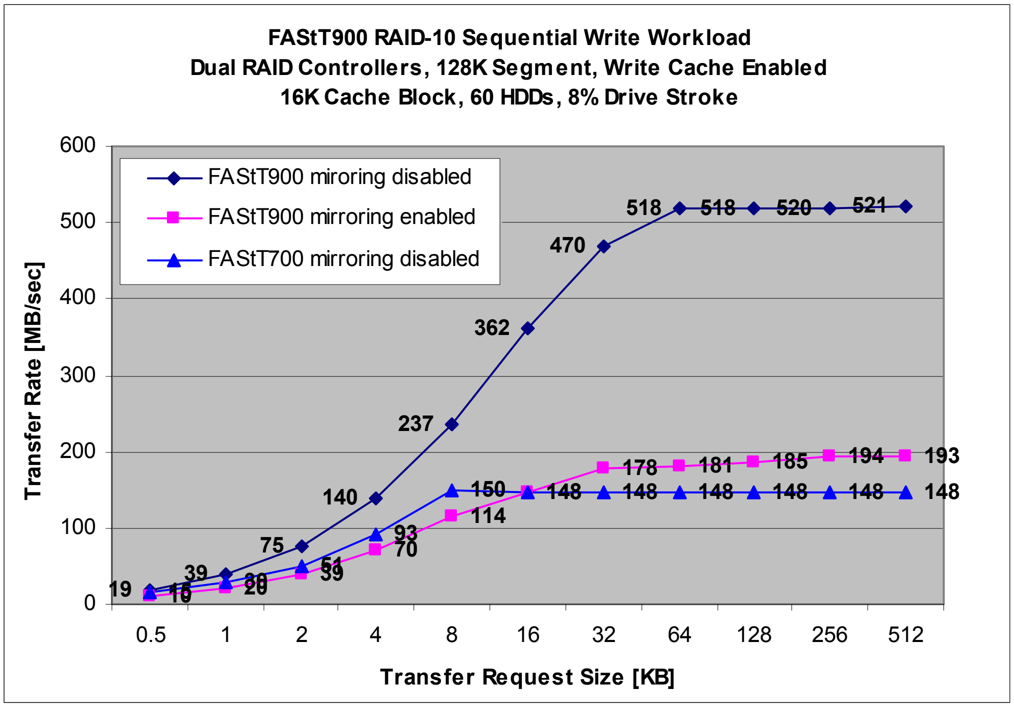

The FAStT700 transfer rate for 64K transfer requests is 148MBps, and the FAStT900 transfer rate for 64K transfer requests is 518MBps. After boosting the FAStT700 performance by 5% because of the older drive technology, the FAStT900 performs approximately 230% better than the FAStT700 for 64K transfers for the RAID-10 Streaming Writes workload.

The FAStT900 Storage Server performs approximately 40% to 230% better than the FAStT700 Storage Server for the RAID-10 Streaming Writes workload over a transfer request range of 512 bytes to 512KB utilizing 60 HDDs.

RAID-10 File Server and Static Web Server Workload Results

Table 7 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.

Table 7. FAStT900 RAID-10 File Server and Static Web Server Transfer Rate*

Workload		File Server [MBps]	Static Web Server [MBps]
RAID-10			
20 HDDs	All Caches Disabled	52	89
	Write Cache Enabled	54	N/A
60 HDDs	All Caches Disabled	155	238
	Write Cache Enabled	156	N/A
100 HDDs	All Caches Disabled	173	246
	Write Cache Enabled	171	N/A
140 HDDs	All Caches Disabled	188	246
	Write Cache Enabled	186	N/A
180 HDDs	All Caches Disabled	189	242
	Write Cache Enabled	191	N/A

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

Chart 10 illustrates a comparison of the FAStT900 and FAStT700 executing the File Server workload. Sampling the data at an average response time of 20 ms yields a transfer rate of 186MBps for the FAStT900 and 163MBps for the FAStT700. The FAStT900 performs approximately 14% better than the FAStT700 for the RAID-10 File Server workload with 140 HDDs. The FAStT700 was measured using the same hardware and software as the FAStT900 for this comparison.

Chart 10. FAStT900 RAID-10 File Server Workload vs. FAStT700 File Server Workload

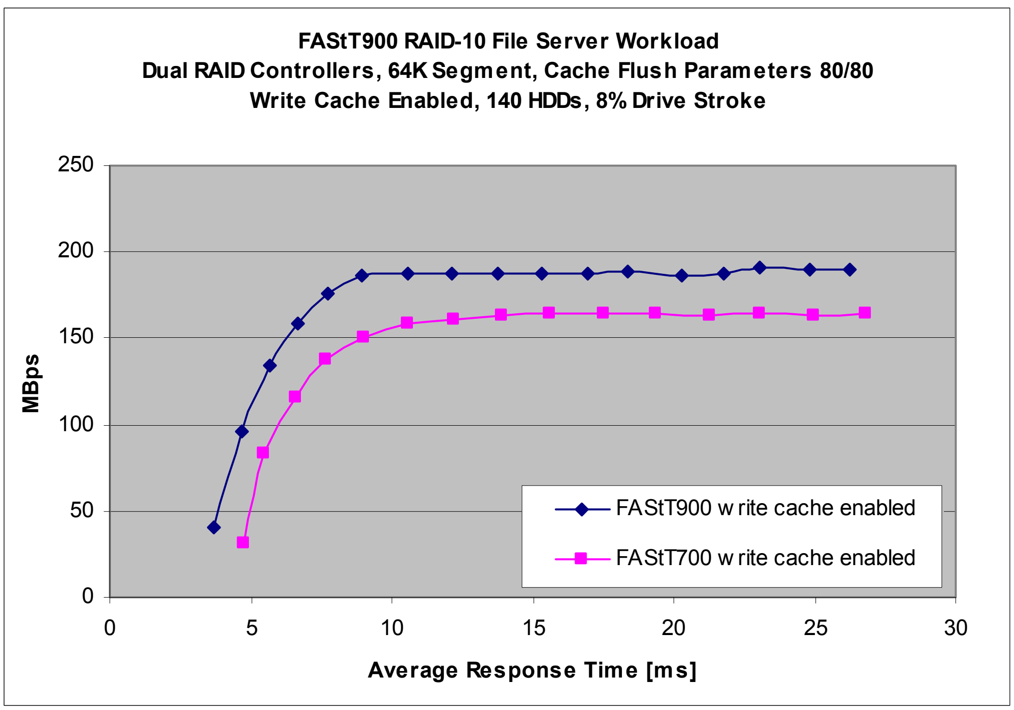

The FAStT900 Storage Server performs approximately 5% to 15% better than the FAStT700 Storage Server for the RAID-10 File Server workload depending on the number of HDDs utilized.

RAID-10 Random Reads and Random Writes Workload Results

Table 8 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 8. FAStT900 RAID-10 Random Reads and Random Writes IOps*

Workload		Random Reads 4K		Random Reads 8K		Random Writes 4K		Random Writes 8K	
RAID-10		IOps	MBps	IOps	MBps	IOps	MBps	IOps	MBps
20 HDDs	All Caches Disabled	8600	33	8300	64	2800	10	2770	21
	Write Cache Enabled	N/A	N/A	N/A	N/A	5895	23	5965	46
60 HDDs	All Caches Disabled	25,170	98	24,330	190	8630	33	8400	65
	Write Cache Enabled	N/A	N/A	N/A	N/A	9360	36	9050	70
100 HDDs	All Caches Disabled	36,430	142	30,260	236	11,000	42	9690	75
	Write Cache Enabled	N/A	N/A	N/A	N/A	12,875	50	11,175	86
140 HDDs	All Caches Disabled	42,500	166	30,060	234	16,245	63	13,900	108
	Write Cache Enabled	N/A	N/A	N/A	N/A	17,170	67	15,960	124
180 HDDs	All Caches Disabled	40,760	159	29,485	230	13,160	51	11,615	90
	Write Cache Enabled	N/A	N/A	N/A	N/A	15,275	59	12,775	99

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

Chart 11 illustrates the performance of the FAStT900 for a RAID-10 Random Writes 8K workload.

Chart 11. FAStT900 RAID-10 Random Writes 8K Workload

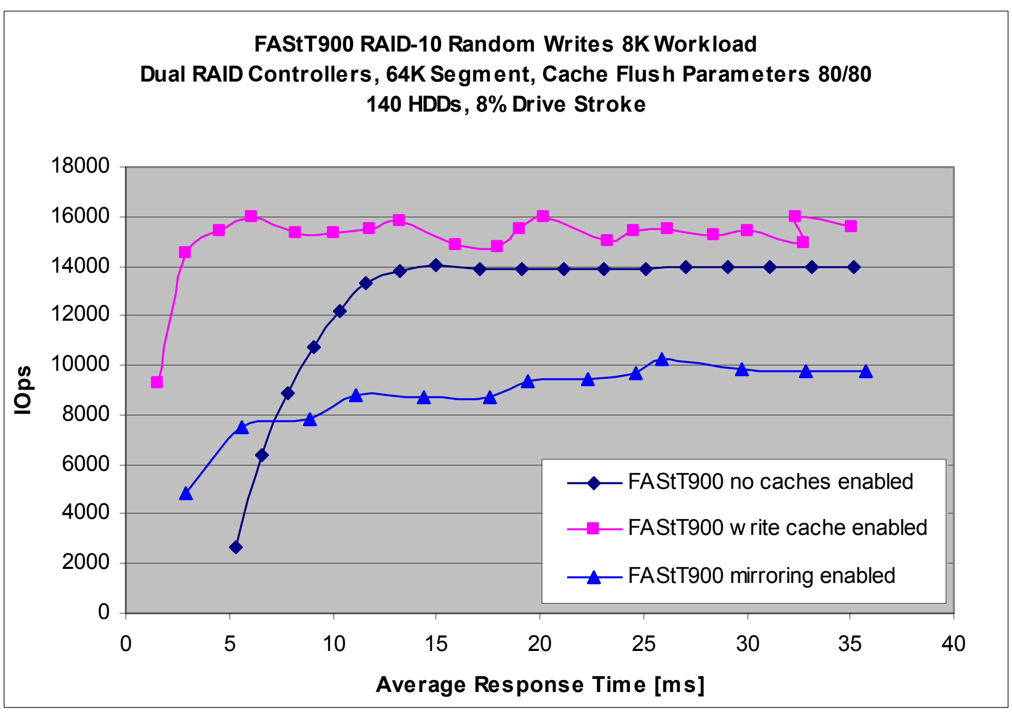

The Random Reads workload can be used to compare the performance of the FAStT900 to the FAStT700. Chart 12 illustrates this comparison.

Chart 12. RAID-10 FAStT900 8K Random Reads vs. FAStT700 8K Random Reads

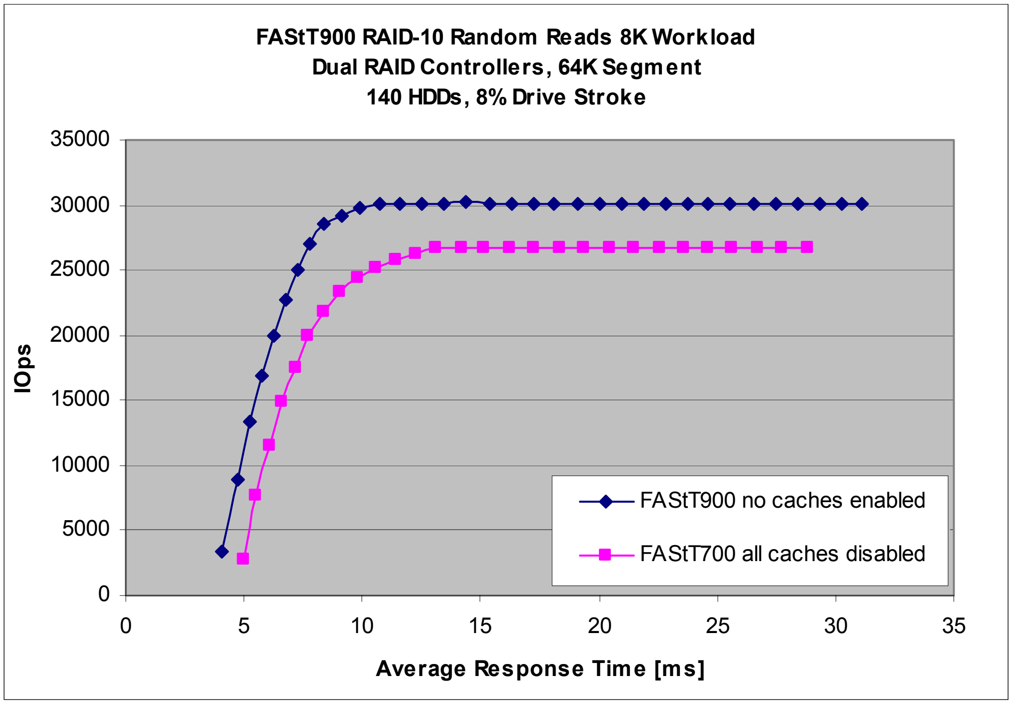

With 140 drives attached, the FAStT900 is capable of achieving approximately 30,060 8K IOps. The FAStT700 is capable of achieving approximately 26,700 8K IOps. Therefore, the FAStT900 performs approximately 13% better than the FAStT700 for the RAID-10 Random Reads workload with 140 HDDs.

The FAStT900 Storage Server performs approximately 5% to 15% better than the FAStT700 Storage Server for the RAID-10 Random Reads 8K workload depending on the number of HDDs utilized.

Summary

The FAStT900 offers substantial performance improvement over its predecessor, the FAStT700. Some examples of the approximate performance gains are:

RAID-5

- 5% to 30% for OLTP 8K workload
- 50% to 75% for Streaming Reads workload (60 HDDs)
- 70% to 120% for Streaming Writes workload (60 HDDs)
- 5% to 15% for File Server workload
- 5% to 15% for Random Reads workload

RAID-10

- 5% to 45% for OLTP 8K workload
- 50% to 85% for Streaming Reads workload (60 HDDs)
- 40% to 230% for Streaming Writes workload (60 HDDs)
- 5% to 15% for File Server workload
- 5% to 15% for Random Reads workload

IBM Systems and Technology Group

Department MX5

Research Triangle Park NC 27709