

Advanced Format Technology Brief

Table of Contents

What is Advanced Format?	1
How Does Advanced Format Technology Work?	1
How Does Advanced Format Technology Maintain Compatibility?	1
512-byte Read	2
512-byte Write (Read-Modify-Write)	2
How Does Advanced Format Technology Maintain Performance?	2
Partition Alignment	3
Advanced Format Aware Operating Systems	3
Legacy Operating Systems	4
Using a Partition Alignment Tool	4
Partitioning Software	4
Additional Sources of Information	4

What is Advanced Format?

Advanced Format is a new technology for increasing hard drive capacities while maintaining data integrity. In order to create higher capacity drives, manufacturers have begun the transition to Advanced Format. HGST is now introducing new, high capacity drives with Advanced Format.



Advanced Format drives incorporate several changes to optimize the actual data structure on the hard drive, including increasing the physical sector size from 512 bytes to a more efficient 4096 (4K) byte sector size. In order to ease the transition, current Advanced Format drives provide 512-byte emulation (512e)¹ at the drive interface for backward compatibility with legacy applications.

Most modern operating systems have been designed to work efficiently with Advanced Format drives. For optimum performance, it is important to ensure that the drive is partitioned correctly and that data is written in 4K blocks by both the operating system and the application. Recent operating systems handle this automatically. When using Advanced Format drives in legacy environments, special tools may be required to optimize read/write performance.

How Does Advanced Format Technology Work?

Historically, hard drives have stored data in 512-byte sector sizes (see Figure 1: Advanced Format Layout). In addition to the user's data, overhead data is written on the disk for each sector, including the error correction code (ECC) and drive format information (Gap, Sync, Data Address Mark).

Advanced Format drives use longer sectors that contain 4096 (4K) bytes. This is the equivalent of putting eight historical (512-byte) sectors into one new 4K sector. This approach provides two benefits illustrated in Figure 1. First, by optimizing the overhead associated with each smaller sector, the drive uses less space to store the same amount of information resulting in a format efficiency improvement (see blue arrow in Figure 1). The second benefit is that a larger and more powerful error correction code (ECC) can be utilized, providing better integrity of user data.

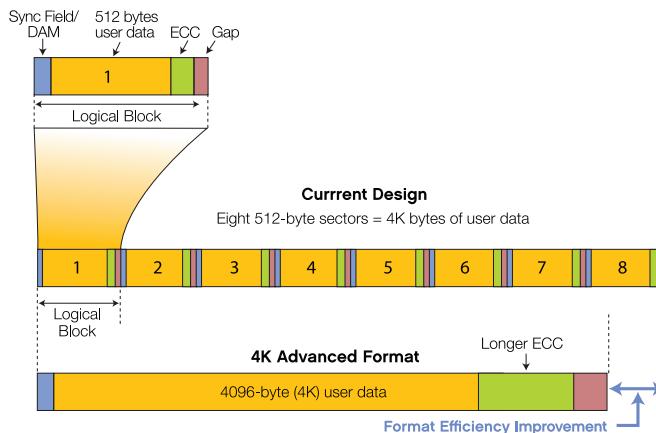
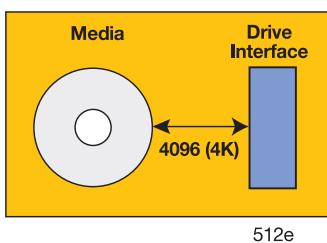


Figure 1: Advanced Format Layout

How Does Advanced Format Technology Maintain Compatibility?

Many existing hardware and software components were designed around 512-byte sector (or block) sizes and expect data to be sent and received in 512-byte segments. In order to maintain compatibility with these legacy applications, the hard drive industry will provide 512-byte emulation at the drive interface. 512e hard drives will transparently map all 512-byte logical blocks into the drive's 4K physical sectors (see Figure 2: Software Stack).

512e Hard Drive



Host System

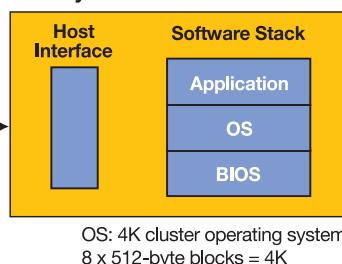


Figure 2: Software Stack (Potential areas using 512-based code)

¹ 512 byte emulation is sometimes referred to as 512e

512-byte Read

When the host requests to read a single 512-byte logical block, the hard drive will actually read the entire 4K physical sector containing the requested 4K block containing the 512 bytes. The 512-byte is extracted and sent to the host. This can be done very quickly.

For example (see Figure 3), a single read of Logical Block #3 will cause Physical Block #1 to be read, and Block #3 is extracted and sent quickly to the host.

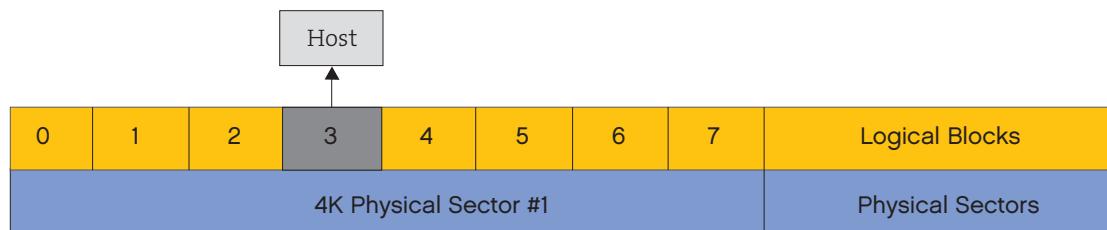
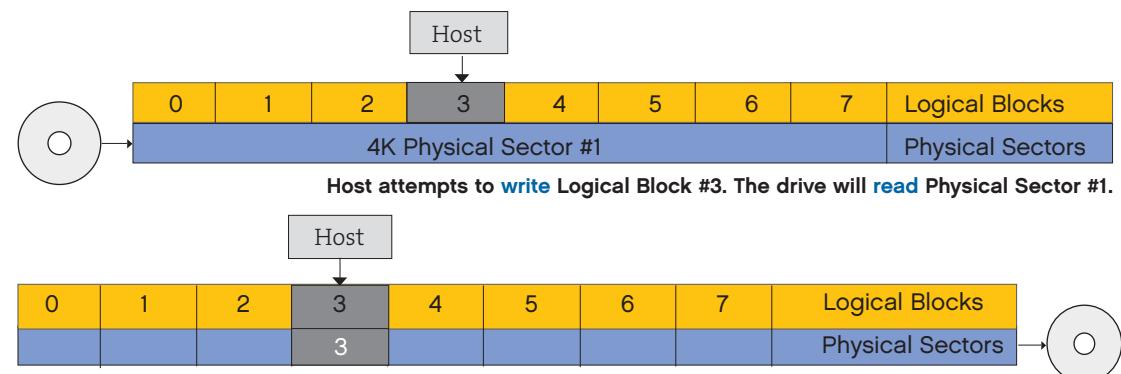


Figure 3: 8 x 512-byte Logical Blocks Mapped to 1 x Physical Sector

512-byte Write (Read-Modify-Write)

When the host attempts to write a single 512-byte logical block, the hard drive will first read the 4K physical sector containing the 512 bytes that are to be overwritten. Next, it will insert the 512 bytes of new data and write the entire 4K block of data back to the media. This process is called a “Read-Modify-Write”. The drive must read the existing data, modify a subset, and then write the data back to the disk. This process can require additional revolutions of the hard disk.

For example (see Figure 3), a single read of Logical Block #3 will cause Physical Block #1 to be read, and Block #3 is extracted and sent quickly to the host.



How Does Advanced Format Technology Maintain Performance?

In order to maintain top performance, it is important to ensure that writes to the disk are aligned. Ideally, writes should be done in 4K blocks, and each block will then be written to a physical 4K sector on the drive. This can be accomplished by ensuring that the OS and applications write data in 4K blocks, and that the drive is partitioned correctly.

Most modern operating systems use a file system that allocates storage in 4K blocks or clusters. In a traditional hard drive, the 4K block is made up of eight 512-byte sectors (see Figure 4: 512-byte Emulated Device Sector Size)

4K block = 8 x 512-byte blocks								OS File System	
0	1	2	3	4	5	6	7	Logical Blocks	
4K Physical Sector #1								Physical Sectors	

Figure 4: 512-byte Emulated Device Sector Size

Logical Blocks [0 to 7]								Logical Blocks [8 to 15]							
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
4K Physical Sector #1								4K Physical Sector #2							

Figure 5: 512-byte Emulated Device Sector Size—Two “aligned” sectors

Partition Alignment

Since most modern operating systems will write in 4K blocks, it is important that each 4K logical block is aligned to a physical 4K block on the disk (see Figure 5). This is especially important because the 512e feature of the drive cannot prevent a partitioning utility from creating a misaligned partition. When misalignment occurs, a logical 4K block will reside on two physical sectors. In this case, a single read or write of a 4K block will result in a read/write of two physical sectors. The impact of a “read” is minimal, whereas a single write will cause two “Read-Modify-Writes” to occur, potentially impacting performance (see Figure 6)..

4K logical block = 8 x 512-byte blocks [4:11]								4K logical block = 8 x 512-byte blocks [4:11]							
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
4K Physical Sector #1								4K Physical Sector #2							

Figure 6: 512-byte Emulated Device Sector Size (MISALIGNED)

If a partitioning or disk imaging utility is used, the partitions must be aligned such that the logical partition starts on a physical 4K sector boundary. For a logical block size of 512 bytes, this can be accomplished by ensuring that the partition starts on a multiple of 8. In the “misaligned” example above, the 4K block starts at logical block #4, which does not align with a physical 4K sector (logical block #0 or #8).

Advanced Format Aware Operating Systems

Recent operating systems have used 4K file systems for some time. These operating systems will naturally use 4K blocks.

Operating System	4K file systems	Automatically aligns partitions during installation
Microsoft® Windows® Vista SP1 or later	Yes	Yes
Microsoft Windows 7	Yes	Yes
Microsoft Server 2008	Yes	Yes
Mac OS® X 10.4+	Yes	Yes
Linux Ubuntu 8.04+, SUSE, Linux kernel 2.6.34+	Yes	Use Linux Partitioning Utility

Even when using a 4K operating system, it is possible for individual applications to ignore OS settings and attempt to write in 512-byte blocks. With 512-byte emulation, these applications will continue to work. However, for optimum performance, applications should work in conjunction with the operating system and use 4K blocks when writing to the file system.

Legacy Operating Systems

Operating System	4K file systems	Automatically aligns partitions during installation
Microsoft Windows XP	Yes	No
Microsoft Server 2003	Yes	No
Microsoft Windows Home Server V1	Yes	No

Using a Partition Alignment Tool

For best performance, it is recommended that the hard drive be partitioned correctly upon first use. By creating aligned partitions, a single 4K block read or write will result in a single physical sector read or write on the hard drive.

For users who have created a misaligned partition, HGST recommends using a disk partitioning utility to correct this condition by re-aligning the partition.

Advanced Format Aware Partitioning Software

Microsoft Windows OS installations have built-in partitioning software. Please see table above for compatibility.

Many third party imaging and partitioning utilities can correctly create aligned partitions. Several options are listed below.

Partitioning Software:

Linux

GPARTED

Gnu Parted 2.1+ use “-a optimal” or “-a minimal” options

<http://www.gnu.org/software/parted/manual/parted.html>

Windows-Based

Acronis Disk Director Home 11

Paragon Partition Manager 11

Mac-Based

Disk Utility (provided with Mac OS X 10.4+), creates a GPT (GUID Partition Table) partition.

Additional sources of information:

www.idema.org (Advanced Format Documents, including specifications)

www.en.wikipedia.org/wiki/Advanced_Format

The ATA9-ACS and SBC-3 standards have provisions for a disk drive to report Advanced Format sector sizes and other performance optimization information. These standards are used for SATA, SAS, USB, and IEEE 1394 based interface technologies.

www.t13.org (Document: ATA8-ACS (ATA Command Set)

www.t10.org (Document: SBC-3 (SCSI Block Commands)

Microsoft Windows 7 and Advanced Format Hotfix (KB981208)

<http://support.microsoft.com/kb/982018>

Intel Rapid Storage Technology (RST)

<http://www.intel.com/support/chipsets/imsm/sb/CS-031502.htm>



HGST trademarks are intended and authorized for use only in countries and jurisdictions in which HGST has obtained the rights to use, market and advertise the brand. Contact HGST for additional information. HGST shall not be liable to third parties for unauthorized use of this document or unauthorized use of its trademarks.

References in this publication to HGST's products, programs, or services do not imply that HGST intends to make these available in all countries in which it operates.

Product specifications provided are sample specifications and do not constitute a warranty. Information is true as of the date of publication and is subject to change. Actual specifications for unique part numbers may vary. Please visit the Support section of our website, www.hgst.com/support, for additional information on product specifications. Photographs may show design models.

© 2014 HGST, Inc.. All rights reserved.

HGST, Inc.
3403 Yerba Buena Road
San Jose, CA 95135 USA

Produced in the United States 110/10. Rev. 12/11, 3/13, 3/14.

Other trademarks are the property of their respective companies. Microsoft, Windows and Windows Vista are trademarks or registered trademarks of Microsoft Corporation in the US and/or other countries. Mac and Mac OS are trademarks of Apple, Inc. Other trademarks are property of their respective companies.