Contribution ID: **47** Type: **poster**

Multi-Terabyte EIDE Disk Arrays running Linux RAID5

Wednesday 29 September 2004 10:00 (1 minute)

High-energy physics experiments are currently recording large amounts of data and in a few years will be recording prodigious quantities of data. New methods must be developed to handle this data and make analysis at universities possible. Grid Computing is one method; however, the data must be cached at the various Grid nodes. We examine some storage techniques that exploit recent developments in commodity hardware. Disk arrays using RAID level 5 (RAID5) include both parity and striping. The striping improves access speed. The parity protects data in the event of a single disk failure, but not in the case of multiple disk failures.
We report on tests of dual-processor Linux Software RAID5 arrays and Hardware RAID5 arrays using the 12-disk 3ware controller, in conjunction with 300 GB disks, for use in offline high-energy physics data analysis. The price of IDE disks is now less than $1/GB. These RAID5 disk arrays can be scaled to sizes affordable to small institutions and used when fast random access at low cost is important.

Authors: PETRAVICK, D. (Fermilab); SANDERS, D. (UNIVERSITY OF MISSISSIPPI); SUMMERS, D. (UNIVERSITY OF MISSISSIPPI); CREMALDI, L. (UNIVERSITY OF MISSISSIPPI); JOY, M.D. (UNIVERSITY OF MISSISSIPPI); GODANG, R. (UNIVERSITY OF MISSISSIPPI); ESCHENBURG, V. (UNIVERSITY OF MISSISSIPPI)

Presenter: SANDERS, D. (UNIVERSITY OF MISSISSIPPI)

Session Classification: Poster Session 2

Track Classification: Track 4 - Distributed Computing Services