



Contribution ID: 210

Type: poster

## Optimized tertiary storage access in the dCache SE

*Wednesday, September 5, 2007 8:00 AM (20 minutes)*

The dCache software has become a major storage element in the WLCG, providing high-speed file transfers by caching datasets on potentially thousands of disk servers in front of tertiary storage. Currently dCache's model of separately connecting all disk servers to the tape backend leads to locally controlled flush and restore behavior has shown some inefficiencies in respect of tape drive utilization. This paper reports on enhancements to dCache and its interface to Hierarchical Storage Management Systems through a new module to centrally control the tape backend. The focus is to optimize the mount behavior, parallelization of requests and utilization of the available tape drive data rate. dCache is widely deployed in production instances, so improvements will be phased to maintain backwards compatibility. The first phase has been achieved by introducing the Flushmanager, a module to collect and coordinate all tape write operations. The second phase of implementation will be about collecting metadata and restore requests respectively. This information will be provided to user-added plug-ins via a clean interface, for managing restore requests using customized strategies. The third phase will focus on resource knowledge and drive steering with a generic and feature-rich interface to unify access to common Tertiary Storage Systems. The final phase will be the aggregation of the existing Flush- and Restore-Manager to implement a single module for controlling tape backends through dCache, offering a extensible programming framework to the community.

**Primary author:** Mr RADICKE, Martin (DESY Hamburg)

**Co-author:** Dr FUHRMANN, Patrick (DESY Hamburg)

**Presenter:** Mr RADICKE, Martin (DESY Hamburg)

**Session Classification:** Poster 2

**Track Classification:** Grid middleware and tools