Contribution ID: **372** Type: **oral presentation**

PetaCache: Data Access Unleashed

Monday, September 3, 2007 4:30 PM (20 minutes)

The PetaCache project started at SLAC in 2004 with support from DOE Computer Science and the SLAC HEP program. PetaCache focuses on using cost-effective solid state storage for the hottest data under analysis. We chart the evolution of metrics such as accesses per second per dollar for different storage technologies and deduce the near inevitability of a massive use of solid-state storage in the near future. We report on the latency and access-rate performance of a DRAM-based prototype constructed in 2005 using commodity hardware and a Flash-based prototype constructed in 2007 using purpose-built hardware. We describe the use of xrootd to cluster individual servers in a highly scalable and fault tolerant approach, and present tests of scalability. A study of access to ATLAS AOD data is reported as a first step in understanding the software issues that will be encountered in achieving unfettered access to objects within HEP events. Finally we examine the cost-benefit outlook for the use of solid-state storage in HEP experiments.

Primary author: Dr MOUNT, Richard (SLAC)

Co-authors: HANUSHEVSKY, Andrew (SLAC); BOEHEIM, Chuck (SLAC); Prof. LEITH, David (SLAC); BECLA, Jacek (SLAC); MELEN, Randal (SLAC); Dr PULLIAM, Teela (SLAC); WEEKS, William (SLAC)

Presenter: Dr MOUNT, Richard (SLAC)

Session Classification: Computer facilities, production grids and networking

Track Classification: Computer facilities, production grids and networking