SandForce™ Enterprise Solid State Drive Processor with DuraClass™ Technology
SSDs Come of Age

By James E. Bagley
Consulting analyst and
Deni Connor
Principal analyst, Storage Strategies NOW

The Billion Dollar Bet is when (not if) Solid State Drives (SSDs) will be the primary high speed storage in enterprise and mobile computer applications. SandForce has moved that date much closer than conventional wisdom would predict.

SandForce is a start-up financed by private equity (Storm Ventures, DCM and leading storage firms) coming out of stealth mode in Q2 2009. A self-described fabless silicon purveyor, SandForce is about to dramatically change the price-speed-performance assumptions related to storage.

The basic premise behind SandForce is that ultra-cheap commodity multi level cell (MLC) Flash Devices (heretofore relegated to throw-away USB keychain memories and the like) can be reliably integrated into the base fabric of enterprise server and storage technology. While all currently deployed enterprise SSDs utilize the more expensive single level cell (SLC) flash silicon, SandForce's strategy is to provide products that support both devices, with price performance emphasis on MLCs. To do this, numerous critical problems had to be solved, with extensive research and development costs and corresponding investment.

These problems are complex but solvable. The founders of SandForce, starting with the knowledge of the technology necessary, have been developing this science into working products for nearly three years and the results are impressive. To begin with, commodity flash (MLC) represents tremendous advantages at the cost per GB, the speed of operation and energy consumption data points. The problems with these devices have to do with three major areas: endurance, retention and reliability. Throw in the security issue related to current SSD implementations (no encryption) and you have big problems for storage devices that need to be able to read what they have written. Let's delve into each problem and the science, dubbed 'DuraClass,' by SandForce, involved with the solution.

Endurance
Flash memory has a finite number of program-erase cycles per addressable area. The devices literally wear out as you write to them. Strategies for dealing with these characteristics include wear-leveling and daily write restrictions. The latter is a simple matter of limiting the number of writes the user is permitted to write in one day or else you void the warranty. DuraClass works to reduce the "write amplification" factor of SSDs with a proprietary technique called DuraWrite™. The flash memory can operate about 80 times longer compared to standard controllers and enables MLC memory to operate in high duty-cycle enterprise environments. DuraClass also eliminates daily write restrictions.

Write Speed and Overall Performance
We have analyzed the published specifications of various SSD implementations and find many vendors shy away from publishing the extreme performance differential between reads and writes. This differential is typically from 30K IOPS read rate to 3K IOPS write rate, a 10x decrease. Further, existing flash implementations may have a quick first time write, but a huge decrease in subsequent writes. This erodes SSD

write performance by an enormous factor. The problem is generally caused by a high write amplification factor forcing the SSD to write so much additional data that it cannot maintain a high throughput. The solution implemented with DuraClass results in greatly reduced write amplification and enables much higher throughput for writes. The initial release of products includes the SF-1500 Processor for enterprise applications and the SF-1200 Processor for mobile use with the supported speeds as shown:

OPERATION	ENTERPRISE (SF-1500)	MOBILE (SF-1200)
Sequential Read 128KB	250MB/sec.	250MB/sec.
Sequential Write 128KB	250MB/sec.	200MB/sec.
Random Read 4KB	30K IOPS	5K IOPS
Random Write 4KB	30K IOPS	5K IOPS

Retention
Somewhat like that old Lithium Ion Battery in your cell phone that appears to recharge in a really short time and keeps the phone running for a similarly short time because it has been 'cycled out', the areas in flash that have been subject to wear retain their information for a shorter and shorter period of time. DuraClass algorithmic assessment of each area and periodic refreshment maintain the retention of data for the full use-life of the system, which is warranted at five years regardless of duty cycle, or how much data is written.

Reliability
SandForce uses a very advanced ECC engine that enables even the low cost MLC devices to operate at a dramatically low failure rate of less than one unrecoverable read error per 3.9 years. This is as compared to 14.1 days for current SSD implementations given the same duty cycle and conditions. They have also coined the term RAISE™ (Redundant Array of Independent Silicon Elements) which is an automated recovery mechanism, similar in concept to RAID that can recover user data in the case of a flash block or even a whole die failure. In a nutshell, the odds of losing data to component failure are reduced by a factor of 100x. Variable forward error correction, which increases error correction data based on algorithmic analysis, provides reliability that cannot be matched by other SSD solutions today and even surpasses the correction ability of rotating magnetic media.

Security
SSD implementations have been identified as security risks because the encryption in current implementations is external to the devices. DuraClass features encryption of all data at AES (128bit) level without impact on performance.

Interface and compatibility
The SF-1000 Family of SSD Processors implement the DuraClass features for both enterprise and mobile applications. Implementing Serial ATA 2.6 at 3Gb/sec. and 1.5Gb/sec. (SAS available through third-party bridges), S.M.A.R.T. command transport, integrated data security and compatibility with most major flash memory manufacturers (for both SLC and MLC), the SF-1000 family processors are available in both 361pin-TFBGA and 128pin-TQFP packages. Incorporating features to ease OEM integrations, the SF-1000 series include I^2C, RS-232 Serial Debug Port, numerous GPIO and a JTAG Boundary Scan Port.

Over-Provisioning
All SSDs today use the difference between GB (1024^3) and billion bytes (1000^3) to serve as spare blocks for failures over time. Beyond that difference is something called over-provisioning. It is a way to offset the reliability and wear by increasing the amount of memory installed versus that available for the user. In a current implementation at a major server manufacturer, the amount of 'extra' memory is 14GB per 50GB available, a whopping 30% of over-provisioning (on top of the GB vs billion bytes difference). While this may help the wear problem, it actually decreases reliability because there are more devices that can fail - not to mention the increased cost impact. DuraClass achieves its high performance and throughput without over-provisioning.

SSG-NOW Assessment
The pre-emptive price-speed-performance of the SandForce SSD technology will move the benchmark of entry into SSD servers and mobile devices to new, significantly lower price points. How much lower? Divide current SLC SSD device cost per GB by three and you'll be in the ballpark. Other enabling technologies, such as Gigabit Ethernet and Fibre Channel over Ethernet (FCoE) will provide low-cost super-bandwidth to take advantage of the speed of these devices. Dramatic commodity SSD cost reduction will proceed into the future (aka Moore's Law) and cause an increasing percentage of storage allocation to SSDs to be a continuous trend. The day when rotating devices are relegated to the same heap as the floppy disk drive is still not on our calendar, but it is much closer thanks to SandForce DuraClass technology.