



## Understanding Solid State Drives (part two - performance)

**One of the big questions around SSD concerns performance. Are they faster than current drive technology? Can they deliver data more efficiently? Finding these answers presents a challenge as there is no straightforward comparison between the two technologies. Let's take a look at what we do know today.**

There are no moving parts in a SSD which means that the drive doesn't need to spin to locate the block where the data is written. Heads do not have to be moved across the drive to where the data is in order to read it or write it. This means there is effectively no seek time. Why is this important?

For standard hard drives we use the seek time as one of the two primary measures of performance. Speeds differ based on the underlying drive interface SAS or SATA but the process is the same. With the exception of track-to-track movement, which can be between 0.2ms and 0.8ms, all other measurements are in the range 3ms-9ms for average seek time. With no moving parts, SSD performance is much faster, typically anywhere between 0.08ms and 0.16ms.

The other primary measurement we use is the read/write transfer rates. This is particularly important when moving large files. With standard hard drives we again get differences based on the underlying technology SAS/SATA, as well as where on the drive the data is located. It takes less time to read or write data on the outside of the disk than the inside.

For SSD there is no such 'inside/outside' distinction. There is also a marked difference between the read and write operations. It has taken some time, but the very latest generation of SSD now match or exceed the read performance of both SATA and SAS with speeds in excess of 120MB/sec.

The major challenge has been on the write cycle and until recently, SSD's were less than 50% the write speed of standard hard drives. That has also changed. Latest generation SSD's boast write speeds of 170MB/s which is better than both SAS and SATA hard drives.

Despite this, there are still things to be taken into account. Speed is not everything. This is not about a Ferrari versus an Aston Martin. Disk performance is linked to various issues including applications, the data bus, the way the drives are connected together in arrays, the size of the data block and the file system.

In a read operation the most important of these is the size of the data block. Smaller blocks are read much faster with SSD than with standard hard drives. As block sizes increase, the differential between the technologies reduces. However, the overall performance for read is still in favour of SSD.

Write operations are less cleanly differentiated. SSD doesn't just overwrite sectors in the way a hard drive does – instead it merges the data with what is already there using the erase block. This takes away performance from the SSD. At the moment, vendors are looking at what can be done to improve this and make a drive faster. As a result, the more writes a SSD does, the slower its overall performance becomes.

Applications tend to have preferred block sizing for optimal performance. Those applications that are predominately reading data, such as video streaming, will make big gains from the introduction of SSD, even with the large block size they prefer. Applications such as transactional banking systems that are almost all write operation where small blocks may not be as successful.

The file system is extremely important in how we address storage. It determines the default block size and how the data is read/written to the storage media. A lot of work is currently going on to see how file systems need to be changed in order to

take the best advantage of SSD. One issue is around transactional file systems, such as Microsoft's NTFS, that use a log file to track changes. This means that there are lots of small writes which will slow down the first generation of SSD.

There are file systems beginning to appear that are optimised for working with SSD. But these are unlikely to make much in the way of mainstream impact, unless any of the Linux or Unix vendors adopt them. Microsoft has just delivered Windows Server 2008 and is not planning any modifications to the way that NTFS works for SSD. That said, there is a new version of the Windows Storage Server due soon. Depending on how Microsoft sees the take-up of SSD and the feedback from those partners who deliver devices using Windows Storage Server we may see a change before the next revision to the operating system.

One way that we massively speed up hard drives is by assembling them into disk arrays. These arrays allow for data to be spread over multiple disks so that when we read/write information we are pulling just a small amount from each drive in the array. The logic is that the more heads and spindles you have, the faster the performance.

We have yet to see the first drive arrays ship using SSD. Several drive controller vendors have expressed caution about the potential for SSD arrays and believe that we may need another technology revision to get the same performance uplift as that achieved with hard drives.

So in conclusion, the currently shipping SSD's are an ideal option for the latest generation blade servers where they are running the operating system and applications software. HP's BL495c blade server is the first blade server to benefit from SSD's with two clear benefits: a power saving v traditional drives and more efficient use of space inside the blade allowing the doubling of the number of memory slots available in a half-height blade. As SSD technologies mature further the benefits will extend far beyond the server.