

With flash, it's not just about the software

Clive Longbottom. Service Director

Quocirca Comment



Disk is an ageing technology

The magnetic disk was invented over half a century ago. Apart from higher storage densities and spinning speeds, not a lot has changed over that period. It is time for a new prime storage medium to take over – and flash fits that bill.

The leap in performance from spinning disk to flash is such that anyone not looking to shift primary workloads over to a flash-based system risks being left behind as their competition makes the move.

The perceptions of flash

However, perceptions remain regarding possible issues with flash-based media. The two main ones are around longevity and price.

Let's take a look at both of these in turn.

Firstly – longevity. With a spinning disk, its main failure mode is due to catastrophic failure. A head touches a disk platter due to vibration, or an actuator fails. The drive becomes completely useless and needs replacing. Sure, sectors on the disk can become corrupted and unusable, but modern arrays can automatically compensate for this and mark those sectors as unusable – apart from where this is an MBR (Master Boot Record) sector.

With flash-based media, there are no moving parts. Catastrophic failure is not going to happen. However, flash can only be written to a certain number of times. It can be read from pretty much as many times as anyone wants, but the act of writing or erasing data (known as a program/erase or P/E cycle) leads to the degradation of a memory cell until it can no longer successfully go through another P/E cycle. Different types of flash memory have different P/E cycle characteristics – but for the purposes of this article, we'll leave that for now.

What is true is that every flash-based array manufacturer is dealing with the same basic problem – flash will eventually wear out. However, there are a range of things that can be done via software and firmware to improve

matters. This is where the battle for flash storage supremacy has been fought to date.

Without going too far into the technical aspects, there are areas such as wear levelling, garbage collection and write amplification (Google them if you want to be totally confused) which can be manipulated to improve the longevity of the flash substrate. This has led to the emergence of highly proprietary approaches to the management of flash systems, aimed at ensuring that P/E cycles are equally distributed across the whole of the flash available and in making sure that P/E cycles are managed to the greatest possible extent.

This means that whereas first generation flash systems tended to be guaranteed for around 3 years or less, newer flash arrays are being run that will happily exceed seven years usage. Violin Memory, an early innovator of flash storage systems, for example, has never had a flash system wear out in the ten years of its history.

Smoke, mirrors and pricing

Secondly, from a price perspective, it is still true that at a byte-for-byte level, flash storage is still considerably more expensive than its spinning magnetic equivalent. Again, though, flash vendors have addressed this and focussed on how they can provide an effective approach to precisely managing how the data is stored on the flash media to ensure that on an actual data volume under management basis, the cost of flash is around the same for an equivalent magnetic disk array while providing better overall benefits and value.

“on an actual data volume under management basis, the cost of flash is around the same for an equivalent magnetic disk array while providing better overall benefits and value”

Sounds a bit like smoke and mirrors? For some vendors, it has been quite close to being the case. However, the leading flash storage array vendors can demonstrate improved overall value of a

known, specific storage workload on flash than on spinning magnetic disk.

How do they do this? Again, through advanced software. The deployment and use of technologies such as in-line data de-duplication and compression means that data volumes are reduced before they are stored at the flash layer. Additionally, the adoption of approaches such as thin provisioning and thin cloning means that usage of the available storage resources is highly optimised.

Over and above the software and firmware needed to manage a specific storage array is the software required to manage a storage infrastructure. In this space, Quocirca sees a much wider difference between offerings, with some point vendors not having a strong story in how a mixed and distributed storage system can be effectively managed to provide an optimised environment. A fabric-based approach, which is workload aware and can interoperate with the compute and network aspects of the overall IT architecture, is required.

In the 'dash to flash', it is therefore apparent that the main battleground amongst the flash storage vendors has been around who has the best software. The software helps decide overall longevity and performance of the underlying flash, and will also have a deep impact on how cost-effective the flash is when compared directly with an equivalent magnetic disk array.

It's not about the software...

So, it is all down to who has the best software, then?

Well, no. As with the physical magnetic disk, all flash media is not equal. What a flash storage system vendor chooses to use is important. Super-fast flash, such as Single-Level Cell (SLC), is very expensive but can withstand a lot more P/E cycles than other flash. Multi-Level Cell (MLC), as found in consumer devices such as tablet computers, cameras and so on, is relatively cheap – but is slower and has around a third of the longevity of SLC. Triple-Level Cell (TLC) was touted for a while as a very cheap flash approach, but its longevity is very low and its performance lags that of MLC and SLC.

The focus has been on MLC (and its 'enterprise' version eMLC – although the move for vendors

has been to make the most of the economies of scale of using pure consumer flash).

There are only a few original equipment manufacturers of flash-based disks (solid state disks, or SSDs). Therefore, it is difficult for those using SSDs as the basis of their storage arrays to differentiate at the SSD level. Thus, some vendors have moved towards a more self-contained system, including not just the storage array, but also compute and network resources to create a 'hyper-converged' storage system. However, the majority of these systems are still based on the concept of the disk – technology essentially over half a century old.

...Or the hardware

Where Quocirca sees greater innovation is through the use of different storage architectures: something which avoids the constraints of disk-based logical unit numbers (LUNs) and standards that were created for a world of spinning magnetic platters, actuators and heads. Flash works in a different way, and much greater use can be made of its capabilities through taking a different approach to how it is utilised.

The first layer of intelligence in how flash is used as a storage medium is through the software, as discussed above. The second layer is through the flash controller, which includes a flash translation layer that essentially takes what the client side expects to see (a magnetic disk) and ensures that the flash storage responds accordingly.

“we do not write to the disk as if it were a sequential storage system for our primary storage, as it would not meet the needs of a modern business.”

However, there is no real reason why the future should try to adhere to an outmoded storage approach – after all, although we can write to magnetic disk as if it were magnetic tape via virtual tape libraries (VTLs), we do not write to the disk as if it were a sequential storage system for our primary storage, as it would not meet the needs of a modern business.

Therefore, Quocirca has seen the emergence of different storage architectures based on flash, from PCIe server-side flash through to Violin Memory's Violin Intelligent Memory Modules

(VIMMs). Such a break from the concept of a disk drive offers much for the future – workloads can be optimised specifically for flash without being constrained by needing to utilise SATA-, SAS- or iSCSI-based calls. Those wedded to the SSD commit themselves to continuing to support the concept of the disk – and will constrain their performance capabilities in the long term.

The right tool for the right job

So, the use of flash as a storage medium is not as straightforward as it at first appears. Different storage workloads have different needs: the flash substrate chosen is increasingly providing little differentiation as MLC becomes the basic substrate of choice. However, magnetic disk is not up to competing with flash as a future medium, and even SSD-based approaches are dependent on backward-looking technology rather than being truly innovative.

The future battle for storage supremacy is around who can build the most effective overall storage system.

One that is built for purpose at the hardware, software and firmware levels; one that is contextually aware of the workloads being thrown

at it; one that can operate as a peer system within a broader data management platform.

Storage is never going to be simple, yet hiding the complexities from the user is key. With systems built up from essentially disconnected hardware and software, this is never going to happen. By choosing a system where the flash has been architected from the ground up to work with the rest of the hardware around it – cpus, interconnects, networks – and has the firmware and software built specifically to ensure optimised operations, storage becomes a business facilitator, rather than a constraint.

For the buyer, flash must be the obvious choice. However, the questions remain as to what type of flash, how it is architected and built and in what manner is it pulled together by intelligent software at the array and platform levels. This is where the future lies: those selling flash-based storage have to demonstrate and provide on the promises they make for storage workloads not just now, but also for the future.

Flash is not just about the software; nor is it just about the hardware: it is all about the mix.

About Violin Memory

Violin Memory transforms the speed of business with continuous data protection through high performance, always available, low cost management of critical business information and applications.

Violin's All-Flash optimized solutions accelerate breakthrough CAPEX and OPEX savings for building the next generation data centre. Violin's Flash Fabric Architecture (FFA) speeds data delivery with chip-to-chassis performance optimization that achieves lower consistent latency and cost per transaction for Cloud, Enterprise and Virtualized mission-critical applications. Violin's All-Flash Arrays empowered by our enterprise data management software solutions enhance agility and mobility while revolutionizing data centre economics.

Founded in 2005, Violin Memory is headquartered in Santa Clara, California.

Further details are available at <http://www.violin-memory.com>



About Quocirca

Quocirca is a primary research and analysis company specialising in the business impact of information technology and communications (ITC). With world-wide, native language reach, Quocirca provides in-depth insights into the views of buyers and influencers in large, mid-sized and small organisations. Its analyst team is made up of real-world practitioners with first-hand experience of ITC delivery who continuously research and track the industry and its real usage in the markets.

Through researching perceptions, Quocirca uncovers the real hurdles to technology adoption – the personal and political aspects of an organisation's environment and the pressures of the need for demonstrable business value in any implementation. This capability to uncover and report back on the end-user perceptions in the market enables Quocirca to advise on the realities of technology adoption, not the promises.

Quocirca research is always pragmatic, business orientated and conducted in the context of the bigger picture. ITC has the ability to transform businesses and the processes that drive them, but often fails to do so. Quocirca's mission is to help organisations improve their success rate in process enablement through better levels of understanding and the adoption of the correct technologies at the correct time.

Quocirca has a pro-active primary research programme, regularly surveying users, purchasers and resellers of ITC products and services on emerging, evolving and maturing technologies. Over time, Quocirca has built a picture of long term investment trends, providing invaluable information for the whole of the ITC community.

Quocirca works with global and local providers of ITC products and services to help them deliver on the promise that ITC holds for business. Quocirca's clients include Oracle, IBM, CA, O2, T-Mobile, HP, Xerox, Ricoh and Symantec, along with other large and medium sized vendors, service providers and more specialist firms.

Full access to all of Quocirca's public output (reports, articles, presentations, blogs and videos) can be made at <http://www.quocirca.com>