

TP555 • From: Global Product Marketing • July 2006

Technology Paper

Seagate Workload Management for Business-Critical Storage

Introduction

Serial ATA (SATA) hard drives have rapidly found favour in many enterprise storage solutions, primarily due to their remarkably low cost-per-gigabyte and solid performance. While fundamentally rooted in the desktop, these budget-priced drives have become popular additions in datacentres. Augmenting mission-critical parallel SCSI and Serial Attached SCSI (SAS) drives in dense storage applications, desktop-class SATA drives deliver maximum capacity in external storage or low-cost server environments.

In an effort to contain storage expenses and work within ever-tighter budgets, IT professionals began specifying these low-cost hard drives in a multitude of enterprise environments. In some cases, they deployed desktop-class SATA drives in more demanding *business-critical* applications for which they were not designed; as a result, some drives experienced unusually high failure rates.

It soon became clear that storage managers needed a new hybrid class of *business-critical* drives, combining the low cost-per-gigabyte of desktop-class SATA drives with the enterprise-class reliability of expensive mission-critical drives.

Seagate® has cooperated closely with customers and key partners to better understand usage environments and customer requirements, in order to design business-critical storage solutions that cost-effectively deliver superior reliability in a variety of high-density storage applications. Investigating the root causes of desktop-class SATA failure rates, Seagate has analysed specific environmental, application and drive characteristics that can contribute to premature failure.

By committing its unparalleled enterprise expertise and technology to analysing the underlying factors that impact SATA reliability in enterprise environments, Seagate has pinpointed **workload** as a key determinant of SATA drive longevity and has applied this conclusion to the new Seagate Barracuda ES family of high-capacity enterprise drives. Incorporating a comprehensive suite of workload management (WLM) tools, Seagate Barracuda ES drives boast the highest reliability of any 7,200-RPM drive in the industry, and are purpose-built for 24x7, business-critical applications.

Challenge: High SATA Failure Rate

Laboratory testing has confirmed the field experience of many IT professionals: When deployed in demanding enterprise applications (with their severe power-on hours and temperatures), desktop-class (PS) drives show more than twice the AFR of enterprise-class (ES) drives (figure 1). Not surprising, given that in such applications, desktop-class drives run hotter, with higher duty cycles and for more power-on hours than in desktop applications.

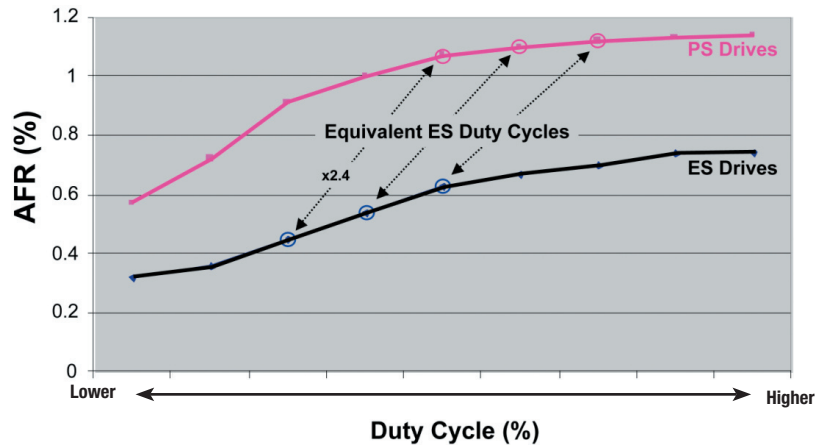


Figure 1. AFR vs. Duty Cycle in Enterprise Application

Enterprise-class drives have been traditionally identified by their interface (FC, SCSI or SAS), and are *architected* and built for IOPS and performance, while desktop-class drives (PATA, SATA) are optimised for single-user, low-load server and consumer electronics duties. In these less-demanding environments, a desktop-class drive need only deliver about half the peak IOPS of an enterprise-class drive.

However, the *bursty* nature of enterprise applications stresses the IOPS capability of desktop drives, ensuring they will see significantly more duty than they were designed for, with deleterious effects on reliability (figure 2).




Figure 2. Annualised Failure Rate (AFR) vs. IOPS

Duty cycle and IOPS are just two important components of a drive's workload. *Workload* refers to the usage pattern of the drive and encompasses a complex set of interacting metrics, including:

- Duty cycle (percentage of time a command is pending)
- IOPS
- Seek-number and seek-length distributions
- Read/Write ratio
- *Burstiness* (low activity punctuated by brief periods of very high activity)

With far lower IOPS than high-performance SAS and FC drives, desktop-class SATA drives are forced to work harder and longer to execute the same drive commands. This means higher duty cycles—and higher temperatures. Heat is the enemy of drive reliability, and business-critical workloads are a key contributor to elevated drive temperatures and higher failure rates.

Moreover, business-critical environments present additional thermal challenges. Their dense servers, large storage farms and rack-based data centres impair airflow and chassis cooling, which in turn raises drive temperatures still higher. Faster, hotter-running processors only exacerbate the problem.

Against this backdrop of significant reliability issues, business-critical hard drives are being developed and optimised for server and external storage (NAS, SAN, DAS RAID) solutions. Poised between mission-critical enterprise drives and desktop-class drives, this new breed of purpose-built SATA drives must deliver superior reliability if it is to succeed in the enterprise marketplace.

Solution: Seagate Workload Management

Business-critical hard drives are designed to be the workhorse infrastructure solution for server and storage environments. Higher capacity and lower cost are the hallmarks of these innovative drives, but enhanced reliability in enterprise applications is the foundation of their value proposition. Workload management technology makes this possible.

Workload management (WLM) is an optimised monitoring and management tool that tracks hard drive temperature and activity in business-critical environments. The significance of WLM is multifold:

- WLM helps optimise hard drive reliability in server and storage environments.
- WLM ensures business-critical drives do not overheat when workload spikes, ensuring higher reliability and longer life.
- WLM is not available in standard desktop products, where workload is not as demanding.
- WLM is a key new feature offered in business-critical SATA products.

Seagate Barracuda ES drives are specifically designed to address the needs of business-critical applications: high capacity at low cost-per-gigabyte, speedy performance and outstanding reliability. To achieve the latter, the Seagate Barracuda ES SATA drive employs an innovative WLM feature, **read after write (RAW)**.

RAW is enabled if the drive's temperature is less than 18°C* or greater than 58°C*. A Write command will be converted to a Write-Verify, which follows the error recovery path for writes and reads.

If the write portion of the Write-Verify fails, the drive performs write recovery and posts a write error if unsuccessful. Should the verify portion of the Write-Verify fail, the Write-Verify is retried. If the verify still fails, a write error is reported to the host as a 03/0C00/11.

The FC version of the Seagate Barracuda ES drive offers these additional WLM features:

1/3 STROKE IOPS LIMITER

IOPS performance is limited for low Q depths where high duty-cycle seeks are present.

LOW POWER MODE (LPM)

This feature disallows faster JIT (just-in-time) seek options when the temperature of the drive exceeds the specified trip temperature. LPM reduces heat generation and helps prevent temperatures from rising even higher, thus increasing reliability. Higher temperatures cause drive heads to fly lower; the combination of lower fly height and fast seeks greatly increases the risk of head/disc contact.

SERVO DEMAND ESTIMATOR MEASUREMENT (SDEM)

SDEM is a measurement of the voice coil motor's (VCM) resistance, which is used to determine the VCM temperature. When the VCM threshold temperature is reached, less current is driven through the coil, thus reducing its temperature. This enhances drive reliability by preventing the drive from operating outside its specified temperature range.

Conclusion

The reliability issues of desktop-class SATA drives are elegantly remedied with workload management. By challenging the status quo with core technologies, such as WLM, Seagate enables storage managers to confidently embrace high-capacity SATA drives as viable business-critical enterprise storage solutions.

Seagate has applied WLM to its new line of business-critical Seagate Barracuda ES drives, retaining the many benefits of SATA drives while significantly improving their core thermal capabilities. And WLM is only the first of many enhancements coming from the leader in enterprise storage solutions.

*Default set points, subject to change.