



Contribution ID: 502

Type: **Poster**

DPM: Future-proof storage

Tuesday, 22 May 2012 13:30 (4h 45m)

The Disk Pool Manager (DPM) is a lightweight solution for grid enabled disk storage management. Operated at more than 240 sites it has the widest distribution of all grid storage solutions in the WLCG infrastructure.

It provides an easy way to manage and configure disk pools, and exposes multiple interfaces for data access (rpio, xroot, nfs, gridftp and http/dav) and control (srm). During the last year we have been working on providing stable, high performant data access to our storage system using standard protocols, while extending the storage management functionality and adapting both configuration and deployment procedures to reuse commonly used building blocks.

In this contribution we cover in detail the extensive evaluation we have performed of our new HTTP/WebDAV and NFS 4.1 frontends, in terms of functionality and performance. We summarize the issues we faced and the solutions we developed to turn them into valid alternatives to the existing grid protocols - namely the additional work required to provide multi-stream transfers for high performance wide area access, support for third party copies, credential delegation or the required changes in the experiment and fabric management frameworks and tools.

We describe new functionality that has been added to ease system administration, such as different filesystem weights and a faster disk drain, and new configuration and monitoring solutions based on the industry standards Puppet and Nagios. Finally, we explain some of the internal changes we had to do in the DPM architecture to better handle the additional load from the analysis use cases.

Primary author: BRITO DA ROCHA, Ricardo (CERN)

Co-authors: ALVAREZ AYLON, Alejandro (University of Cadiz); BECHE, Alexandre (SUPINFO International University (FR)); FURANO, Fabrizio (CERN); BAUD, Jean-Philippe (CERN); KEEBLE, Oliver (CERN)

Presenter: BRITO DA ROCHA, Ricardo (CERN)

Session Classification: Poster Session

Track Classification: Distributed Processing and Analysis on Grids and Clouds (track 3)