



Contribution ID: 376

Type: **Oral presentation to parallel session**

dCache: Big Data storage for HEP communities and beyond

Tuesday, 15 October 2013 13:53 (20 minutes)

Storage is a continually evolving environment, with new solutions to both existing problems and new challenges. With over ten years in production use, dCache is also evolving to match this changing landscape. In this paper, we present three areas in which dCache is matching demand and driving innovation.

Providing efficient access to data that maximises both streaming and random-access workloads has always been a challenge. Various proprietary protocols, such as dcap and xrootd, have been developed to meet this challenge. Maintaining a property protocol incurs costs as the client must be written and supported on different platforms. With the arrival of NFS v4.1/pNFS, there is an industry standard that maximises throughput, so the need for property protocols has gone. Realising this, the dCache team has been involved with NFS v4.1 standard since its initial draft stage and all available dCache versions provide NFS support. We present the results of running NFS in production environment.

Managing storage is an evolving field. Formally, the SRM protocol was key; however, more recently HEP communities have shunned SRM in favour of other protocols, such as WebDAV. In a related move, various groups have investigated using cloud infrastructures, either wholly or using a cloud-bursting model to satisfy surges in demand. While continuing to support WebDAV, dCache is introducing support for CDMI, the ISO standard protocol for managing cloud storage. In common with WebDAV, CDMI is based on HTTP but provides better support for common management operations. In this paper, we present a comparison between WebDAV and CDMI, the current status of support in dCache, early results from experiments and future plans for this protocol.

Finally, one constant complaint from people using the grid is related to X.509 certificates. Various solutions have been presented, including EMI's Security Token Service to allow portal-like access to grid resources and to simplify the "grid login" step. Within Germany, the LSDMA project is investigating a more radical solution. In addition to the established grid methods, storage services may be accessed directly using an identity asserted by the user's home institute. This would allow scientists to use dCache securely by logging in with the user-name and password from their home university or research institute. Details of this development work are presented along with possible deployment scenarios.

Primary authors: Dr ROSSI, Albert (FNAL); Dr PETERSEN, Antje (DESY); BERNARDT, Christian (Deutsches Elektronen-Synchrotron (DE)); Dr LITVINTSEV, Dmitry (FNAL); SCHWANK, Karsten; Dr FUHRMANN, Patrick (DESY); Dr MILLAR, Paul (Deutsches Elektronen-Synchrotron (DE)); MKRTCHYAN, Tigran Mkrtchyan (Deutsches Elektronen-Synchrotron DESY)

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

Session Classification: Data Stores, Data Bases, and Storage Systems

Track Classification: Data Stores, Data Bases, and Storage Systems