

# IBM smartens up SSDs with automated data classification

**IBM is moving to add intelligence to solid-state drives (SSDs) by putting SSD-specific data placement and monitoring capabilities into software across its storage and server platforms.**

May 22, 2009 | **By:** Dave Raffo, Senior News Director | StorageSearch

Industry analysts and storage administrators agree that intelligent management needs to be the next phase in solid-state drive technology now that all of the major storage vendors support SSDs in their arrays. EMC CEO Joe Tucci calls it Phase 2 for solid-state drives, but IBM is the first major storage vendor to add these features widely across its product line by developing new applications and upgrading existing hierarchical storage movement software to work specifically with SSDs.

Clod Barrera, chief technical strategist for storage at IBM, said the software addresses "how one determines data placement, how you decide what data belongs on those devices, and how to move data to SSDs with minimum human intervention and minimum disruption to applications. Nobody is eager to add more complexity to their storage environment."

IBM supports Flash drives from STEC in its **DS8000** storage arrays, its zSeries mainframe servers and pSeries Unix servers. It has also been working with Fusion-io and will likely add solid-state products from the startup soon.

## Software upgrades for SSD

On the software front, IBM today disclosed it has added SSD support for its Data Facility Storage Management Subsystem (DFSMS) software that places most critical data on the highest-performing tiers of DS8000 arrays. Other applications with solid-state drive support include **Softek Transparent Data Migration Facility (TDMF)** software that migrates data between drives, and FileMon that monitors data on pSeries servers. A new SSD Data Balancer application for the pSeries lets an administrator move frequently accessed data to solid-state drives and rarely accessed data to hard drives.

Barrera said automated data migration and data monitoring will also be added to IBM software for Windows and Linux systems.

"Today it's a two-step process," Barrera said. "There are tools for monitoring and measuring, and tools that allow you to place data and migrate data while in use. Eventually, there will be integration between measurement devices and placement and migration, so all of this can be more automated."

With most storage vendors supporting the same STEC family of drives, management capabilities can be the differentiator for SSDs in the enterprise. Vendors have set this as their next priority for solid-state drives. For instance, EMC Corp. is working on what it calls FAST – fully automated storage tiering – capability for automatically moving data between tiers of SSD and hard drive storage.

"This is a direction everybody's going to have to take," said David Hill, founder and analyst at Mesabi Group. "If you use the same drives, it's how you put it together and arrange the data. Software is the distinguishing factor, the secret sauce. It has to be automated. If you do it manually, you can identify some of the hot spots, but not all of them."

## IBM's solid-state drive vision

IBM began shipping solid-state drives in its DS8000 storage systems earlier this year, more than a year after EMC said it would offer SSDs in its Symmetrix arrays. IBM has advised customers to proceed with caution on solid-state

drives because of the high prices, but the firm's Barrera said solid state makes sense in some situations today where wide striping is used.

For instance, he said firms that need fast online transaction processing (OLTP) could benefit from SSDs as an alternative to placing small amounts of data on hard drives and wasting most of the capacity to reduce response times.

"Solid-state drives today are a cost-effective alternative to that approach," Barrera said. "Not only do they give you better performance and cost, but they give you a smaller package, much better energy efficiency and better reliability."

He said when higher capacity, lower-cost multi-level cell (MLC) Flash becomes enterprise-ready – possibly next year – SSDs will begin to replace high-performance 15,000 rpm Fibre Channel and SAS drives. "We don't believe bulk drives – today's SATA – are likely to be replaced by Flash technology ever," Barrera said.

He does see non-Flash technologies such as phase-change memory and magnetoresistive random access memory (MRAM) as alternatives to SATA because of higher densities, lower cost and better latency than Flash-based solid-state drives.

"Then you might see spinning magnetic drives go away completely," Barrera said. "But those are probably at least four or five years away."