

## SDSC, UC San Diego, LBNL Team Wins SC09 'Storage Challenge' Award

*Team Highlights Flash-Memory of SDSC's New "Dash" and "Gordon" Systems*

November 20, 2009

Jan Zverina

A research team from the San Diego Supercomputer Center (SDSC) at UC San Diego and the University of California's Lawrence Berkeley National Laboratory has won the Storage Challenge competition at SC09, the leading international conference on high-performance computing, networking, storage and analysis being held in this week in Portland, Oregon.

The research team based its Storage Challenge submission for the annual conference based on the architecture of SDSC's recently announced *Dash* high-performance compute system, a "super-sized" version of flash memory-based devices such as laptops, digital cameras and thumb drives that also employs vSMP Foundation software from ScaleMP, Inc. to provide virtual symmetric multiprocessing capabilities.

*Dash* is the prototype for a much larger flash-memory HPC system called *Gordon*, which is scheduled to go online in mid-2011 after SDSC was awarded a five-year, \$20 million grant from the National Science Foundation (NSF) this summer to build and operate the powerful system. Both *Dash* and *Gordon* are designed to accelerate investigation of a wide range of data-intensive science problems by providing cost-effective data performance more than 10 times faster than most other HPC systems in use today.

The hypothesis of the team's Challenge, called "*Data Intensive Science: Solving Scientific Unknowns by Solving Storage Problems*", was that solid state drives (SSDs) based on NAND (Not AND) flash technology are "ready for prime time" in that they are reliable and cheap enough to improve input/output density (I/O rate) by more than 10 times, or greater than one order of magnitude.

"The current data storage I/O rate is far lower than the ever-increasing rate of enthusiasm among researchers and scientists, who are now drowning in a sea of data because of this differential," said SDSC's Arun Jagatheesan, team leader for this year's competition. "With the SC09 Storage Challenge, our team demonstrated the prototype of a data-intensive supercomputer that can bridge this gap. Our mission in this challenge was to design, build, deploy and commission a prototype of such a supercomputer at challenging construction and operational costs, without compromising the data-intensive performance of the scientific applications."

"A major challenge for the scientific user community is to deal with storage latency issues in our systems," said SDSC Interim Director Mike Norman, who is also the principal investigator for the center's upcoming *Gordon* supercomputer. "Even though not all scientific problems are data-intensive, many of them are, and this challenge illustrated that we can overcome latency issues with innovative approaches and partnerships. We're looking forward to helping the NSF and others meet the needs of this new generation of data-intensive science."

"Moving a physical disk-head to accomplish random I/O is so last-century," said Allan Snavely, associate director of SDSC, co-principal investigator for SDSC's *Gordon* system and project leader for *Dash*. "Indeed, Charles Babbage designed a computer based on moving mechanical parts almost *two* centuries ago. With respect to I/O, it's time to stop trying to move protons and just move electrons. With the aid of flash SSDs, we can do latency-bound file reads more than 10 times faster and more efficiently than anything being done today."

The Storage Challenge team considered several ideas, which in the process resulted in making changes to the traditional storage architecture, including both hardware and software, to achieve its goal. The changes included incorporating a large (750 GB) RAMFS (random-access memory file system) with 1 TB (terabyte) of flash SSD file system to dramatically accelerate scientific database searches such as those used in the Palomar Transient Factory database, a fully automated, wide-field survey aimed at a systematic exploration of the optical transient sky using a new 8.1 square degree camera installed on the 48-inch Samuel Oschin telescope at the Palomar Observatory in southern California.

"In the SC09 Storage Challenge, we presented the architecture of our *Dash* prototype and the excellent results we have already obtained," noted Jagatheesan. "We believe our approach provides cost-effective data performance, and we hope that other academic and non-academic data centers throughout the HPC community can benefit from our approach and experiences.

In addition to Jagatheesan and Norman, SDSC team members included Jiahua He, Allan Snavely, Maya Sedova, Sandeep Gupta, Mahidhar Tatineni, Jeffrey Bennett, Eva Hocks, Larry Diegel and Thomas Hutton from SDSC; Steven Swanson (UC San Diego); Peter Nugent and Janet Jacobsen (Lawrence Berkeley National Laboratory); and Lonnie Heidtke, of Instrumental Inc., a Bloomington, Minn.-based provider of professional services focused on advanced and high-performance computing.

The Storage Challenge is a competition showcasing applications and environments that effectively use the storage subsystem in high-performance computing, which is often the limiting system component. Submissions were based upon tried and true production systems as well as research or proof-of-concept projects not yet in production. Judging was based on present measurements of performance, scalability and storage subsystem utilization as well as innovation and effectiveness.

Media Contacts: Jan Zverina, SDSC Communications, 858-534-5111 or [jzverina@sdsc.edu](mailto:jzverina@sdsc.edu) Warren R. Froelich, SDSC Communications, 858-822-3622 or [froelich@sdsc.edu](mailto:froelich@sdsc.edu)

