

## Comparing 1 HyperDrive4 to 10 Western Digital Raptor 740ADFDs in RAID6

The latest WD Raptor, the WD740ADFD, is the fastest commercially available SATA Hard Disk at present (December 2006). And the Areca 12xx series RAID cards, being full Hardware RAID, and using the Intel IO33x series chip are generally acknowledged to be the fastest or thereabouts of the server RAID cards. So the best the mankind can achieve in performance (given an acceptable level of data security) with hard disks is represented by around 10 WD Raptors in RAID6.

The Areca 1280 card uses the Intel IO341 chip and can carry up to 24 SATA Disks. But it can only manage a maximum Sustained Read/Write rate of 771 MB per second - see [http://www.areca.com.tw/news/html/release\\_9\\_27\\_06.htm](http://www.areca.com.tw/news/html/release_9_27_06.htm). Now since one WD Raptor can produce a Sustained Transfer Rate of 90 MB per second, any more than 8 of them in RAID0 would push the Areca RAID card over its STR limit. So there will be no real performance advantage in striping more than 8 raptors at present. Then we add the two parity disks for RAID6 and get 10 Raptors in RAID6 as an optimum performance real world RAID solution.

It is fairly easy to deduce what performance advantages true hardware RAID with 8 and 16 and 32 and 64 disks could theoretically produce from the performance advantages that 2 and 4 disks in RAID0 have over 1 disk in RAID0. Generally, in the absence of read ahead caches affecting performance, if 2 disks are 50% faster than one disk at a task, then 4 disks will be 25% faster than 2 disks, and 8 disks will be 12.5% faster than 4 disks, and 16 disks will be 6.25% faster than 8 disks, etc etc. So the vast majority of RAID performance advantage comes from the first 4 or 8 Striped disks. Striping 16 or more disks puts you well into the law of exponentially diminishing returns.

Furthermore, with Hard Disks, the more you use, the more likely you are to have a mechanical failure. So the maintenance cost of very large arrays of Hard Disks means that striping more than 8 is really not worth it for the small percentage performance gain.

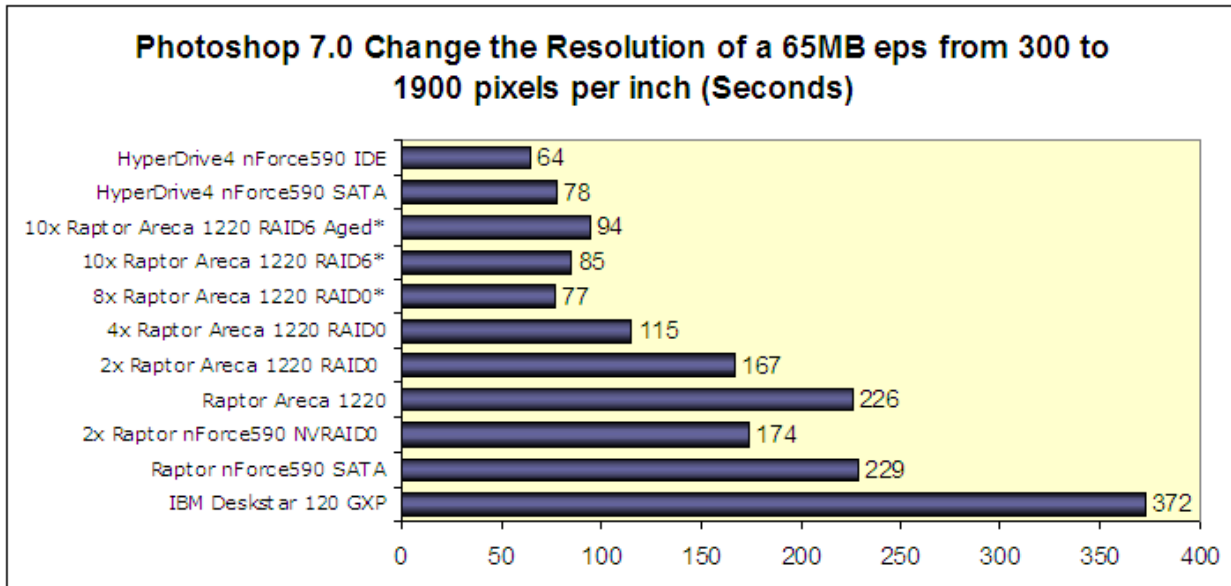
At this time we have only tested 1, 2 and 4 Raptors in RAID0 on an Areca 1220 card. But from these figures we can fairly accurately deduce the theoretical RAID0 performance of 8 Raptors, and the Areca card is good enough to achieve that performance. RAID6 is recognised by Areca as producing a 9% performance hit over a RAID0 system with two less disks (Sustained read rates are 885 MB/s vs 811, sustained write rates are 847 MB/s vs 776). The two extra disks required for the 2 disk failure safe rotating parity overhead of RAID6 cause this 9% performance loss. So we can deduce RAID6 performance from RAID0 performance. The one being 91% of the other (less 2 disks).

In all our charts below we have marked theoretical results with an asterisk.

So now we can ask the 64 million dollar question...

**Which would you say is faster?  
1 HD4 all by itself or 10 WD740ADFDs in RAID6?**

Well, here is the result of the Photoshop contest...



The test system was an ASUS CrossHair Mobo with a dual core 3800 Athlon 64 processor and 1GB of 667 MHz DDR2 and a crummy Geforce6 Graphics card. We ran Windows XP and Photoshop and the paging file entirely on the HD4 and then entirely on the Raptors. We used the IBM Deskstar as the boot disk, with HyperOs 2006SE SuperGeek to swap to the other copies of Windows on the HD4 and the RAID Raptors. The boot disk only uses the MBR, BOOT.INI, NTLDR and NTDETECT.COM, when booting XP resident on another disk, so it did not cause a measurable performance hit.

The SATA results are significantly worse than the IDE results for the HD4 because it is a native UDMA6 ATA133 device which employs a SATA conversion chip. This chip is fine for hard disk performance levels, but slows down the HD4 by around 10-20%!!


But even with the SATA connection, the HD4 is still 20% faster than 10 WD Raptors in RAID6 aged for normal disk fragmentation (a 10% performance hit). The HD4 is a random access device and so suffers no performance degradation due to fragmentation. It does not care where in its DDR the file is that you are looking for, or how fragmented that file is. It finds every bit of data at the same speed from all over the RAMdisk. So it does not age with time as all Hard Disk based PCs do.

So for Photoshop scratch disk operations the HD4 is faster than 10 RAID6 Raptors and in fact would be faster than 18 or 34 or 66 RAID0 Raptors. And if one could build a RAID card that would take 1024+2 RAID6 Raptors, it would be faster than them as well, due to the exponentially diminishing performance return one gets from doubling the number of hard disks in ones RAID array. True RAID doubles the STR, the sustained transfer rate of a disk array, it does not affect the seek time, that is the time the RAID array takes to find a file. Being able to borrow loads of books from a library is not much use if the librarian takes a long time to find each book. This mechanical latency means that 1026 Hard Disks in RAID6 would only be perform around 300%-400% faster than one Hard Disk in most IO intensive situations.

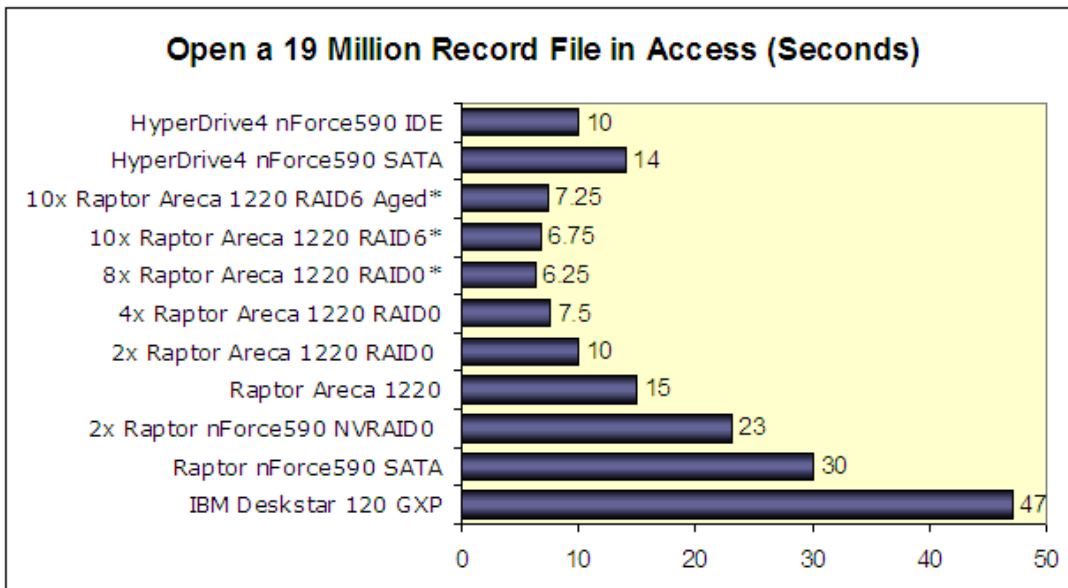
This is not true with HyperDrive4s however, because they have a seek time of 40 microseconds rather than 8 milliseconds (The WD740ADFD has 8.3 ms seek time as measured by HDTach3.0 or H2BenchW). So 8 RAID HD4s would actually be 800% faster than one HD4. The seek time fall off would not happen until we have around 200 HD4s in RAID presumably since the seek time is 200x faster than a Hard Disk.

So a machine with 2 RAID HD4s would be twice as fast as a machine with any number of RAID6 Raptors for IO intensive applications.

## Large File Copying

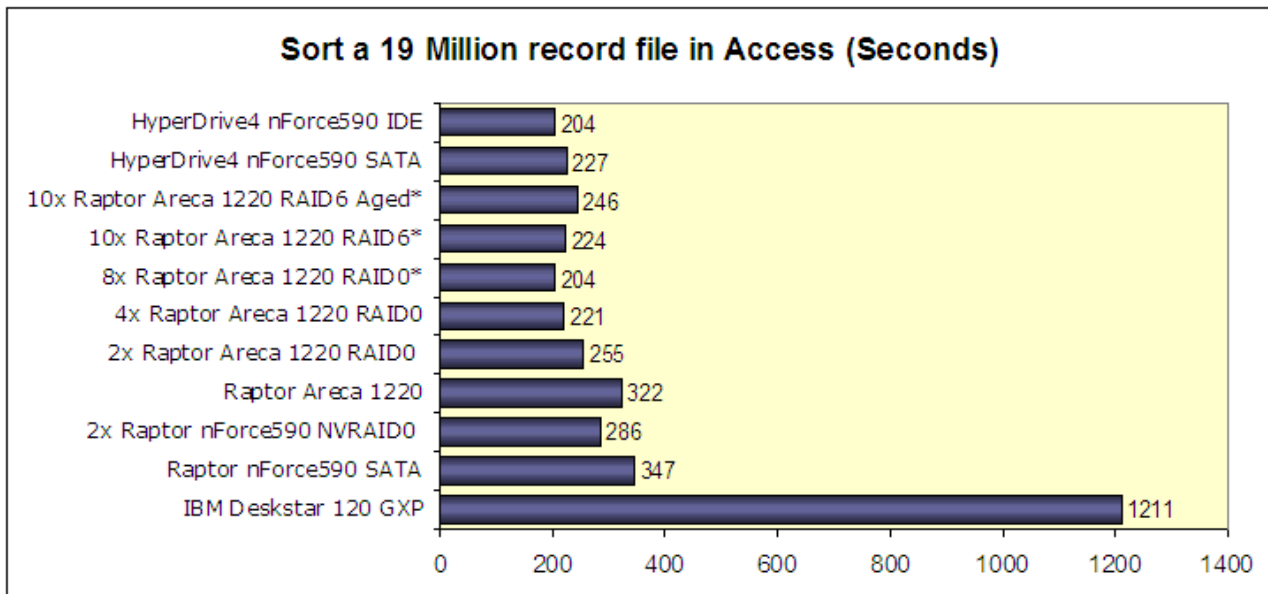


The Areca 1220 has a good look ahead read cache so it performs better than normal when reading one large file which is not fragmented on the disk. But in the real world that rarely happens. Here are the results for example when opening a large 19 million record table in Access....

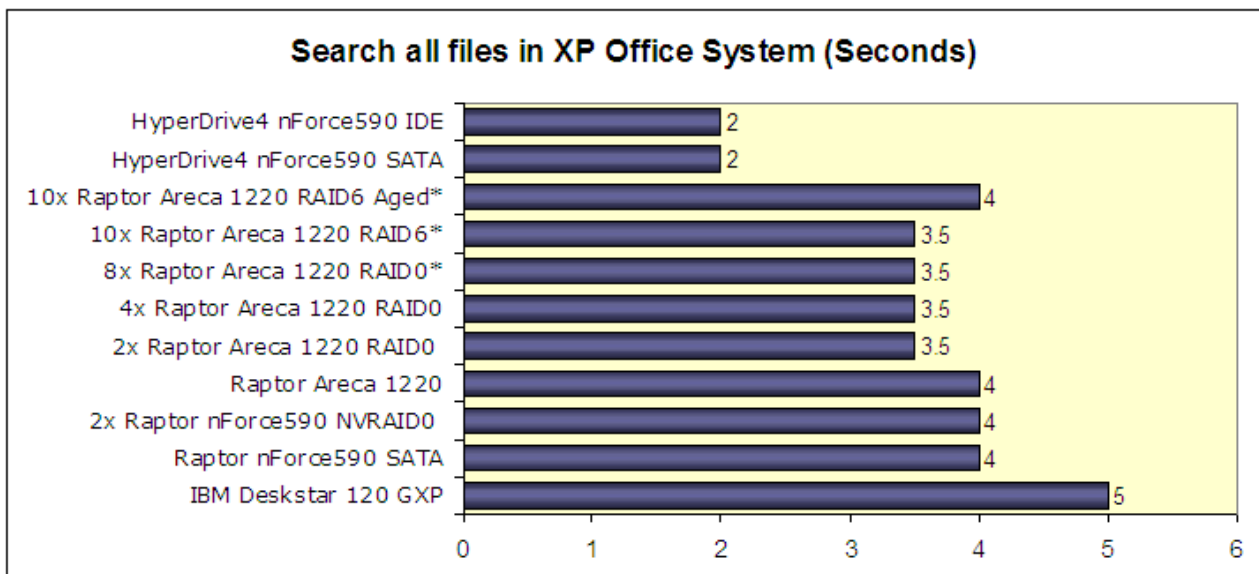


The performance differential between the HD4 and the 10 RAID aged raptors is reduced. But then here are the results of sorting that 19 million record file...

## Database operations




Since the sort operation requires a lot more head movement than reading a large table into memory, the HD4 overtakes the Raptors again. Both SATA and IDE connected HD4s are faster than 10 RAID6 Raptors. The performance differential is even larger when searching through loads of small files...




## Instant Desktop

The real performance advantage to the user however, is the incredibly fast feel of his desktop. You have to experience this to appreciate it, or to believe it! The whole desktop comes alive. It resonates before your eyes. It jumps to the task to which you point it. You really do feel like it can think faster than you can click or type - which is a new experience with a Windows machine. It might sound like a silly thing to say for a 4GHz processing PC, but that is the reality of gramophone technology computing. The famously slow Photoshop opens a large file in 4.5 seconds on an HD4. That is saving 3 to 8 seconds of your real time. This is not processing time, it is your time.



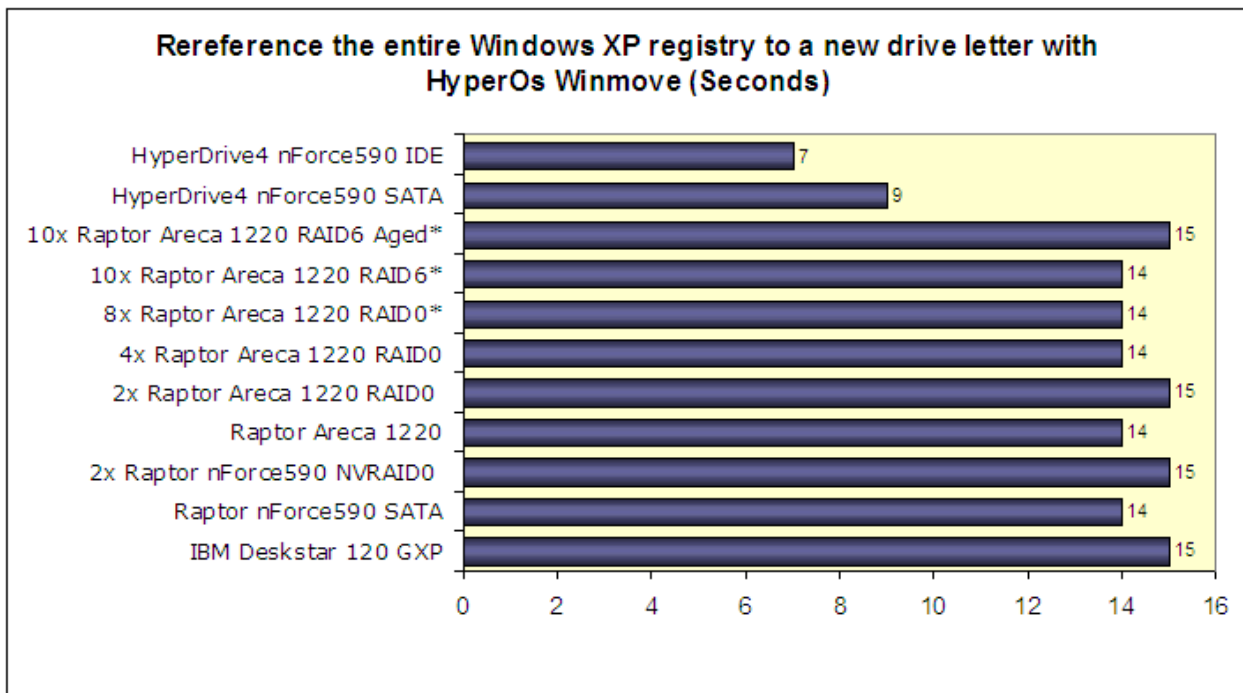
When surfing the net on an HD4, your PC feels like you are on a much faster broadband. It feels like your telecoms company has only one customer. And any nasties you may catch will disappear after the session if you are using HyperOs and a cloned system. Alternatively you can just unplug the HD4 and everything disappears forever.

## Booting up XP



## Complex Registry Operations

You cannot cache registry operations. So reading from and writing to the Windows XP registry is an easy win for the HyperDrive4.



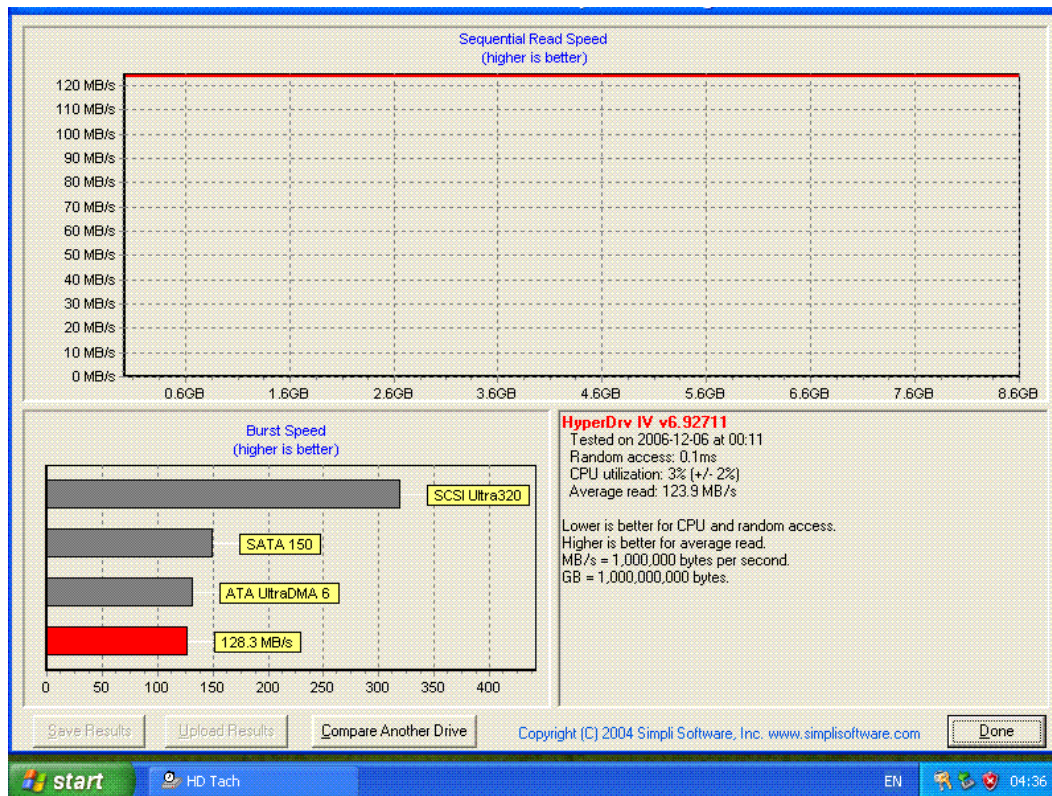
## Multiple Users

All of the above is for one user, computing on one HD4. If the single user wants really unattainable performance he must go to 2 or more RAID HD4s. But if a multiple user group wants unattainable performance then one HD4 is enough. If 4 people use an HD4 at the same time, then it will be 4 times as fast as 10 RAID6 Raptors in an IO intensive environment. And it will be much faster than 10 RAID6 Raptors in any environment at all. Because RAID does not help 2 users to find two files any faster. It only helps to deliver those files once found. So really any network with multiple concurrent users must use HD4 for performance. Then the well known phrase: *I am waiting for my screen to come up*, will no longer be heard in forward thinking organisations!

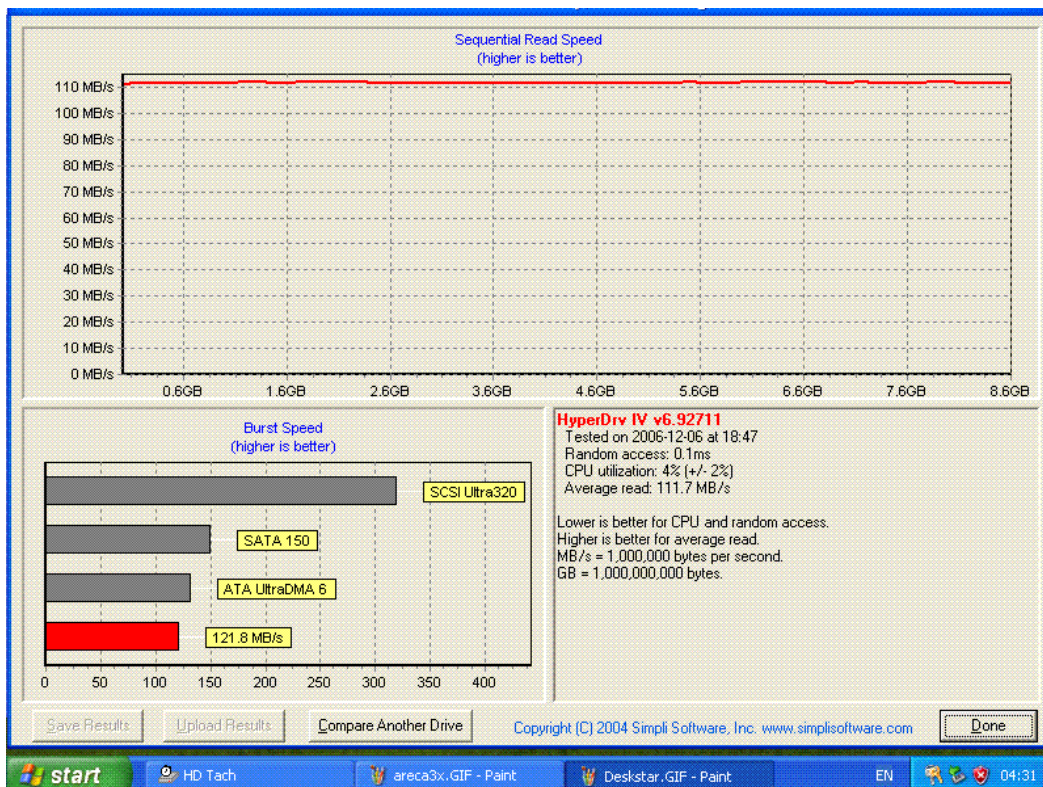
## HDTach Results for all the Storage Systems

**HD4 IDE**  
**HD4 SATA**  
**Areca 4x**  
**Areca 3x**  
**Areca 2x**  
**Areca 1x**  
**nVidia 2x**  
**nVidia 1x**  
**Deskstar**

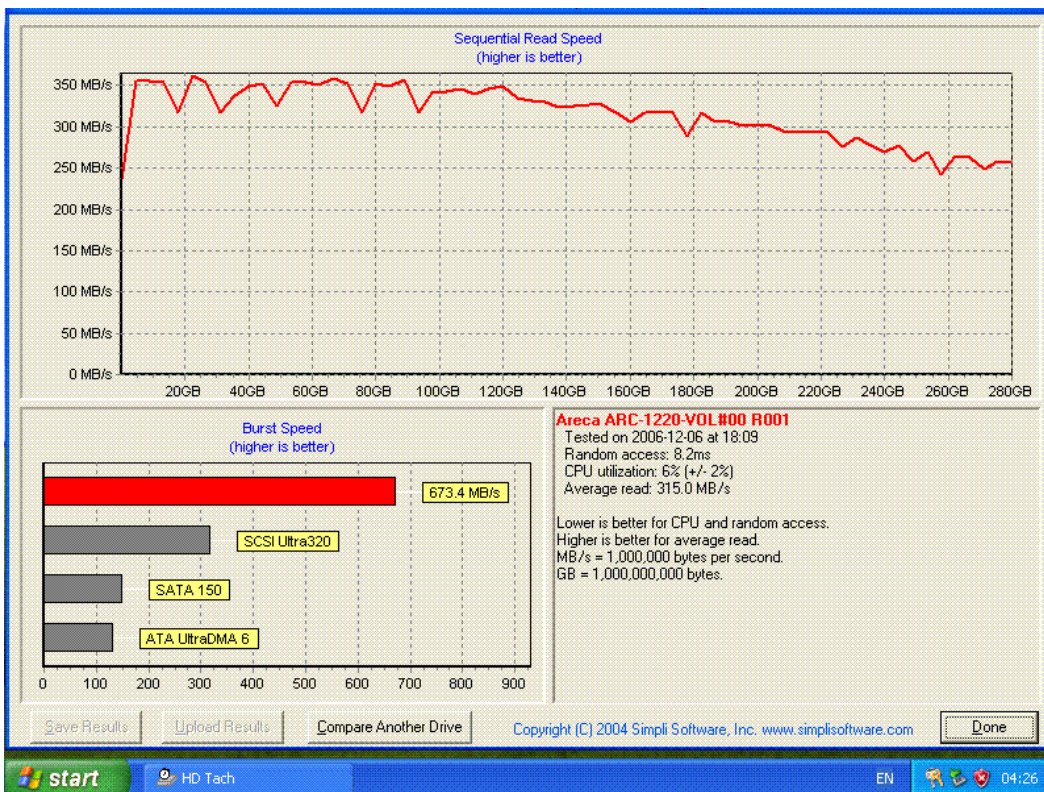
## HD4 IDE




## HD4 SATA




## Areca RAID 4x




## Areca RAID 3x




## Areca RAID 2x




## Areca 1x



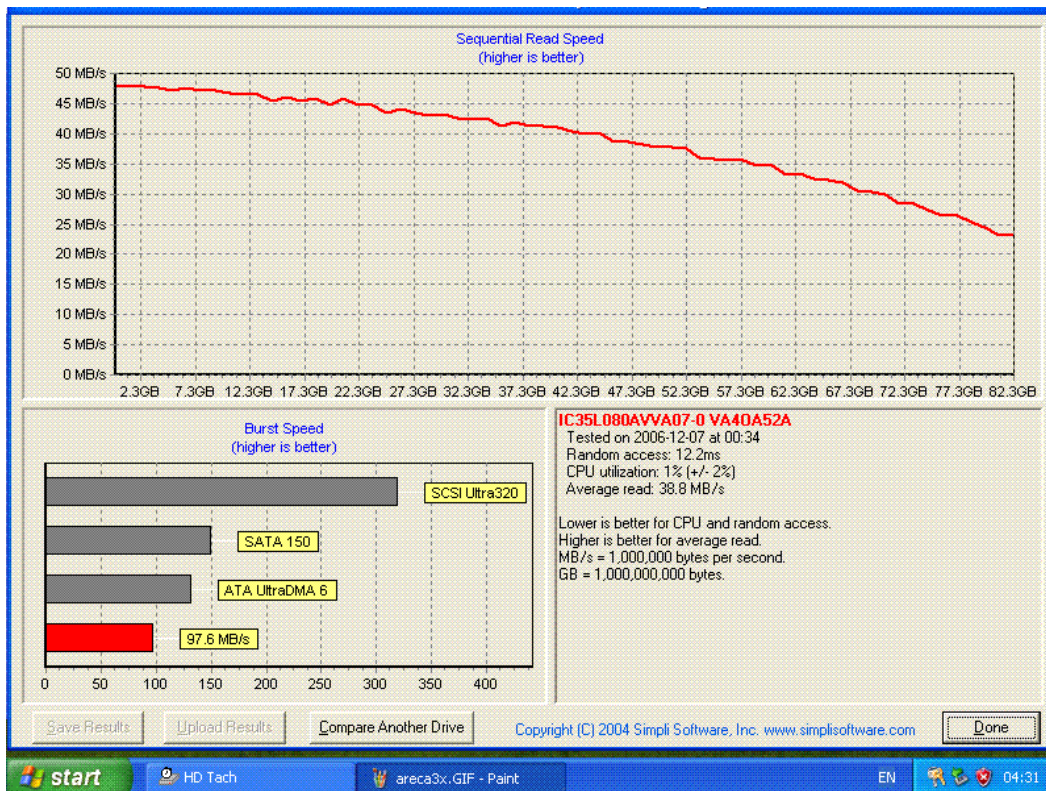
## nVidia SATA 2x



## nVidia SATA 1x



## Deskstar 120 GXP



Look out for the next two performance reports on RAID HD4s and on multiple users

Testing with Intel IPEAK!

Copyright HyperOs Systems London December 2006