

A Chip to Slash the Power in Data Centers

Michael Kanellos | July 20, 2009

Flash memory drives use far less power than conventional drives. The problem has been the price and short life. A team of silicon vets at **Sandforce** claim they've cracked the code.

It's like taking iPod economics to the datacenter.

Sandforce, a startup in Saratoga, has devised a processor that it claims will effectively allow search companies, banks and other companies with large datacenters to swap out storage systems made out of hard drives with drives made of flash memory, which only use about 5 percent of the power.

And not any kind of flash memory. Sandforce's chip is particularly designed to allow multi-level cell (MLC) flash to take on corporate jobs. MLC flash can hold two or more bits in a memory cell and is the cheapest kind on the market. It's the stuff you find in cameras and MP3 players.

"Storage will be free if they [corporate owners] take the total cost of ownership into account," said CEO Alex Naqvi. "Over five years, they amount of power saved will pay for itself in five years."

A traditional datacenter with 240 high-end drives holding approximately 73 gigabytes each could be replaced by one with 9 Sandforce-enabled drives, he said. The energy cost for storage in the traditional datacenter would cost \$50,000 over five years. The Sandforce-enabled one would have a cumulative energy bill of \$250.

Put another way, total storage costs per gigabyte, including energy, would come to \$3.16 for hard drives and \$1.43 per gigabyte. The math assumes fewer flash drives would be needed because of better input/output, he argued. On a straight comparison, a high-end 73-gigabyte hard disk would need \$211 in power-including air conditioning power-over five years. A 73-gigabyte flash drive would get by on \$52 of power.

Considering that Best Buy sells several terabyte hard drives, you can imagine the storage needs of someone like Google. Meanwhile, IBM, Hewlett-Packard, Sun Microsystems, Cisco, Dell, EMC and others have all also sketched out plans to devise flash-based storage products.

By 2013, solid state datacenter drives could curb energy consumption by over 57,000 megawatt hours, according to iSuppli, assuming that they would be installed only 10 percent of the time.

Naturally, hard drive manufacturers will hotly contest this and pour over the beta results reported by Sandforce. The first chips came out in April and the first complete drives made with its SF-1000 processors come out in the fourth quarter.

Nonetheless, the company seems to be generating buzz and interest in the industry. Founded in 2007, it has raised money from DCM, Storm Ventures and two unnamed storage vendors. It will sketch out its plans at the Flash Memory Summit, which begins August 11 in Santa Clara.

Sandforce is also one of the few Silicon Valley companies hiring. The person who greeted me at the door thought I arrived for a job (not a news) interview.

"I'm interviewing at least 10 guys a day," said Naqvi.

The key to Sandforce's processors is that they effectively mask the technological shortcomings of MLC. Like all flash memory, MLC essentially stores data by brute force. A flash memory cell consists of a sealed silicon dioxide, or glass, tube, similar to the flask inside of a thermos. To write data to a flash cell, an electrical charge is applied that leaves a set number of electrons inside the sealed tube. To erase data, another charge is applied. Flash engineers describe the process as violent.

The problem is that the tube is getting thinner and smaller with each turn of the Moore's Law crank. Flash chips made on the 20 nanometer process, coming in a few years, may only accommodate about 20 electrons, said Naqvi. Escaping electrons in turn leads to errors and data corruption.

"Every time you hit the oxide, you damage it," he said. As a result, a conventional flash drive (without the company's chips) would have to be replaced every 30 days because of the number of data rewrites.

The company's DuraClass technology gets around the problem by only rewriting the cells that require an update. An entry in a Social Security file, hypothetically, might contain 100,000 bytes. An update to the data might only affect 100 bytes. Traditional flash drives would rewrite the whole entry. DuraClass organizes the data and functionality of the chip so that only 100 bytes get impacted. In all, a DuraClass-enabled drive could last over five years with equal or better performance when it comes to data reliability under ordinary corporate pounding.

The overall market for flash drives is expected to come to around \$10 billion by 2012, up from about half of that today, according to statistics from Gartner. The semiconductor element of the drive market could account for one-tenth of the revenue.

"This is the only market in the last few years that is over a billion dollars," he said.

Will competitors come along? Sure, but designing a processor can take three to four years. Moreover, the chip is targeted at flash memory makers, which have been brutalized by price cuts. There may not be a rush of followers.

Although it ultimately will aim at data center storage systems, Sandforce initially will see its chips go into flash drives for notebooks. The qualification cycles are less stringent. While the chip will work in flash drives for notebooks, it won't shine. Conventional notebook drives are actually quite advanced right now. Paying off the flash drive through energy savings would take about 50 (that's five zero) years, he acknowledged.

Correction: We misspelled his Naqvi's name originally.