

Introduction of High-Capacity Hard Drives

While manufacturers begin introducing the latest generation of high-capacity drives, the industry has been working to resolve known capacity addressing limitations that date back to original design of the PC. These limitations are caused by 32-bit definitions used for partition size and logical block addresses (LBA)¹. Now that Hitachi and others offer hard drives with capacity greater than (\geq) 2.2 Terabytes (TB), users will need to implement a combination of operating systems, BIOS, drive partition tables, and Hard Disk Drive (HDD) drivers that support $\geq 2.2\text{TB}$ in order to gain trouble-free access to the higher capacities offered. When working in legacy environments, more effort will be required to ensure that all components are capable of supporting high-capacity drives.

The purpose of this technology brief is to help you understand what is required to successfully choose components that work with high-capacity drives.

PC Configurations

In a typical PC (see Figure 1: PC Configurations), a drive may be connected as the primary boot drive, or a data drive. A data drive could be a second internal SATA drive, or an external drive like a USB device. The boot drive will contain the operating system such as Microsoft® Windows®, Linux™, or Mac® OS® X. The operating system is required in order to start the computer from the hard drive. The data drive is a secondary drive and does not contain an operating system.

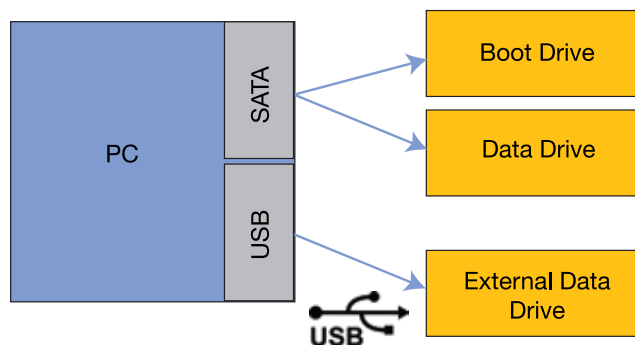


Figure 1: PC Configurations

Data Drive

A data drive greater than or equal to 2.2TB requires an operating system that supports long LBA addressing, such as Windows 7, a GUID partition table (GPT) and HDD drivers that support $\geq 2.2\text{TB}$ drives. More information can be found in the Component Details section.

Boot Drive

In addition to the requirements for a data drive, a Windows boot drive requires a UEFI BIOS (Unified Extensible Firmware Interface) and a 64-bit version of Windows Vista or Windows 7. Linux does not require a UEFI BIOS. More information can be found in the Component Details section below.

USB Drive

USB drive enclosures must be configured to work with high-capacity drives so it is important to confirm with the manufacturer that the enclosure will support capacities greater than 2.2TB. The enclosure may require a driver or firmware update to support the larger capacity.

Component Details

Partition Type

In order to store data on a hard drive, it must first be partitioned. A partition table tells the computer how data on the drive is laid out, including the size of the drive. Historically, PCs have partitioned the hard drive using a Master Boot Record (MBR). Unfortunately, the MBR is limited to partitions smaller than 2.2TB.

¹ 32 bit limitation: 512 byte sectors x 2^{32} bits = 2.199 TB

To address this problem, a new partition table scheme was created, called the GUID Partition Table (GPT). With the introduction of Windows Vista®, Microsoft added support for GPT. Using a GPT table to tell the computer how data is laid out on the disk drive, the computer can access data beyond 2.2TB. Microsoft also enabled support of GPT partitions with Windows XP 64-bit Service Pack 2. However, 32-bit versions of Windows XP do not support GPT partitions. Other operating systems that can handle GPT-partitioned drives include Windows 7, Linux and Intel-based Mac OS X systems.

Using a GPT-partitioned hard drive will ensure that your data remains intact in a legacy MBR system. GPT-partitioned drives contain a special “protective MBR”, designed to protect your data from legacy software that is not aware of GPT. The protective MBR is located in the first sector and will make the drive appear as a single, non-useable partition to legacy software and operating systems.

BIOS

When powering up the PC, the PC BIOS will search for the boot drive and boot the operating system (OS) contained on that drive. For this process to work, the traditional PC BIOS relies on data stored in the partition table (i.e. the MBR) in order to boot the PC. When GPT was created, a new boot method was defined. This method was included in the UEFI BIOS standard. To boot from drives that use the GPT technology, a host PC also needs to be running the UEFI BIOS. If you attempt to boot a GPT-partitioned drive with the traditional PC BIOS, the BIOS will see a single, non-bootable partition, and will not attempt to boot from that drive. The UEFI BIOS replaces the PC BIOS, and is managed by the Unified EFI Forum (UEFI). Linux boot loaders can get around this limitation, and do not require a UEFI BIOS.

Operating Systems and HDD device drivers

Most recent operating systems can correctly handle high-capacity hard drives. In some cases, the storage drivers must be updated within the OS. As of this writing (October, 2010), Intel-based SATA controllers must use the latest device drivers to address data beyond 2.2TB. Microsoft currently ships an HDD driver with Windows 7 and Windows Vista that addresses the issue. The expectation is that Intel will update their device driver in the near future to support high capacity drives. You should verify high capacity drive support with your vendor before using any storage device drivers. The OS must support GPT partitions, and when booting from the high-capacity drive, must also support a UEFI BIOS.

Microsoft Windows:

Windows Vista² and Windows 7 provide support for partitions that use GPT. However, when creating a boot drive, only the 64-bit versions support the UEFI BIOS, which is required for booting.

Windows XP (32-bit) does not support GPT, nor does it support UEFI BIOS. Therefore, Windows XP cannot support high-capacity hard drives through traditional methods.

Apple Mac:

Intel-based Macs, running OS X, support both boot and data drives. Apple uses an EFI-based³ BIOS and OS X supports GPT partition tables.

Linux:

Linux-based operating systems are more flexible in working with high-capacity drives. Common Linux distributions, including Ubuntu and SUSE (32- & 64-bit), can access high-capacity drives. The Linux boot loaders, GRUB 2 and patched versions of GRUB Legacy, support GPT partitions, as do the latest versions of popular disk partitioning utilities, such as GNU Parted or GParted.

High-Capacity-Aware Operating Systems Summary

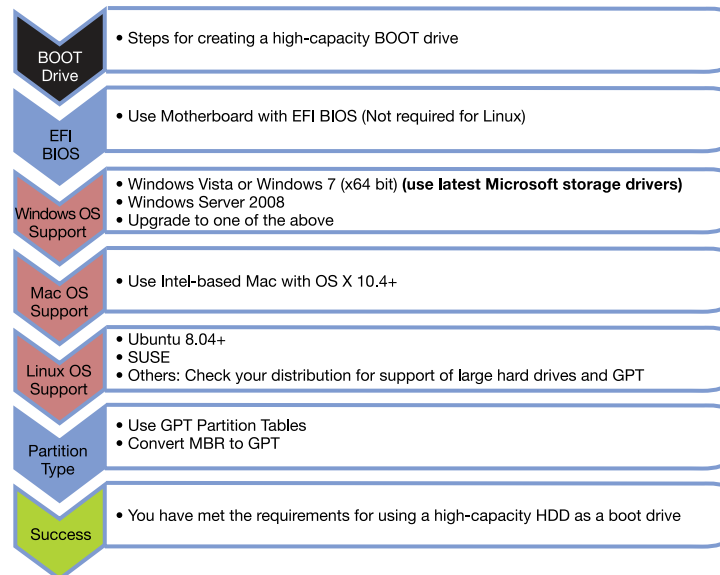
The following table is a summary of operating system support for high-capacity hard drives. All configurations assume that partitioning is done using GPT partition tables (rather than MBR).

Operating System	Boot Disk	Data Disk
Microsoft Windows XP (x32)	No	No
Microsoft Windows Vista SP1 or later x32	No	Yes
Microsoft Windows Vista SP1 or later x64	Yes, with UEFI BIOS	Yes
Microsoft Windows 7 x32	No	Yes
Microsoft Windows 7 x64	Yes, with UEFI BIOS	Yes
Intel-based Mac OS X 10.6+	Yes	Yes
Linux UBUNTU 8.04+ / SUSE x32, x64	Yes	Yes

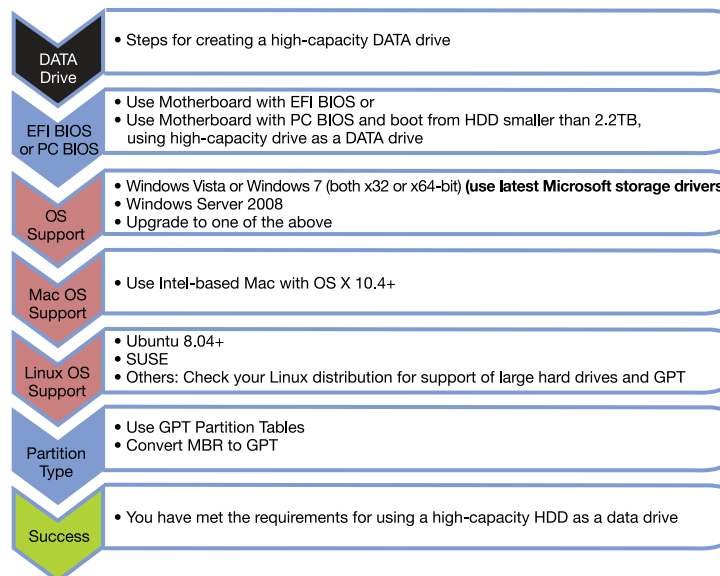
² Vista SP1 or later

³ The UEFI standard incorporates and supersedes the original Intel EFI standard. Both UEFI and EFI BIOS support ≥ 2.2TB drives.

Decision Tree: Steps for Creating a **BOOT** Drive



Decision Tree: Steps for Creating a **DATA** Drive



Server HBA/RAID configurations

Other common configurations in which $\geq 2.2\text{TB}$ drives could be used are servers, storage arrays and workstations where the use of a host bus adapter (HBA) or RAID controller will typically connect two or more hard drives. HBA and RAID controllers typically take several drives and present them to the host as a single large drive. This can be accomplished by aggregating the drives (JBOD, Just a Bunch Of Disks), or creating a complex RAID system using the capacity of more than one drive.

In both cases, one can create a larger logical drive than the individual physical drives connected to the controller. For example, if two 2TB drives are connected to the HBA, then the HBA can present both drives as a single 4TB drive to the host PC. Most HBA vendors have already created firmware and host drivers that allow this configuration to work in current operating systems, supported by the controller.

With the introduction of high-capacity drives, the HBA controller must now be able to handle individual drives larger than 2.2TB. In this case, the firmware on the controller must be able to recognize the high-capacity drive. For example, if two 3TB drives are connected to the HBA, then the HBA could present both drives as a single 6TB drive (see Figure 2).

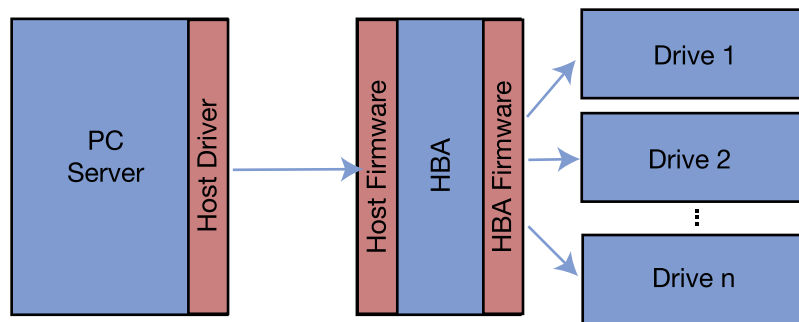


Figure 2: Server HBA/RAID Configuration

Before using a high-capacity hard drive with an HBA/RAID controller, please check with the controller vendor to determine if high-capacity hard drives are supported.

Additional sources of information:

EFI & GPT Sources

Unified Extensible Firmware Interface Forum

<http://www.uefi.org/>

GPT information

http://en.wikipedia.org/wiki/GUID_Partition_Table

Mac OS X

Technical Note TN2166 Secrets of the GPT

<http://developer.apple.com/library/mac/#technotes/tn2006/tn2166.html>

Microsoft

Windows and GPT FAQ

http://www.microsoft.com/whdc/device/storage/GPT_FAQ.msp

Hitachi, Hitachi Global Storage Technologies and Hitachi Inspire the Next logos are trademarks of Hitachi, Ltd in the U.S., Japan and/or other countries. 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.

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

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

© 2010 Hitachi Global Storage Technologies

Hitachi Global Storage Technologies
3403 Yerba Buena Road
San Jose, CA 95135 USA

Produced in the United States 11/10, revised 12/10.
All rights reserved.