

## Free Disk-Imaging Tool for Windows 8

16/08/2012

Microsoft has introduced a shiny new Resilient File System, or ReFS, in Windows 8 Server Beta. Being the successor of NTFS, ReFS has a new storage engine that promises to protect against storage failures, resist data corruptions, support ultra-large file and volume size, and grant much higher performance. It looks like Microsoft intends to make ReFS the primary file system for storing data, so it's no surprise that everybody displays a keen interest in trying this innovative file system.



Paragon Image Backup for Windows 8 is a disk-imaging tool for Windows 8 and Server 2012. It employs Microsoft VSS and Paragon's patent-pending technologies for consistent point-in-time copies of the whole disk system or separate volumes. For recovery purposes it includes a powerful Linux/DOS environment. But its power is its support of the innovative ReFS. This makes Image Backup for Windows 8 the only stand-alone freeware backup and recovery tool for Windows 8 and Server 2012 that supports Resilient File System!

### Key Benefits:

- **Support of the latest Windows 8 and Windows Server 2012**
- **Support of the brand-new ReFS (Resilient File System)** for backup, restore, and browsing under Windows and Linux
- **Live Windows backup.** Create an image-based copy of the whole disk system, or separate volumes and place it locally, to CD/DVD/BD, or on the net. The use of the disk-imaging mechanism and Microsoft VSS ([Volume Shadow Copy Service](#)) tackles the problem of backing up running applications and system locked files
- **Flexible recovery.** Restore the whole backup image, or only particular files under Windows or from the Linux/DOS recovery environment
- **Additional maintenance functions.** The Linux/DOS recovery environment includes basic functions for initializing, partitioning and formatting hard disks
- [Download Image Backup for Windows 8 for FREE](#)

### Calendar 2016

May 2019

M	T	W	T	F	S	S
		1	2	3	4	5
6	7	8	9	10	11	12
13	14	15	16	17	18	19
20	21	22	23	24	25	26
27	28	29	30	31		

« Aug

### Blog Stats

23,868 hits

### Post Index

[Free Disk-Imaging Tool for Windows 8](#)

[Partitioning GPT Disks on UEFI systems](#)

[3TB Hard Disk Utilities](#)

[Table for Operating System – File Cache Alignment](#)

[Paragon Drive Backup Professional 9.5 Free Edition](#)

[Partition Alignment for SSD's](#)

[AS SSD Benchmark](#)

[My Workstation PC \("The Force"\)](#)

### Recent Comments

梶具 on [Partition Alignment for S](#)

The Data Cloud

Open a store.  
Launch a business.  
Build any website.

WordPress.com

REPORT THIS AD

Share this:

Print & PDF Email

Like

Be the first to like this.

FILED UNDER [TECHNOLOGIC](#) TAGGED WITH [BACKUP](#), [DISKIMAGE](#), [MICROSOFT WINDOWS](#), [NTFS](#), [SOFTWARE](#)

## Partitioning GPT Disks on UEFI systems

14/03/2012 2 COMMENTS

[Tom's Hardware \(December 30, 2009 by Patrick Schmid and Achim Roos\)](#)

"You need a 64-bit version of Windows XP Professional (with the same core operating system version as Windows Server 2003 SP1), Vista, or Windows 7, to get full GPT support for drives that offer more than 2TB capacity. GPT offers the following features:

- Maximum raw partition capacity of 18 exabytes
- The Windows NTFS file system is limited to 256TB each: with up to 2<sup>64</sup> logical blocks in length and a physical cluster size of 65536-bytes (formatting with an allocation unit of 64 KiB sector groups required!)
- Implementation restriction of 128 partitions per disk (limited by the amount of space reserved for partition entries in the GPT)
- Primary and backup partition tables available for redundancy
- Well-defined and self-identifying partition format
- Each partition has a unique ID to avoid identifier collisions ("GUID" partition table)

Here is an overview of your possible partitioning options and actions on GPT and MBR:

	32-bit Windows		64-bit Windows	
	GPT	MBR	GPT	MBR
Boot	No	Yes	Yes	No

- Registry on [Partition Alignment for S...](#)
- Create a free website or blog at [WordPress.com](#).
- Lorene Cooper on [Table for Operating System...](#)
- Marisol Loter on [Table for Operating System...](#)
- Doyan on [Partitioning GPT Disks on UEFI...](#)
- Raymond Bodenhausen on [3TB Hard Disk Utilities](#)
- Sasha Christiani on [Partition Alignment for S...](#)

### ▾ Data Cloud

[Alignment Backup DiskImage](#)  
[EFI System Partition GPT Hardware Intel](#)  
[Microsoft Windows NTFS](#)  
[Partitioning RAID](#)  
[Software SSD System](#)

### ▾ Archives

[August 2012](#)

[March 2012](#)

[November 2011](#)

[January 2011](#)

[June 2010](#)

[May 2010](#)

### ▾ Me-TA MARA. (More Access Resources Available)

[Register](#)

[Log in](#)

[Entries RSS](#)

[Comments RSS](#)

[WordPress.com](#)

### ▾ Email Subscription

Enter your email address to subscribe to this blog and receive notifications of new posts by email.

Read	No	Yes	Yes	Yes
Write	No	Yes	Yes	Yes

[UEFI support](#) becomes relevant for drives >2TB. It must be present to enable booting from such large partitions provided the other prerequisites (64-bit Windows and GPT) are met.

#### GPT Details on Windows x64

GPT will automatically install the EFI System Partition (ESP), which contains the boot loader, EFI drivers, and all other necessary files for booting the system (such as boot.ini, HAL, and NT Loader). It utilizes the GUID Partition Table rather than the master boot record. The ESP is 1% of the drive capacity, or a minimum size of 100MB and a maximum size of 1,000MB.

GPT systems will also be equipped with an MSR partition, which stands for Microsoft Reserved. Since GPT partitions don't allow hidden sectors, Windows utilizes this reserved space for operating system use. If you decide to convert a basic disk into a dynamic disk, Windows will use the MSR partition, shrink its size, and create the dynamic disk database using available space. On drives smaller than 16GB, the MSR partition will be only 32MB. For larger disks, it will consume 128MB."

[Chatwin says](#): Even if your system is GPT-compliant and created using the necessary partition, you will still need [UEFI support](#) to be able to boot from a large volume (UEFI = Unified Extensible Firmware Interface).

In the second place support for 48bit [LBA](#) also requires Advanced Format CDB block re-alignment, enabled with an internal PCI-E Host Bus Adapter Card (by using the native Windows AHCI drivers only), your RAID controller and/or HDD device drivers. Here is where Long LBA addressing comes in...

Long LBA addressing increases the number of bytes used to define the LBA in the Command Descriptor Block (CDB). System must support at least 16-byte Command Descriptor Blocks (CDBs) or 32 bytes (Protection Information commands only) and Dynamic Sense Data to access the LBA counts above 2.1TB. Eight LBA address bytes are provided in 16- and 32-byte CDBs. This is a doubling of the 4 bytes in 10-byte commands ([http://www.seagate.com/docs/pdf/whit...\\_readiness.pdf](http://www.seagate.com/docs/pdf/whit..._readiness.pdf)).

#### Installation requirements


1. Windows XP Professional with SP2 for Windows Server 2003 64-bit, Vista 64-bit, Windows 7 64-bit for GPT.
2. UEFI booting support since Windows Vista with Service Pack 1 (64-bit only).
3. Motherboard with UEFI 2.x implementation for boot drives on a partition larger than 2.2 TB.
4. Windows 7 (SP1) is supposed to understand AF drives, and correctly aligns partitions. Without the updates in a hotfix, Windows can't detect the drive's physical sector size, and Windows Update (and some other features) fail. Microsoft explanation and hotfix link: [An update that improves the compatibility of Windows 7 and Windows Server 2008 R2 with Advanced Format Disks is available](#).
5. SP1 with hotfix aren't correctly updating the storage driver, so it can't report sector size to Windows. The clue here is that fsutil still reports "bytes per physical sector" as "<Not Supported>". You can fix this issue to manually install the latest Intel driver (9.6 or later): [Intel® Rapid Storage Technology](#).
6. Alignment instructions for installing XP64:
  - Use [LSoft Active Boot Disk Suite](#) (based on WinPE 3.1) to create a partition with an offset of 2048KB, format with 4K sectors and reboot.
  - Replace with XP64 SP2 installation DVD and install without altering the partition parameters.
  - After final installation reboot and use [Paragon Partition Manager Server](#) (WinPE 3.0 recovery environment) to convert cluster size from 4KB to 64KB.
  - Shutdown and start from primary HDD partition.

Advertisements

**AUTOMATTIC**

**We're hiring PHP developers anywhere in the world. Join us!**

[APPLY](#)



REPORT THIS AD

#### Paragon just introduces a new solution: [Migrate to UEFI](#).

"It's a simple four-step Copy Hard Disk Wizard that transfers a 64-bit Windows Vista/7 initially installed on an MBR disk in the BIOS mode to a GPT disk and the UEFI mode, thus opening up all benefits of the GPT+UEFI configuration, including support of a high-capacity 2.2TB+ hard drive for use as the primary system HDD.

[Become an early adopter and download this component for free>>](#)


Home and small-office users who try to move their 64-bit Windows Vista or Windows 7 software environments from the dated MBR/BIOS configuration to a modern GPT/UEFI platform find that you can't just change the partitioning scheme and enable UEFI without reinstalling the operating system from scratch. The problem is that Windows Disk Management supports conversion to GPT for empty drives only, so you've got no way of converting the system drive to GPT without losing data. Besides Windows Management tools do not provide an option to adjust boot files for UEFI, making it impossible for users transferring established data or system components within a single computer directly from a BIOS boot mode to a UEFI-based mode.

Paragon Migrate to UEFI can help you tackle both issues with minimal effort. Some important benefits:

- Automatically adjusts Windows OS to start up in the UEFI mode
- Automatically creates a GPT partitioning scheme during migration
- Supports AFD (Advanced Format Drive) disks (mixed 512B/4K sector scheme only, including virtual)
- The source system remains intact and can be loaded in the BIOS mode at any time
- Copy operation is run without the system restart by employing Microsoft VSS technology"

---

#### Share this:



Be the first to like this.

---

FILED UNDER [TECHNOLOGIC](#) TAGGED WITH [ALIGNMENT](#), [EFI SYSTEM PARTITION](#), [GPT](#), [HARDWARE](#), [PARTITIONING](#)

---

## 3TB Hard Disk Utilities

26/11/2011 21 COMMENTS

[Support for Disk Drives Beyond 2.2 TeraBytes \(TB\) and 4K Advanced Format Sectors \(Seagate Knowledge Base, ©2010 Seagate Technology LLC\)](#)

"Most legacy systems built before 2011 have a traditional PC BIOS. This type of BIOS uses a Master Boot Record (MBR). The MBR Partitions can define a disk drive capacity up to 2.2TB. Windows operating systems that boot from an MBR are therefore limited to 2.2TB per MBR. A 3TB disk drive in a legacy BIOS and Window system will need a DiscWizard device driver to access the full capacity of a 3TB disk drive. Two partitions will be necessary because of the MBR limitation. The device driver

mounts the capacity above 2.2TB with another MBR which looks to the system as a second virtual “physical” device.

**PCWorld:** The problem with deploying 3TB drives relates to older PCs (those more than a few months old, in most cases), and stems from the formula  $2^{32} \times 512 = 2,199,023,255,552$ , or 2.2TB – a hard-drive addressing scheme found in legacy BIOSs and operating systems. In that formula, 2 indicates binary, 32 is the number of bits allowed in a legacy disk address, and 512 is the number of bytes in a legacy hard-drive data block. If the BIOS, drivers, I/O card, or operating system in your PC still plays by rules that involve this formula, you’ll have issues installing and using a 3TB drive.

This situation could have been avoided if the entire computer industry had future-proofed after enduring the [137GB \(28-bit\) limit problems](#) that cropped up around the turn of the millennium. In truth, most vendors did, with the notable exception of Microsoft. The company chose not to implement support for anything larger than 2.2TB drives in any of its 32-bit consumer operating systems – including Windows 7.

**Installing a 3TB Drive:** Note that NTFS is limited to  $2^{32}$  clusters, also known as groups of sectors. That means you must format the drive with at least 1024-byte clusters, or you’ll fall prey to the formula I talked about earlier. The default is 4096, which allows up to 16TB; however, converting FAT partitions to NTFS often results in 512-byte clusters. If you’re trying to move a FAT partition to your new drive—don’t. You’re better off creating a new partition, reinstalling the OS, and restoring the files you backed up (it’s always wise to back up before doing any partitioning operations). This is because the FAT partition will in all likelihood be misaligned with the Advanced Format of your new drive.

**Tip:** With drivers or programs that don’t allow a full 3TB partition, you can use Windows dynamic volumes to combine two partitions into a single drive letter.

Windows Vista and Windows 7 introduced a new partition-table scheme, dubbed GPT (for GUID Partition Table). The GPT blasted past the previous limitations by supporting up to 8 zettabytes. ( $2^{64}$  sectors, which, when using 512 bytes per sector, equates to 8 zettabytes). For perspective, consider that 1024 terabytes equal 1 petabyte, 1024 petabytes equal 1 exabyte, and 1024 exabytes equal 1 zettabyte.

Though not strictly a 3TB issue, all of the shipping [3TB hard drives](#) also use Advanced Format, a low-level storage scheme employed on newer drives. AF uses larger 4KB data sectors, an approach that improves performance and diminishes the number of addresses required for any given amount of data. If you transfer legacy partitions via image backups to an AF drive, or if you format an AF drive with XP, the older 512-byte sectors might not align correctly with the new scheme.

UEFI BIOS desktop systems are new since 2011. Windows 7 64-bit and Vista 64-bit operating systems support booting from UEFI and GPT without the need of a non-Microsoft device driver. This is the Windows native solution for booting a 3TB drive to a single partition.

#### **Quick facts about Windows and 3TB drives:**

- Windows 7 and Vista support GPT 3TB single partitions.
- Windows 7 and Vista can only boot GPT on systems with UEFI BIOS.
- Windows 7 and Vista can mount a GPT non-booting data drive.
- Intel RST device drivers before v10.1 do not support 3TB disk drives.
- Windows systems with Legacy BIOS and MBR boot drives are limited to 2.2TB partitions.
- Windows XP x32 does not support GPT. Windows XP x64 edition can use GPT disks for data only. Only Windows XP for Itanium-based systems can boot from GPT partitions.
- While 32-bit Windows XP systems can’t handle GPT natively, there is a solution: Paragon Software Group’s [Paragon GPT Loader](#), that allows you to utilize all 3TB in a single data

partition.

- DiscWizard software (Acronis®) can install a device driver which opens the full capacity of a 3TB. You can use it to create a second partition for the capacity above 2.2TB.

**DiscWizard v13** with support for 3TB drives is [now available](#).


- For a Flash tutorial: Cloning your old drive to a 3TB using DiscWizard. [Part 1](#), [Part 2](#).
- For a Flash tutorial: [Using DiscWizard to access the full capacity of a 3TB after you built a Windows 7 or Vista system](#).
- For a Flash tutorial: [Using DiscWizard to access the full capacity of a 3TB after you built a Windows XP system](#).

**DiscWizard v11** for Windows is available today for all lower capacities. [Click here to access DiscWizard v11.](#)"

Overview for the best free disk imaging programs ([Gizmo](#)):

1. Acronis for Seagate/Maxtor HDDs <http://www.seagate.com/www/en-us/support/downloads/discwizard>
2. Acronis for Western Digital HDDs <http://support.wdc.com/product/downloaddetail.asp?swid=119>
3. 1-Click Restore Free [http://www.keriver.com/oneclick\\_free.html](http://www.keriver.com/oneclick_free.html)
4. Comodo Backup <http://backup.comodo.com/>
5. Drive Image XML <http://www.runtime.org/driveimage-xml.htm>
6. PING <http://ping.windowsdream.com/>
7. Clonezilla <http://clonezilla.org/>
8. Paragon Backup and Recovery 2010 <http://www.paragon-software.com/home/br-free/>
9. O&O DiskImage 4 Express Edition <http://www.oo-software.com/home/en/products/oodiskimageexpress/>
10. Macrium <http://www.macrium.com/reflectfree.asp>
11. EASEUS <http://www.todo-backup.com/business/free-backup.htm>

Share this:



Be the first to like this.

FILED UNDER [TECHNOLOGIC](#) TAGGED WITH [BACKUP](#), [GPT](#), [MICROSOFT WINDOWS](#), [NTFS](#), [PARTITIONING](#), [SOFTWARE](#)

## Table for Operating System – File Cache Alignment

05/01/2011 38 COMMENTS

[Geeking Out on SSD Hardware Developments \(Linux Magazine, February 9th, 2010\)](#)

### Intel and Smaller Cells

"According to this [link](#) the theoretical lower limit on NAND cells appears to be around 20nm. One of the big reasons for this limitation is that as cells shrink in size they are, naturally, closer together. However, the voltage required to program a cell remains about the same (typically 12V). This means that the power density increases (amount of power in a given area) increasing the probability that

the voltage will disturb the neighboring cells causing data corruption. Consequently, increasing the density of cells can be a dicey proposition, hence the lower limit.

Recently, IMFT (Intel-Micron Flash Technologies) LLC announced that it had begun sampling 2 bits per cell (MLC) NAND flash chips that have been manufactured using a 25nm process. This announcement is significant because of the increased density and how close the density is getting to the theoretical limit. Plus the fact that one of the best performing SSD drives is from Intel, one of the participating companies in the LLC, and we can see a fairly significant shift in technology.

IMFT is a joint project by Intel and Micro to develop new technologies around Flash storage (primarily NAND). The project started with production several years ago with a 72nm process. They then moved to a 50nm process in 2008 followed by a further reduction to a 34nm process in 2009. It is the 34nm process that current Intel SSDs utilize (Intel X25-M G2). The 34nm process produces a 4GB MLC NAND chip with a die size of 172 mm<sup>2</sup>. The new 25nm process is targeted for a first product which is an 8GB MLC NAND flash chip with a die size of 167 mm<sup>2</sup>. So going from 34nm to 25nm doubles the die density.

In addition to the doubling of die density the new chips will have some other changes. The current 34nm generation of chips have a page size of 4KB and 128 pages per block resulting in a block size of 512KB. The new chip will have a page size of **8KB** and **256** pages per block. This means that the new block size is 8KB \* 256 = 2,048KB (2MB). This change in block size can have a significant impact on performance.

Recall that a block is the smallest amount of storage that goes through an erase/write cycle when any single bit of the block is changed. For example, if any bit within the block is changed then the entire block has to first have the unchanged data copied from the block to cache and then the block is erased. Finally the updated information is merged with the cache data (unchanged data) and the entire block is written to the erased block. This process can take a great deal of time to complete and also uses a rewrite cycle for all of the cells in the block (remember that NAND cells have a limited number of rewrite cycles before they can no longer hold data).

The new 25nm chip will switch from **512KB** blocks to **2MB** blocks (2,048KB) increasing the amount of data that has to go through the read/erase/write cycle compared to the 34nm chips. To adjust to this change, Intel will have to make adjustments to the firmware to better handle the larger blocks. In addition, it is likely that Intel will have to at least double the cache size to accommodate the larger block sizes. It may have to increase the number of spare pages as well since a single bit change could cause a greater number of blocks to be tagged for updating. However, on the plus side, with larger pages, the controller can do much more optimization for writes including a much greater amount of write coalescing. But this too could increase the amount of cache needed."

Chatwin's Table

	<u>Conventional Allocation</u>	<u>Improved Allocation</u>
	4 KiB <u>Alignment</u> 1 MB	64 KiB <u>Alignment</u> 2 MB
OS – 32bit File Level Cache	Page – 32 K Size – 128 MB	Page – 64 K Size – 256 MB

OS – 64bit File Level Cache	Page – 64 K Size – 256 MB	Page – 128 K Size – 512 MB
OS – 32bit Storage Level Cache	Page – 64 K Size – 256 MB	Page – 128 K Size – 512 MB
OS – 64bit Storage Level Cache	Page – 128 K Size – 512 MB	Page – <b>256 K stripe unit</b> Size – 1 GB
	<u>Cluster Unit Aligned</u> 256 Kibit ( <b>32 KB</b> )  Windows XP x32 (8*4) System RAM: 2 GB  Physical Page Size: 32 Kibit (4 KB) Logical Block Size: 128 KBytes  Sector Size – 512 bytes	<u>Cluster Unit Aligned</u> 512 Kibit ( <b>64 KB</b> )  Windows XP x64 (8*8) System RAM: 4 GB  Physical Page Size: 64 Kibit (8 KB) Logical Block Size: 256 KBytes  Sector Size – 1024 bytes

Share this:

 Print & PDF  Email

 Like

Be the first to like this.

 FILED UNDER [TECHNOLOGIC](#)  TAGGED WITH [ALIGNMENT](#), [HARDWARE](#), [INTEL](#), [RAID](#), [SSD](#), [SYSTEM](#)

## Paragon Drive Backup Professional 9.5 Free Edition

 01/06/2010  7 COMMENTS

[Paragon giving away Drive Backup 9.5 until Tuesday – TechSpot News](#)

“[Paragon Software Group](#) is offering up a free copy of its Drive Backup Professional tool. The software usually retails for \$50 and includes several easy to use features to backup an image of either your entire hard drive or a partition. Backups can be performed in multiple ways, each with a different purpose, from simple backup of files to more advanced scheduled backup operations. You’ll get a very reliable and powerful backup solution with a lot of tools and features for getting your stuff secured against a possible drive failure in the future. If anything goes wrong and you need to restore your data, you can do it from a single backup image at any time without requiring you to install further software. A bootable recovery CD or USB stick can also be created.



Paragon Drive Backup 9.5 Professional is [available for download here](#). You must register to get the free serial."



From the website of Paragon: Take complete control of your PC's safety. Based on solid commercial backup and recovery software from Paragon, the new Backup & Recovery 10.2 Free Edition will give you a rich set of features that you can trust. [Free Download](#).

### Share this:



Be the first to like this.

FILED UNDER [TECHNOLOGIC](#) TAGGED WITH [BACKUP](#), [DISKIMAGE](#), [PARTITIONING](#), [SOFTWARE](#)

## Partition Alignment for SSD's

30/05/2010 25 COMMENTS

[Microsoft Support: Article ID: 929491 – Last Review: June 8, 2009 – Revision: 4.0 \(Summary\)](#)

"Disk performance may be slower than expected when you use multiple disks in Microsoft Windows Server 2003, in Microsoft Windows XP, and in Microsoft Windows 2000. For example, performance may slow when you use a hardware-based redundant array of independent disks (RAID) or a software-based RAID.

This issue may occur if the starting location of the partition is not aligned with a stripe unit boundary in the disk partition that is created on the RAID. A volume cluster may be created over a stripe unit boundary instead of next to the stripe unit boundary. This is because Windows uses a factor of 512 bytes to create volume clusters. This behavior causes a misaligned partition. Two disk groups are accessed when a single volume cluster is updated on a misaligned partition.

To verify that an existing partition is aligned, divide the size of the stripe unit by the starting offset of the RAID disk group. Use the following syntax:

**((Partition offset) \* (Disk sector size)) / (Stripe unit size)**

Example of alignment calculations in bytes for a 256-KB stripe unit size:

$$(63 * 512) / 262144 = 0.123046875$$

$$(64 * 512) / 262144 = 0.125$$

$$(128 * 512) / 262144 = 0.25$$

$$(256 * 512) / 262144 = 0.5$$

$$(512 * 512) / 262144 = 1$$

Example of alignment calculations in kilobytes for a 256-KB stripe unit size:

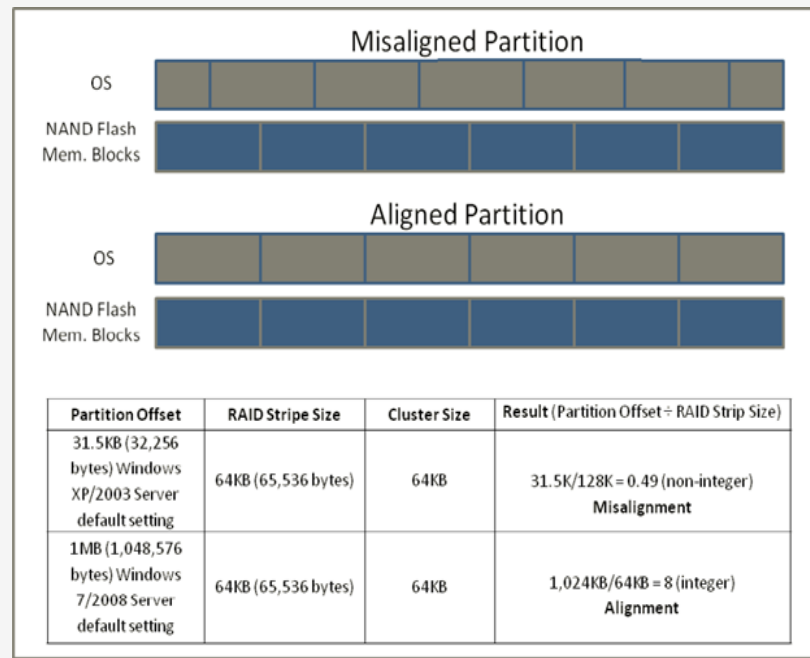
$(63 * .5) / 256 = 0.123046875$   
 $(64 * .5) / 256 = 0.125$   
 $(128 * .5) / 256 = 0.25$   
 $(256 * .5) / 256 = 0.5$   
 $(512 * .5) / 256 = 1$

These examples shows that the partition is not aligned correctly for a 256-KB stripe unit size until the partition is created by using an offset of 512 sectors (512 bytes per sector)."

SMART Modular Technologies (author: Esther Spanjer), April 2010

"NAND flash devices are divided into erasable blocks composed of multiple pages (up to **256** pages per block, and up to **8KB** per page). A flash block must be fully erased prior to re-writing, and a single-block erase process can take up to several milliseconds. The write speed may suffer a great deal if the SSD controller has to perform unnecessary block erase operations due to partition misalignment. Proper partition alignment is one of the most critical attributes that can greatly boost the I/O performance of an SSD due to reduced read modify-write operations.

Windows XP or Windows Server 2000/2003 start partition offset at 31.5KB (32,256 bytes). Due to this misalignment, clusters of data are spread across physical memory block boundaries, incurring read-modify-write penalty. As a result, the host ends up writing at least 2X more I/O for every write as illustrated in this figure:



When choosing a partition starting offset, SMART Modular recommends system integrators to correlate the partition offset with the RAID stripe size and cluster size to achieve the optimal SSD I/O performance. The figure shows an example of a misaligned partition offset and an example of an aligned partition offset for Windows Server."

Chatwin says: Based on the conclusions as mentioned above, I would recommend the following alignment and cluster settings for NAND flash devices (unit = 512 bytes):

**((Partition offset / Flash unit) \* (Disk sector size)) / (Erase block)**

$((2048 / 2) * .5) / 512K = 1$

$((4096 / 4) * .5) / 512K = 1$

With a flash unit of 2 for SLC flash (aka. 1K, 1 byte/cell group [32×32]) and 4 for MLC flash (aka. 2K, 2 bytes/cell group [32×32]), this results in an offset of 1 MB (2048) for SSD's with Single Layer Cell memory, and 2 MB (4096) for Multi Layer Cell memory, when you want to align your partitions on erase block boundaries. The best cluster size (allocation unit) can be calculated with this formula:

**((Single erase block / Flash unit) \* (Disk sector size)) = Cluster Size**


$((128 / 2) * .5) = 32K$  (SLC),

$((512 / 4) * .5) = 64K$  (MLC, "The Force")

With 8K as the ideal sector size for flash memory (4K SLC), this is in my opinion the best choice when you want the fastest performance for your SSD (Indilinx, Samsung, JMicron controllers), unless you're more concerned about maximizing your storage capacity at any price. And not to forget: a longer lifespan, for reducing a lot of overhead and fragmentation compared to the standard NTFS formatting of 4096 bytes. Keep in mind that this setting was default in a time that hard disks were limited to 2-4 GB, instead of 2 TB in this time and age.

---

### Share this:



Be the first to like this.

FILED UNDER [TECHNOLOGIC](#) TAGGED WITH [ALIGNMENT](#), [PARTITIONING](#), [RAID](#), [SSD](#)

---

## AS SSD Benchmark

26/05/2010 7 COMMENTS

[Benchmark Reviews](#) (© 2010):


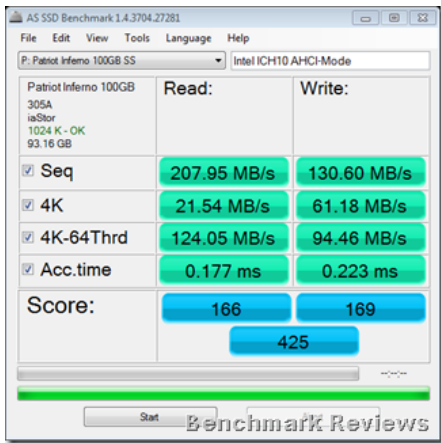
"Solid State Drives have traveled a long winding course to finally get where they are today. Up to this point in technology, there have been several key differences separating Solid State Drives from magnetic rotational Hard Disk Drives. While the DRAM-based buffer size on desktop HDDs has recently reached 32 MB and is ever-increasing, there is still a hefty delay in the initial response time. This is one key area in which flash-based Solid State Drives continually dominates because they lack moving parts to "get up to speed". However the benefits inherent to SSDs have traditionally fallen off once the throughput begins, even though data reads or writes are executed at a high constant rate whereas the HDD tapers off in performance. This makes the average transaction speed of a SSD comparable to the data burst rate mentioned in HDD tests, albeit usually lower than the HDD's speed. Comparing a Solid State Disk to a standard Hard Disk Drives is always relative; even if you're comparing the fastest rotational spindle speeds. One is going to be many times faster in **response** (SSDs), while the other is usually going to have higher throughput **bandwidth** (HDDs)."

"The biggest mistake PC hardware enthusiasts make with SSDs is grading them by their speed. File transfer speed is important, but only so long as the operational IOPS performance can sustain that bandwidth under load. Benchmark Reviews tests the 100GB Patriot Inferno SSD, model PI100GS25SSDR, against some of the most popular storage devices available and demonstrates that 4K IOPS performance is more important than speed. For decades, the slowest component in any computer system was the hard drive. Most modern processors operate within approximately **1 ns** (nanosecond = one billionth of one second) response time, while **system memory** responds between **30-90 ns**. Traditional Hard Disk Drive (HDD) technology utilizes magnetic spinning media, and even the fastest spinning desktop storage products exhibit a 9,000,000 ns – or **9 ms** (millisecond = one thousandth of one second) initial response time. In more relevant terms, The processor receives the command and waits for system memory to fetch related data from the storage drive. The difference a SSD makes to operational reaction times and program speeds is dramatic, and takes the storage drive from a slow 'walking' speed to a much faster 'driving' speed. Solid State Drive technology improves initial response times by more than 450x (45,000%) for applications and Operating System software, when compared to their HDD counterparts. Alex Schepelijanski of Alex Intelligent Software develops the free AS SSD Benchmark utility for testing storage devices. The AS SSD Benchmark tests sequential read and write speeds, input/output operational performance, and response times. Because this software receives frequent updates, Benchmark Reviews recommends that you compare results only within the same version family. Beginning with sequential read and write performance, the Patriot Inferno Solid State Drive produced 207.95 MB/s read speed, and 130.60 MB/s write performance. The sequential file transfer speeds have traditionally been low with this benchmark tool, especially for SandForce controllers, which is why we will concentrate on the operational IOPS performance for this section. Single-threaded 4K IOPS performance delivers 21.54 MB/s read and 61.18 MB/s write, which is among the highest results we've recorded. Similarly, the 64-thread 4K reads recorded 124.05 MB/s while write performance was 94.46... both earning the Patriot Inferno SSD a spot at the very top of our charts."

Chatwin's opinion: The Patriot Inferno (on the left) has excellent random R/W values, thanks to the SandForce controller, which uses complex algorithms to compress small user-data into 4K flash pages. The manufacturer also implements some kind of data redundancy, to ensure data integrity. All of this increases the IOPS and efficiency of the SSD and reduces the need for high quality (SLC) NAND flash or large DRAM buffers. The Indilinx controller of the Vertex Turbo (on the right) has a straight forward page level mapping (with 64 MB memory cache to combine write requests). In combination with a logical/physical 1:1 mapping (LBA/SSD), the Vertex (RAID 0, specifications: "[The Force](#)") outperforms every other MLC SSD I've seen so far in access time (except the Intel X25-M G2 Postville, with a narrowed write bandwidth of 95 MB/s and 32K clustering on-the-fly, these SSD's are simply unbeatable...). The sequential read/write speed is quite impressive. The two OCZ Vertex Turbo's score lower IOPS values than the Patriot, but the bandwidth is much higher, even for RAID configurations.

Besides an exceptional response time, the result is better write coalescing when using my workstation for day-to-day stuff: email checking, editing office documents, internet browsing (caching redirected to hard disk of course). An average of 30 MB/hour at the most (zero to 3 when left idle). I use Diskeeper HyperFast for maintenance, but so far it never initiated an automatic (free space) defragmentation. Even after several months of installing software and restoring system backups.

*Note: AS SSD Benchmark was run in 2 sessions (first: Seq./Acc. Time, then 4K (-64) Thrd). No tricks or manipulations involved...*



Share this:

Be the first to like this.