



Technology Paper

## Deploying the Next Generation of High-Capacity Hard Drives

**Overcoming Legacy Storage Architecture Limits With Hard Drives Larger Than 2.1TB**

### A Brief History of Storage Architectures

It has been 30 years since a small software company made an agreement with IBM to produce the first operating system for what would become the IBM PC. That company was, of course, Microsoft and the operating system was DOS or PC-DOS, as it became branded by IBM. DOS became the basis for the operating system that would be broadly adopted as the standard in the PC industry and would also serve as the foundation for many aspects of the Microsoft Windows operating system.

In those early days, no one really envisioned the incredible advancements in the computer industry that would ultimately lead to today's super fast processors, massive solid state memory and huge hard drives, now commonly available in sizes up to 2TB. Some of the basic design decisions made in the original computer architecture, both hardware and software, left the industry with inherent limitations. One of those limitations is impacting the storage industry today: the ability to address hard drives that exceed 2.1TB capacities.

One of the most fundamental elements of storage architecture is the sector. A sector is the smallest physical block of data represented on a hard drive. Over 30 years ago, the sector was defined as 512 bytes in length and that definition has persisted to this day. Today there is a move to transition to a larger 4K sector (also called Advanced Sector Format), but this transition will not be implemented at the host level (hard drive controllers, OS, BIOS, etc.) for a number of years. In the meantime, the storage industry still must contend with a base sector size of 512 bytes even with the growing need for more capacity.

It turns out that another key architectural decision was made related to hard drive sectors. This was related to the space set aside to address sectors. Each sector is assigned a unique address that defines where the data is located on a hard drive. This is called a logical block address (LBA). Back in the late 1970s and early

# Deploying the Next Generation of High-Capacity Hard Drives



1980s, no one could fathom a hard drive even approaching 1TB, so limiting the LBA range to 2.1TB was thought to be more than enough. As a result, operating systems, BIOS controllers, HDD controllers and device drivers have used the same basic limitation of 2.1TB for the maximum size of a hard drive or logical storage device.

Now that the industry is on the verge of releasing hard drives that exceed this 2.1TB capacity limitation, attention is being focused on how to deploy this new generation of high-capacity storage solutions in light of these historical architectural design limits. It is best to do this in the context of looking at two distinct market segments—the desktop PC and workstation segment, and the enterprise segment focusing on multi-drive storage arrays and servers.

## The Desktop Market

In evaluating how to deploy hard drives larger than 2.1TB in a desktop environment, three distinct installation conditions must be considered:

1. Deployment of any hard drive with a native capacity larger than 2.1TB using Microsoft Windows XP or older, or a using another operating system that is not capable of long LBA addressing
2. Deployment of a non-bootable hard drive with a native capacity larger than 2.1TB using and operating system that is capable of long LBA addressing
3. Deployment of a bootable drive where native drive size larger than 2.1TB

### Condition One

Unfortunately, there is no solution to successfully utilize hard drives larger than 2.1TB using any operating system which does not support long logical block addresses, often referred to as long LBA addressing. This includes Windows XP, which still has a sizable installed base.

Long LBA addressing extends the number of bytes used in a Command Descriptor Block (CDB) to allow access to an LBA range that exceeds the 2.1TB limitation. A CDB is simply a data structure that is used to format data passed between host computers and hard drives. A key element of this data is the LBA, which tells the hard drive which specific sector of data is being addressed.

Operating systems without this long LBA addressing capability cannot recognize a hard drive larger than 2.1TB. In fact, using a hard drive larger than 2.1TB in a Windows XP system can produce unpredictable results, depending on the BIOS used. It may recognize only the capacity up to the 2.1TB limit or it may in fact only recognize the capacity over the 2.1TB limit. So, a 2.5TB hard drive might be recognized by Windows XP as a 400GB drive or a 2.1TB drive, but not correctly as a 2.5TB drive. Using hard drives with a native capacity larger than 2.1TB requires Windows Vista, Windows 7 or another long LBA-capable OS.

### Condition Two

We have now established the basic operating system requirement of support for long LBA addressing in order to use a hard drive larger than 2.1TB. If the desired use of this drive is as a non-bootable drive, only one more requirement remains—the drive must be partitioned using a GUID Partition Table (GPT).

A GPT was defined as part of a more comprehensive specification originally launched by Intel with the purpose of creating a replacement to the original PC BIOS design. This standard was named the Extensible Firmware Interface (EFI) and the specification is now managed by Unified EFI Forum (UEFI). In addition to other enhancements, a GPT allows for a much larger LBA addressing scheme, which enables

# Deploying the Next Generation of High-Capacity Hard Drives



the use of hard drives in excess of 2.1TB in size. Furthermore, the legacy standard MBR (Master Boot Record) will result in the same limitations discussed earlier and limits the hard drive capacity to 2.1TB.

## Condition Three

Clearly, the use of a single hard drive in a desktop PC configuration is very common. In this case, the hard drive must be bootable. Using a hard drive larger than 2.1TB for a bootable configuration has the same requirements as described above in condition 2, with two additional requirements.

First, in order to overcome problems in deploying hard drives larger than 2.1TB (for bootable drives), it is important to look at the legacy PC BIOS. The PC BIOS standard that has been used for many years also has the same fundamental limitation in not being able to address a hard drive larger than 2.1TB. Today the only solution around this limitation is the use of a new BIOS standard, also part of the work of the UEFI Forum. The UEFI BIOS has a number of extensions that go beyond the capabilities of the older PC BIOS system. Among these is the capability to address hard drives larger than 2.1TB. To date, however, the use of a UEFI BIOS in the desktop market is rare. Nevertheless, with the introduction of a new generation of high-capacity hard drives, adoption of UEFI BIOS systems will accelerate.

The second requirement for deployment of a hard drive larger than 2.1TB is the use of a compatible hard drive controller driver. One of the most popular drivers used in the PC marketplace is the Intel Matrix Storage driver. This driver, which is commonly shipped as a component of the Windows Vista and Windows 7 OS configuration, also shares the limitation of not properly addressing hard drives larger than 2.1TB. Intel is planning an update to this driver in the near

future. In the meantime, the standard Windows HDD driver included in the Windows Vista and Windows 7 installation solves this issue.

## Servers and Storage Arrays

Fundamentally, the requirements described for the standard desktop market in use of hard drives larger than 2.1TB also apply to the server and storage array markets. The most significant difference in these two markets is the dominant use of Host Bus Adapters (HBAs) or RAID controllers to manage and control multiple hard drives, common in enterprise solutions.

The use of an HBA or RAID controller simplifies the adoption of these new larger hard drives as the addressing and control issues are essentially all managed by the HBA or RAID controller. The details of the sizes and types of drives behind the controller are, unfortunately, hidden from the OS.

Furthermore, these advanced storage solutions will typically utilize a very small hard drive as the primary boot drive while depending on the HBA or RAID controller to manage the remaining hard drives installed as part of the total solution.

Seagate recommends checking with your RAID or HBA controller manufacturer to explore the use of hard drives larger than 2.1TB, ensuring they are compatible. Several controller manufacturers are limiting this support to the latest 6Gb/s SAS and SATA controllers which, in turn, may require an upgrade if you are currently using a 3Gb/s controller.

If you choose to run a hard drive larger than 2.1TB off of an existing legacy controller without long LBA- addressing capability, then you must *mode select* the drive to de-stroke the drive capacity to match the capabilities of your legacy controller.

## Summary

There are near-term obstacles to overcome to allow broad-based adoption of hard drives larger than 2.1TB in the desktop environment, particularly in the case of bootable drives. The adoption of these new high-capacity drives in the enterprise market has fewer, simpler limitations. While these conditions may seem complex in the context of a detailed explanation, a concise summary is shown in Table 1.

Second Drive (non-bootable) Solution Desktop, Workstation	
Requirement	Explanation
Long LBA-capable Operating System (Windows Vista, Windows 7 or modified Linux)	XP is not capable of >2.1TB
GUID Partition Table (GPT) required	Legacy Master Boot Record (MBR) partitions are limited to 2.1TB in size.

  

Primary Drive (bootable) Solution Desktop, Workstation	
Requirement	Explanation
Long LBA-capable Operating System (Windows Vista, Windows 7 or modified Linux)	XP is not capable of >2.1TB
GUID Partition Table (GPT) required	Legacy Master Boot Record (MBR) partitions are limited to 2.1TB in size.
A Unified Extensible Firmware Interface (UEFI) BIOS is required for bootable drives >2.1TB	Older BIOS do not include support for drives >2.1TB. Updated BIOS with UEFI support for this purpose are in development.
HDD driver support for >2.1TB hard drives	Updates to standard HDD drivers are being modified and tested. The standard Windows driver is ready and Intel is testing modifications.

  

Storage Server or Storage Array Solutions	
Requirement	Explanation
Confirmed HBA or RAID controller support for long LBA Addressing	Controllers manage the HDD addressing
Mode Select to de-stroke drive for use with legacy	Legacy controllers do not support long-LBA Addressing

**Table 1. Requirements Summary**

# Deploying the Next Generation of High-Capacity Hard Drives



## Call to Action

Even today's small businesses need terabytes of storage. As storage demands continue to grow and the storage market continues to make improvements in capacity, performance and cost efficiencies, larger hard drives will continue to be a vital part of these improvements. This is particularly true in the enterprise markets where data centers and cloud computing readily adopt high capacities where they can be deployed quickly with significant leverage in driving operational efficiencies and reducing storage costs. The next generation of large hard drives, extending past the 2.1TB limit for legacy systems will begin shipping in the summer or fall of 2010. In order to minimize integration complexities and to best utilize these new larger hard drives that benefit customers, system builders, integrators and storage solution providers need to act now.

If your product roadmaps focus on storage arrays and servers that depend on RAID controllers and HDD host bus adapters, contact your vendor and ask what their plans are for high-capacity HDD support above 2.1TB.

If the solutions you develop are more focused on desktop, workstation or server solutions that depend on motherboard-equipped HDD controllers and OS support, you should begin having these conversations with your component system suppliers. Intel and Microsoft are key enablers and are already engaged in the discussions to solve these limitations. However, your voice is important in establishing a sense of urgency for these suppliers, as well as the broader ecosystem of suppliers who can help implement the solutions to overcome these legacy limitations for HDD support larger than 2.1TB.

Working together, we can continue to bring high-capacity, high-performance, computing and storage solutions to the market and continue to drive productivity and economic benefits to our industry.

## For More Information

Read the FAQs in the Seagate document, *High-Capacity Storage Readiness (>2.1TB): Frequently Asked Questions*, MB603.

AMERICAS Seagate Technology LLC 920 Disc Drive, Scotts Valley, California 95066, United States, 831-438-6550  
ASIA/PACIFIC Seagate Technology International Ltd. 7000 Ang Mo Kio Avenue 5, Singapore 569877, 65-6485-3888  
EUROPE, MIDDLE EAST AND AFRICA Seagate Technology SAS 16-18, rue du Dôme, 92773, Boulogne-Billancourt Cedex, France 33 1-4186 10 00

© 2010 Seagate Technology LLC. All rights reserved. Printed in USA. Seagate, Seagate Technology and the Wave logo are registered trademarks of Seagate Technology LLC in the United States and/or other countries. All other trademarks or registered trademarks are the property of their respective owners. When referring to drive capacity, one gigabyte, or GB, equals one billion bytes and one terabyte, or TB, equals one trillion bytes. Your computer's operating system may use a different standard of measurement and report a lower capacity. In addition, some of the listed capacity is used for formatting and other functions, and thus will not be available for data storage. Seagate reserves the right to change, without notice, product offerings or specifications. TP614.1-1005US, May 2010