

TP-525 • From: Seagate Global Product Marketing • May 2004

Technology Paper

Economies of Capacity and Speed: Choosing the most cost-effective disc drive size and RPM to meet IT requirements

Introduction

IT departments are under growing pressure to find the most cost-effective means to meet their escalating data storage requirements, including disc drive capacity and performance. Today's storage systems must hold ever-increasing amounts of data, achieve input/output (I/O) performance goals, meet client response-time expectations and minimize physical space (U or ft3) requirements. Furthermore, IT professionals must ensure their storage systems are reliable-maximizing system availability while minimizing drive failures, the number of RAID rebuilds and RAID rebuild time. All of these factors can come into play when determining the optimal capacity and speed of a system's disc drives.

In this paper we will discuss a variety of factors that impact the choice of disc drive capacity and speed for systems that offer a choice between 3.5-inch 10K-RPM drives and 3.5-inch 15K-RPM drives. Let us start by looking at the current trend in drive purchase decisions. While IT departments often speak of buying drives on the basis of cost/capacity (price per GB), the majority of drives they will actually purchase are 36GB 10K (lowest price box cost), as opposed to 146GB 15K (highest initial acquisition price/GB) or 300GB 10K (lowest price per GB). Were cost/capacity truly the dominant purchase criterion used by IT departments, they would instead be buying 300GB 10K drives (see Figure 1).

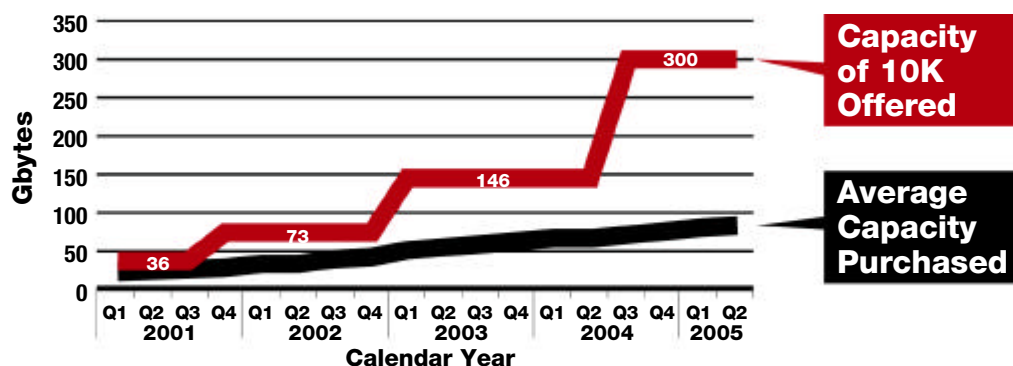


Figure 1. IT departments are not buying the lowest cost/GB drives.

Why have IT departments preferred 10K 36GB drives? The history of disc drives and other system component offerings will shed some light on this.

System Evolution

Over the past two decades enterprise systems have achieved extraordinary gains in performance. Dramatic advances in processor speed, RAM size and RAM speed have combined to accelerate system performance to levels unthinkable just a few years ago. Such powerful hardware resources have made feasible software solutions with increasingly sophisticated and comprehensive capabilities, enabling business productivity to climb at a remarkable rate.

Yet one aspect of system evolution has historically lagged behind: disc drive performance. While impressive advances in areal density have yielded exponential growth in disc drive capacity, disc drive speed has achieved only modest gains over the years (see Figure 2).

	1987	2004	Increase
CPU Performance	1 MIPS	2,000,000 MIPS	2,000,000 x
Memory Size	16 Kbytes	32 Gbytes	2,000,000 x
Memory Performance	100 usec	2 nsec	50,000 x
Disc Drive Capacity	20 Mbytes	300 Gbytes	15,000 x
Disc Drive Performance	60 msec	5.3 msec	11 x

Figure 2. Disc drive performance has not kept pace with other components of the system.

Not only has disc drive performance lagged behind servers, it also hasn't kept up with the dramatic increases in disc drive capacity. These mushrooming capacities mean a drive is expected to do more (access more data) than ever before. For a given environment, there are a given number of operations (Input/Output Operations per Second, or IOPS) accessing each gigabyte of data stored on a disc drive. This is referred to as access density, or IOPS/GB. Every year the capacity offered on a single drive has dramatically increased, yet the number of actuators inside the drive (see Appendix A) that access that additional data has remained the same-one. The drive's performance has not kept pace with the drive's capacity, and as a result, the disc drive's effective performance, or IOPS/GB, has actually decreased over time.

Yet in multi-drive environments, IT departments have been able to meet their performance needs with 10K-RPM disc drives. How has this been accomplished? Rather than buying higher capacity drives that become available with each new generation of drives, many IT departments continue to purchase lower capacity drives (see Figure 1). By selecting lower capacity, 10K 36GB or 73GB drives, they must purchase more drives than if they were to choose 300GB drives. What is the benefit of this strategy? Each 10K 36GB or 73GB drive is required to access less data, thus decreasing the chance that the drive will be a bottleneck in the system. In effect, IT departments are purchasing more actuators to simultaneously access the data, increasing the overall performance of the drives.

COMPARISONS BETWEEN 10K-RPM AND 15K-RPM DISC DRIVES


Server and storage performance significantly increases with 15K-RPM drives. For a typical workload, 15K drives deliver 32% more IOPS and 24% greater response time than the same server configured with 10K drives (see Figure 3). These performance advantages are maintained whether testing 14 drives (as shown) or a single drive.

Drive Comparison	Number of Drives	Performance (IOPS)	Response Time (ms)
10K 146 GB	14	725	22.1
15K 146 GB	14	955	16.7
		32% Faster	24% Faster

All values are calculated based on previous generation drive test results from IBM xSeries 335 with EXP 400:8KB, 2/3 random read, 1/3 random write, queue depth 16.

Figure 3. 10K vs. 15K comparison

15K drives are able to deliver this increased performance by achieving quicker access to data (see Figure 4). 15K drives, of course, spin their internal discs faster, which reduces latency, while the discs' smaller diameter shortens the distance an actuator must move, decreasing seek time (see Appendix B). The smaller discs have the added benefit of yielding power and cooling requirements similar to those of 10K drives.



	10K RPM	15K RPM	15K-RPM Advantage
Latency (msec)	2.99	2.00	33% Faster
Seek Avg R (msec)	4.6	3.3	30% Faster

Figure 4. 10K vs 15K RPM Drive Data Access Time.

The performance 15K drives deliver is particularly significant when one considers that, for the average price of a configured server, 15K drives may represent less than a 10% increase in system cost over 10K drives, yet, as noted above, deliver 32% greater performance and 24% faster response times. This increases productivity and efficiency for every user, paying dividends throughout each year of service operation.

In servers that allow for a single drive or only a handful of drives, lower IOPS/GB drive performance cannot be augmented by using a larger number of drives to simultaneously access data. In such environments 15K drives offer significant risk mitigation benefits to IT managers, reducing the potential for bottlenecks and maintaining fast response times during critical periods of peak activity. Furthermore, 15K drives preserve system investment by ensuring server and storage can keep up with growing performance needs over a longer period of time. For these reasons, 15K drives define the direction storage technology is going.

10K VS 15K RPM DRIVES IN MULTI-DRIVE ENVIRONMENTS

Many enterprise storage systems employ large numbers of 10K-RPM disc drives housed in server racks and/or storage arrays. These deployments utilize advanced interfaces (for example, RAID 5) that enable data to be simultaneously distributed across multiple drives.

In order to attain specific performance goals with 10K drives, it's often necessary to spread data over a multitude of drives, configured to access data simultaneously. While this lessens the extent to which disc drives impose a performance bottleneck in the system, it comes at a considerable cost. Specifically, the expensive infrastructure required to support so many disc drives, combined with the substantial investment in the drives themselves, significantly increases the cost of the overall system. Many IT departments have become accustomed to purchasing 10K 36GB drives, not realizing the benefits (greater performance and capacity) of denser 15K 73GB drives. Or, many IT departments may become accustomed to purchasing 10K 73GB drives, not realizing the benefits (greater performance and capacity) of denser 15K 146GB drives.

For example, as shown in Figure 5, eleven (11) 15K 146GB drives will give the same performance as sixteen (16) 10K 73 GB units, but with 37% more capacity and 31% fewer drives purchased.

Drive Comparison	Number of Drives	Capacity (GB)	Performance (IOPS)	Response Time (ms)	Results With 15K Drives
10K 146 GB	16	2336	830	19.3	Same Performance 31% Fewer Drives
15K 146 GB	11	1606	840	19.1	
10K 73 GB	16	1168	830	19.3	Same Performance 37% More Capacity 31% Fewer Drives
15K 146 GB	11	1606	840	19.1	

All values are calculated based on previous generation drive test results.

Figure 5. 10K vs. 15K comparison

IT departments may choose to increase their 10K drives' performance by storing data only on the outer diameter of the disc, a practice referred to as short-stroking or destroking (see Appendix C). Short-stroking a drive reduces the distance the actuator must move to access the data, and thus improves disc seek time. Of course, such an approach significantly reduces effective disc capacity, increasing the cost per GB of the drive. Even with short-stroking, IT departments would still need to buy more 10K drives

than 15K drives to meet their needs, as shown in Figure 6. For example, an IT manager may typically buy 10K 73GB drives, but elect to purchase 10K 146GB drives and short-stroke them to increase their performance. Nevertheless, for a given capacity and performance level (even when 10K drives are short-stroked), fewer 15K drives than 10K drives are required.

Drive Comparison	Short-Stroked to (GB)	# of Drives	Capacity (GB)	Performance (IOPS)	Response Time (ms)	Results With 15K Drives
10K 300 GB	140 GB	15	2100	930	17.1	Same Performance Same Capacity 7% Fewer Drives
15K 146 GB	Full-stroked	14	2044	965	16.7	
10K 146 GB	110 GB	14	1540	805	19.9	Same Performance Same Capacity 21% Fewer Drives
15K 146 GB	Full-stroked	11	1606	840	19.2	
10K 146 GB	110 GB	18	1980	925	17.3	Same Performance, Same Capacity, 22% Fewer Drives, Fewer Enclosures
15K 146 GB	Full-stroked	14	2044	960	16.7	

All values are calculated based on previous generation drive test results.

Figure 6. Comparison of full-stroked 15K-RPM drives vs. short-stroked 10K-RPM drives

Note that in these lab tests the data was accessed randomly, clearly highlighting the benefit of shorter seek times. But in real-world workloads I/O activity may be localized over a few gigabytes (a small section of the disc radius), making the drive's latency (a function of disc RPM) the primary factor in determining performance. In such cases the slower latency of short-stroked 10K drives significantly undermines performance, requiring even more 10K drives to deliver the same performance and capacity as 15K drives. In short, 15K drives meet IT departments' needs more efficiently.

The net result of deploying 10K drives in typical multi-drive environments is disproportionately expensive, inefficient disc storage systems that still fall short in terms of speed and throughput. 15K-RPM disc drives were specifically designed to address the performance deficiencies of 10K drives, and in doing enable greater cost-effectiveness to be achieved. A multitude of short-stroked 10K drives can be replaced by fewer 15K drives to attain a given level of performance, significantly lowering total cost of ownership (TCO).

Furthermore, 15K-RPM drives offer an additional benefit to IT departments whose requirements may grow: Purchasing 15K-RPM drives up front minimizes the need to buy a whole new system later, as 15K disc drives enable a system to keep up with increasing performance demands over a longer period of time.

Economies of Speed and Capacity

Simply put, the Economies of Speed and Capacity can be defined as the cost savings realized by meeting system performance, capacity and reliability goals with fewer 15K-RPM disc drives.

Fewer drives benefit the customer by enabling:

- Less supporting infrastructure
- Less space required for that infrastructure
- Lower maintenance and storage-management costs
- Higher IOPS/U

Deploying smaller quantities of a single drive type (15K) can reduce the workload on IT staff, while also lowering inventory costs. The cost to maintain a server over a three-year period can be as high as 100,000 USD. The less infrastructure the customer needs, the better (see Figure 7).

	Capacity (GB)	Performance (IOPS)	Drives Required	# of EXP 400s	Space	Reliability (MTBF)
10K 146 GB	2628	925	18	2	6U	67K
15K 146 GB	2044	960	14	1	3U	86K
	Similar	Similar	22% Fewer	50% Fewer	50% Less	28% Greater

All values are calculated based on previous generation drive test results.

Figure 7. 10K vs 15K comparison

Impact of Drive Choice on Overall Reliability

The Seagate 10K-RPM and 15K-RPM disc drives carry the highest reliability rating in the industry. Both have multiple generations of field-proven technology. Both drives offer similar low power and cooling requirements, acoustics and rotational vibration-ensuring dependable performance in high-density, high-demand environments.

Not surprisingly, 15K storage solutions offer greater reliability simply because their lower drive count compared to 10K-RPM configurations means fewer drives in the system to fail. Similarly, fewer drives require less supporting infrastructure, whose failure could cause system downtime. Additionally, high-performance 15K drives decrease RAID rebuild times.

Third-Party Test Results and Applications

15K-RPM drives are designed to excel in a broad range of common applications, from e-mail to business processing to decision support. 10K-RPM drives are better suited for applications where data is infrequently accessed.

Storage Performance Council testing, which includes e-mail in its workloads, verifies the cost efficiency of 15K-RPM disc drives. All top SPC price/performance leaders elected to use 15K drives. (see Figure 8)

Storage Subsystem	Disc Drive	\$ per SPC-1 IOPS™	Total ASU Storage Capacity (GB) Total SPC-1 IOPS™, Data Protection Level, SPC-1 Submission Identifier
IBM Total Storage® FAST 600 with Turbo Option (non-mirrored)	15K RPM	\$9.28	478.43 GB 12102.97 SPC-1 IOPS™, Mirroring A00018
HP StorageWorks (non-mirrored) EVA ZC12D	15K RPM	\$19.99	2596.3 GB 24,006 SPC-1 IOPS™, Non-Mirroring A00021
STK-D280 (mirrored)	15K RPM	\$16.69	1196.092 GB 24507.22 SPC-1 IOPS™, Mirroring A00023

Figure 8. SPC comparison

Disc drive capacity has repeatedly multiplied over recent years, placing increasing demands on drive performance in order to maintain acceptable response times and keep pace with the additional data stored on each drive. Changing from 10K-RPM to 15K-RPM disc drives enables both servers and external storage systems to deliver 30% more performance, while needing fewer 15K-RPM drives to meet performance and capacity requirements. Furthermore, 15K-RPM disc drives have three field-proven generations of reliability, decrease RAID rebuild times and have similar or lower power and cooling requirements, acoustics and rotational vibration compared to 10K-RPM disc drives. For these reasons, the 15K is the drive of choice for the majority of SCSI applications. The bottom line is 15K-RPM disc drives enable higher performance and reduced TCO-key requirements in today's IT departments.

Appendices: The Fundamentals of Disc Drive Performance

APPENDIX A

There is a single actuator (shown in Figure A1) in each drive. This actuator carries multiple heads in order to provide a separate head to read and write data on each side of each disc used in the drive. However, only one head can be active at a time due to limitations in track consistency from disc to disc.




Figure A1.

APPENDIX B

A drive's access time is an essential factor in determining overall drive performance. Access time is defined as the time needed for the read/write head to reach the disc location where the desired data resides. In current storage configurations, this access time is much greater than the time required by the device's electronics to actually read or write the data (referred to as data transfer time). The two major components of a drive's access time are seek time and latency (Figure A2).




Figure A2.

Seek Time: When the drive receives a request, the head moves, or seeks, across the disc to the track containing the desired data. The time it takes for the head to move from its initial position to the data's track is called seek time. Because 15K drives utilize smaller-diameter discs than 10K drives, the distance the head must travel is reduced, resulting in faster seek times.

Latency: After the head is positioned on the appropriate track of the disc, the head must wait on that track until the rotating disc brings the desired sector of data underneath the head. The time it takes for the data to reach the head is called latency time. Of course, the faster the disc is spinning, the shorter the latency time. The total time required for the disc to make one complete revolution is 4.0 msec for a 15K drive (vs. 6.0 msec for a 10K drive).

When a seeking head reaches the target track, the data may already be immediately under the head, resulting in an effective latency of 0 msec. In other cases, where the data has just passed beneath the head, the head must wait a full revolution (4.0 msec) before reading the data. Thus latency on a 15K-RPM drive averages out to 2.0 msec.

Data Transfer Time: For the majority of performance applications, the time required for the electronics to read the data and transfer it to the bus (referred to as the data transfer rate) is negligible compared to the data access time; as such, the data transfer rate has relatively little impact on the drive's performance.

APPENDIX C

Short-Strokeing (or Destroking) 10K-RPM Drives to Improve Access Time: As shown in Figure A4, destroking stores data on the shaded region on the outer diameter of the disc. The lighter area is available data storage that remains unused to keep seek times to a minimum. Such underutilized media results in a very high cost-per-usable-GB and cost-per-IOPS solution for the end user.

15K-RPM disc drives lower access time without forcing users to sacrifice the capacity on the disc. Thus fewer 15K-RPM disc drives of a given capacity are needed to meet the performance requirements of the system.




Figure A3.