Solid-state disks stage a comeback
By Mark Ferelli

Traditional hard disk drives (HDDs) have been the storage technology of choice for decades. They dominate in primary storage from laptops to enterprise arrays. Competing technologies usually encounter skepticism at best, and derision at worst. However, scoffers are becoming less vocal when solid-state disk (SSD) technology is mentioned. Historically relegated to niche markets and applications, SSDs are increasingly making their way into mainstream data-center systems and applications, as well as laptops.

The SSD market is poised for rapid growth, according to Jeff Janukowicz, research director for SSD at International Data Corp. (IDC). Revenues are expected to surge from $77 million this year to $1.2 billion in 2011-a 76% compound annual growth rate.

Janukowicz identifies several drivers that are fueling OEM and end-user adoption of SSDs. In addition to the obvious performance advantages, declining prices are one key factor. The surge of interest in energy conservation is another key driver behind renewed interest in SSDs. Since SSDs have no mechanical parts, energy and cooling costs drop significantly, making SSD technology a natural for “green” environments.

The basics

IDC defines an SSD as “a semiconductor-based block storage device that behaves as a virtual HDD and appears to the host device as a disk drive.” SSD technology includes sub-categories such as SDRAM (volatile) and NAND flash (non-volatile).

SSDs based on volatile memory such as SDRAM are characterized by extremely fast data access times-less than 0.01 millisecond-and are used primarily to accelerate applications that would otherwise be held back by the latency of disk drives.

SDRAM-based SSDs typically incorporate internal battery and backup disk systems to ensure consistent data access. If power is lost, the battery keeps the unit powered long enough to copy all data from RAM to disk.

When power is restored, data is copied back from disk to RAM and the SSD resumes normal operation.

Many SSD manufacturers use non- volatile NAND memory to create more rugged and compact alternatives to SDRAM-based SSDs. These flash memory-based SSDs, also known as flash drives, do not require batteries, allowing vendors to replicate standard disk drive form factors (i.e., 1.8-, 2.5-, 3.5-inch).

In addition, non-volatility allows flash SSDs to retain memory even during sudden power outages, ensuring data retrievability. Like DRAM SSDs, NAND SSDs are very fast since these devices have no moving parts, eliminating seek time, latency, and other electro-mechanical delays inherent in conventional disk drives (although flash SSDs are significantly slower than SDRAM SSDs).

The NAND segment of the market is separated into two sub-segments: single layer cell (SLC) and multi-layer cell (MLC) formats. The principal difference between the two is that SLC stores a single bit of data per cell whereas MLC, by using a larger number of voltage levels, stores two bits of data per cell. SLC, the higher-performing technology, enables streaming video and Internet capability in handheld electronics. MLC NAND flash, on the other hand, is generally selected for shorter lifetime, price-sensitive applications. Since MLC stores twice the data per cell, it is also considered the capacity leader.

SSD in the enterprise

“SSD has near-term opportunities in the enterprise,” says Patrick Wilkison, vice president of marketing and business development at STEC. The NAND-based SSD vendor points to a number of factors leading to enterprise market acceptance.

“Improvements in random performance are increasingly important, and SSD eliminates latency and significantly reduces data access times. Aggregating short-stroking drives is often sub-optimal,” says Wilkison.

Traditional HDDs have improved steadily from the mid-1950s, when 5MB drives had rotational speeds of 1,200rpm. Today, HDDs pack up to 1TB and have rotational speeds up to 15,000rpm. Capacity has grown, but what about I/O?

Woody Hutsell, executive vice president at Texas Memory Systems, points to performance tests that measured HDD data access times at 4 milliseconds (msec) for both reads and writes, 0.2msec for NAND flash read performance, and 0.02msec for SDRAM-based SSD technology. SSD’s random I/Os-per-second performance advantage over HDDs ranges from 5x to more than 2,000x, depending on the specific technology.

STEC's Wilkison also points to low power consumption as an advantage of SSDs. At small capacities, flash SSD offers the lowest power consumption at 0.5 watt. At larger capacities (more than 150GB) flash SSD and HDD draw up to 18 watts, with the interface type impacting power consumption. The annual average utility cost of a 100,000-sq-ft data center is almost $6 million, with power and cooling the major villains. This is a non-trivial operating cost, illustrating budgetary considerations well beyond the high acquisition cost associated with SSDs. Wilkison says cost per gigabyte will be only one cost consideration. Cost- per-I/O and cost-per-power-increment should also be calculated when you are evaluating storage options.

Although there is increased awareness of the benefits of SSDs among end users, OEMs may ultimately be the driving force behind SSDs in the enterprise. Ofer Tsur, head of the SSD operation at SanDisk, points to the company's recent OEM agreement with IBM as an example of increased OEM interest in SSDs. IBM offers SanDisk's 2.5-inch SATA 5000 flash SSD drives in its Blade- Center HS21 XM servers.

"The IBM agreement shows the value of SSDs in blade servers," says Tsur, "and as more density appears in silicon, it will be possible to reduce prices even further for OEMs." Tsur points to applications such as video-on-demand and read-only databases as key applications for SSDs.

Many SSD products are designed to fit slots traditionally filled by HDDs, but Tsur foresees a day when OEMs will design servers specifically for SSDs.

Mix n'match

Mixing SSD drives with traditional HDDs is expected to grow in popularity as SSD prices continue to decline. According to channel consultant Diamond Lauffin at the Lauffin Group, "Using SSDs with SATA arrays can provide an affordable alternative to traditional arrays." Such a combination can make sense by, for example, aggregating the index files and search indexes on solid-state memory while putting the application data on spinning media.

"Blended" HDD-SDD architectures like the ones Lauffin describes will become more common as the popularity of SSD grows. But integrating different types of SSD technologies may also be an emerging trend.

Texas Memory's Hutsell points to an upcoming "cached flash" device, which combines DDR SSD technology with flash technology (see "Get ready for 'cached flash' drives," p. 12).

Regardless of packaging issues, SSD is moving out of niche markets (most notably, military applications) and is taking a place in the enterprise data center. ❑