Understanding Solid State Drives (part one)

Solid State Drives (SSD) are being promoted as energy efficient, fast and the next generation of storage technology. Having first made an appearance in laptops, they are now beginning to be deployed in servers and storage arrays. Over the next three articles, we will take a look at the technology, the benefits and why you should be considering them.

Anything that can lower the energy usage in a datacentre has to be given consideration. Dense computing through blade systems and virtualisation have been the two key technologies for the last three years in this category, but costs are still running ahead of budgets. Now there are other ways to reduce cost.

This next generation of components is something we have talked about in previous articles, and include low wattage processors, more efficient power supplies and low power memory. Low wattage processors have the most impact of cost reduction per component. Reducing the power in an AMD Quad Core processor from 75W to 55W reduces power costs by over 26%. Not all applications will be suited to the lower wattage processors but this is still something we can improve today.

Low power memory - not to be confused with low voltage memory, which is still not mass market - can save up to 2W per DIMM socket. With motherboards now supporting anywhere from 8-24 DIMM sockets, this adds up to a significant saving. Changing the size of the memory chip can also save a significant amount of power. A more expensive 8GB DIMM uses 11.4W. Compare this to a pair of 4GB DIMM, each using 10.7W, and for the same amount of memory you can reduce the power required by over 46%.

Changing the power supply to one that ramps up the efficiency faster and holds that curve will also save power. In a dynamic datacentre where technology is powered up and down based on usage, efficient power supplies will cope with a wider spread of workload.

All of these do come at a cost, but if you look at the differences over a three-year life of the hardware, a slightly higher CapEx now is offset by the OpEx saving from consuming less power. So how can we take this further?

The answer is Solid State Drives - SSD. These are drives using the same technology that you find in a flash pen drive, but with much more sophistication. At the moment, they are relatively expensive but the prices are almost halving year on year. As more and more vendors begin to offer SSD the prices will drop further. But what is SSD and why should we use it?

What is SSD

SSD is a collection of memory chips, combined with additional technologies that allow it to appear as a normal hard drive to the operating system. While this same description also applies to a USB pen drive, a SSD is more than a pimped up USB device. It contains technology that increases the longevity of the drive and improves the performance.

The technology was invented in the 1980s and the first NAND flash memory was manufactured in 1989. Since then, it has gone on to appear in countless devices such as USB memory cards, mobile phones, cameras, MP3 players, laptops and now, hard disks inside servers and storage arrays.

There are several different generations of NAND - SLC (single-level cell), MLC (multi-level cell) and now TLC (triple-level cell).

How it works

SSD work by writing data directly to NAND flash memory. The data is written into the cells on the chip, one bit at a time. SLC stores just one bit, MLC stores two bits and TLC stores three bits. There has been a suggestion that MLC should change its name to DLC (double-level cell) or that TLC should be renamed MLC v2.

The cells are arranged into pages and these are then grouped into blocks. This is similar to a block on a hard disk with vendors using a variety of page and block sizes. Pages range from 512 bytes through to 4,096 bytes and blocks contain anywhere from 32 to 128 pages, meaning that a block can be anywhere from 16kB through to 512kB.

This block size is important for the way that NAND-based SSD work. While the data is written one bit at a time, deletion is done block by block. The frequency at which the data changes will impact the amount of activity and available space.

There are other similarities between block sizes on a SSD drive and a standard hard drive. Both use sections of the drive to record bad blocks. When the performance of a particular block falls below a certain level, it is marked bad and the data moved away.

Writing data to SSD is done sequentially, just like magnetic tape or optical drives. Compare this to a standard hard drive where the disk technology has moved from just writing data in the first available location. Now large blocks are written together to improve the speed with which the data can be retrieved. On drives that are written to regularly, defragmentation programs are used to reassemble files that have been spread all over a drives surface.

In this article we have looked at the SSD technology and benefits, particularly in terms of power savings. In part two of this article we'll take a look at the topic of SSD performance, Are they faster than current drive technologies? Can they deliver data more efficiently? In what environments does SSD technology make sense?