EGI Community Forum 2012



Contribution ID: 159

Type: not specified

dCache, agile adoption of storage technology

Tuesday, 27 March 2012 16:30 (30 minutes)

Description of the Work

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 such protocols supporting high volumes of data. The adoption of standards allows end-users to use their favourite desktop data-transfer clients or unmodified analysis software, keeping dCache competitive with industry solutions.

Hadoop FS (HDFS) provides an easy-to-maintain back-end storage that is showing promise as an easy-tomaintain 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.

Conclusions

This presentation will give 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.

Impact

The adoption of widely used industry standards, the integration of well accepted software systems (e.g. the Hadoop FS) as well as operating new hardware technologies (e.g. Solid State Disks) in an optimal fashion with respect to regular disks and tape, provides dCache customers with an open source storage technology which easily competes with expensive storage systems on the market.

Overview (For the conference guide)

For over a decade, dCache is synonymous with large-capacity, fault-tolerant storage using commodity hardware that supports seamless data migration to and

from tape. It has satisfied the requirements of various demanding scientific user communities to store their data, transfer it between sites and fast, site-local access.

Over time, technology changes. 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, technology changes have driven

down the cost-per-GiB of harddisks. This triggered 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.

Primary author: Dr MILLAR, Paul (DESY/dCache.org)Presenter: Dr MILLAR, Paul (DESY/dCache.org)Session Classification: Data Technologies