Contribution ID: **54** Type: **not specified**

Evaluation of distributed file systems using trace and replay mechanism

Wednesday, 4 May 2011 12:15 (30 minutes)

Reliable benchmarking of file systems is a complex and time consuming task when one has to test against a production environment to achieve relevant results.
In case of the HEP community, this eventually leads to setting up a particular experiments' software environment, which could be a rather complicated task for a system administrator.
To simplify this task, we developed an application for exact replaying of IO requests to reliably replicate an IO behavior of the original applications without a need of installing the whole working environment.
Using the application, we present performance comparison of Lustre, GPFS and Hadoop file systems by replaying traces of LHCb, CMS and ATLAS jobs.

Primary author: Mr HORKY, Jiri (Institute of Physics of Acad. of Sciences of the Czech Rep. (ASCR))

Presenter: Mr HORKY, Jiri (Institute of Physics of Acad. of Sciences of the Czech Rep. (ASCR))

Session Classification: Storage & FileSystems

Track Classification: Storage & Filesystems