



Contribution ID: 253

Type: **Parallel**

## dCache, agile adoption of storage technology

*Thursday, May 24, 2012 4:35 PM (25 minutes)*

For over a decade, dCache has been synonymous with large-capacity, fault-tolerant storage using commodity hardware that supports seamless data migration to and from tape. Over that time, it has satisfied the requirements of various demanding scientific user communities to store their data, transfer it between sites and fast, site-local access.

When the dCache project started, the focus was on managing a relatively small disk cache in front of large tape archives. Over the project's lifetime storage technology has changed. During this period, technology changes have driven down the cost-per-GiB of harddisks. This resulted in a shift towards systems where the majority of data is stored on disk. More recently, the availability of Solid State Disks, while not yet a replacement for magnetic disks, offers an intriguing opportunity for significant performance improvement if they can be used intelligently within an existing system.

New technologies provide new opportunities and dCache user communities' computing models are changing. The traditional data models, in which tape is used as an active storage, are being revised with tape adopting a more archival model. The symbiotic relationship between dCache and the end-users means that dCache is both driven by and facilitating these changes.

Recently, dCache introduced support for WebDAV and the NFS 4.1/pNFS protocols. This move away from bespoke protocols towards standards is the result of the availability of protocols that support large storage systems. dCache's adoption of standards allows end-users to use their favourite desktop data-transfer clients or unmodified analysis software. This keeps dCache competitive with industry solutions.

Hadoop FS (HDFS) provides an easy-to-maintain backend storage that is showing promise as an easy-to-maintain storage system. dCache is adopting HDFS as an alternative to local filesystem storage. Since HDFS doesn't offer file system semantics, integrating support into dCache provides some challenges. Once solved, this work will allow dCache integration with other storage technologies such as object stores and cloud storage.

We present a short summary of what dCache is providing in new long-term support release (the next "Golden Release") and offers a glimpse into the future of dCache with the emerging storage technology.

**Primary authors:** Dr FUHRMANN, Patrick (DESY); MILLAR, Paul (Deutsches Elektronen-Synchrotron (DE))

**Presenter:** MILLAR, Paul (Deutsches Elektronen-Synchrotron (DE))

**Session Classification:** Distributed Processing and Analysis on Grids and Clouds

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