



Faster, quieter and more reliable

Flash memory is used everywhere and while it may have its disadvantages it looks set to become the storage technology of choice. By **Peggy Lee**.

Faster, quieter and more reliable, with better performance. Flash memory based solid state storage devices (SSD) are pushing hard disk drives out of the market and could, one day soon, replace them all together, according to Robert Fan, president of NAND controller developer Silicon Motion.

“SSDs are taking over the whole category now and we’re seeing with Seagate and Western Digital that their hard disk market is declining,” he asserted.

“And flash is now used everywhere – phones, cameras, Satnavs, USB sticks ...,” said Tom O’Neill, chief technical officer at all-flash storage company Kaminario.

Flash stores information in an array of memory cells made from floating-gate transistors. The two

types of flash memory are NAND and NOR, named after logic operations. NAND flash may be written to and read from in blocks. NOR, meanwhile, allows a single byte to be written or read independently.

Fan points to four benefits from using NAND, rather than hard disk drives. The first is improved performance as flash moves the data quicker. The second is higher reliability as it is not mechanically-based, and better durability to withstand shocks, high pressures and temperatures. No moving parts also means less noise and, finally, he claims it consumes much less power.

Silicon Motion offers a range of NAND controllers for SSDs. According to Fan, a flash controller provides the interface between the host controller

or processor and the flash memory device.

“The controller interprets the host microprocessor’s instructions; for example, ‘accept these 10kbit of information for storage in 10k locations’, and it decides where in the flash memory to store the information. It then retrieves the information when requested.”

Flash controllers will also organise the way data is stored to achieve ‘wear levelling’. The locations in a flash memory will gradually wear out with repeated write/erase cycles, so using the available memory cells equally is important to avoid premature failure.

Applying single- and multi-cell technologies in its products, Silicon Motion is also exploring 3D NAND. According to Fan: “3D NAND is the latest generation of flash technology. It enables very high capacity and very low cost. All the NAND vendors – like Samsung, Toshiba – are moving towards supplying 3D NAND.”

Shannon Systems – a Silicon Motion subsidiary – has taken advantage of 3D NAND in Direct-IO G4i, which it calls the first 12.8Tbyte PCIe SSD drive.

“3D NAND technology heralds a new era of flash storage with high performance, high capacity and low cost being available at the same



“[NVMe] allows you to talk to the storage medium more efficiently, with higher levels of parallelism and lower latency.”

Tom O’Neill

time,” said Shannon’s CEO Yang Xueshi. “By using 3D NAND, we can store 12.8Tbyte on a single SSD drive. This is a milestone for the SSD industry.”

Alongside the Direct-IO PCIe SSD upgrades, Shannon also unveiled a 3D-NAND based NVMe-based SSD line. These products offer access latencies of less than 15µs, more than 2Gbyte/s read/write bandwidth, while consuming less than 6W.

While much has been written about the benefits of SSDs in enterprise applications, the technology is also finding uses in other areas. “One of our key growth areas in the past couple of years has been in the client SSDs,” Fan noted. Client SSDs are built for the consumer market, rather than for enterprise applications, and may provide fewer write cycles.

Fan says SSDs are now finding application in many sectors, ranging from automotive and medical to factory automation and point of sale systems.

Targeting the enterprise sector is Kaminario, whose systems store large amounts of data and are targeted at organisations such as banks and hospitals. Its product, K2, can scale from a 4U rack array with two controllers and one shelf of flash media up to eight controllers and 16 shelves of flash media.

“You can start with, say, 15Tbytes and that system can scale non-disruptively to store 4Pbyte,” said O’Neill.

O’Neill is proud of the company’s scale-up or scale-out independent capability, meaning it can scale both the capacity, or the performance, or both independently.

“K2 v6 offers a true scale-out design to increase the performance linearly by adding more and more K-Blocks,” he explained. “As the system grows, it can maintain growing numbers of throughput and input/output operations/s (IOPS) while keeping the same low level of latency.”

The advent of SSDs means the performance bottleneck has moved away from the media – HDDs – to the CPU and networking, O’Neill contended. “However, the scale-up approach is no longer sufficient. Scale-out is not only about performing on a large scale, it’s also about starting small and growing with your needs.”

The latest version of K2 – called v6 – is said to increase scale-out performance linearly by adding K-Blocks. To enable scale up and scale out, Kaminario has separated the storage layer, where the data is kept, from the compute layer, where the metadata is kept.

“Because we have this separation in our storage design, it has enabled us to benefit from the possibilities offered by NVMe over Fabric connectivity and separate the layers,” explained O’Neill.

NVMe – which has been around since 2011, is the first protocol designed specifically to handle non volatile storage device attached via PCI Express. “This protocol allows you to talk to the storage medium more efficiently, with higher levels of parallelism and lower latency,” O’Neill added.

One of the benefits from adopting this separation is that it enables new capabilities, such as disaggregation and dynamically composable systems. A ‘disaggregated’ design, for example, enables a system to be recomposed to meet changing requirements. O’Neill gave the



“SSDs are taking over the whole category now and we’re seeing [the] hard disk market declining.”

Robert Fan

example of a system being used in a retail environment in the run up to Christmas.

“With our system, you can allocate compute power non-disruptively from a different part of the organisation to cope with the extra workload without making any physical changes or without having to do any planning or move things around physically,” he said.

To reduce the risk of failure when the memory reaches the end of its allocated write/read cycles and to protect against unrecoverable block read errors, Kaminario has created its own patented version of a random array of independent disks (RAID) called K-RAID.

Looking to maximise efficiency and performance, along with high levels of protection, Kaminario has implemented an automated pyramidal RAID scheme with 87.5% utilisation in its K2 v5 system.

“Out of a stripe of 24 RAID devices in an array group,” O’Neill noted, “three are used for parity protection, so we can withstand, in a single RAID stripe, up to three components failing, while still not losing data.”

According to O’Neill, K-RAID uses compression and de-duplication technologies to avoid keeping multiple copies of the same data. There is a risk, however, that when doing de-duplication, the technology will mistakenly believe it’s found a duplicate and delete the data. This called a flash-collision.

“We’ve decided not to take that risk. Instead, every time the system thinks it’s found a duplicate, it rereads the original block of data, rather than just relying on the metadata,” O’Neill explained.

For all its benefits, flash has disadvantages, including the limited number of read/write cycles. However, companies such as Kaminario and Silicon Motion are working to improve this and flash memory is set to become the storage technology of choice for a wide range of applications.

FerriSSD integrates Silicon Motion’s controller technology with a NAND flash device and passive components into a BGA package capable of storing up to 256Gbyte of data

