



Contribution ID: 55

Type: **Oral Presentation**

EMI and dCache: the vanguard of standards adoption.

Wednesday, 13 April 2011 14:00 (30 minutes)

Overview

dCache is a hugely successful storage software that is currently deployed in 8 (out of 11) Tier-1 centres in WLCG and providing over 40% of the total storage capacity available to the LHC experiments. With dCache as a founding member, the EMI project has great experience in storage.

The HEP community has a tradition of building their own protocols for LAN and WAN access to data. This is because the existing protocols have proven unable to sustain the required throughput.

Unfortunately, these non-standard protocols introduces a major barrier to new communities adopting grid computing. These communities may have custom software that would need to made “grid-aware”; the end-users may be using commercial off-the-shelf analysis software that cannot be modify.

Recent developments in industry mean that a large storage system, such as dCache, must adopt industry standards to remain competitive with solutions from HPSS, NetApp, EMC², Panasas, Dell-DDN, etc, without sacrificing throughput

Impact

dCache’s adoption of WebDAV means that end-users can access their data on their choice of platform, with their choice of access tool. There are WebDAV clients for all major computing platforms, which often supplied as part of the operating system.

Several sites are deploying dCache with WebDAV specifically to allow end-users outside the HEP community to have access to large data storage; for example, dCache and WebDAV forms the basis of the Swedish National Infrastructure for Computing (SNIC) national storage facility: SweStore. SweStore aims to support non-HEP communities such as climate modelling, bioinformatics and bioimaging.

With the availability of native support for NFS v4.1 in the various operating systems, a site-administrator can simply mount the large data stored in dCache on local analysis machines. Programs running on those machines may access data like any other file, without modification. This permits dCache to support analysis work-flows that uses software that cannot be modified to support HEP-specific protocols, which permits analysis facilities and compute farms to support a much wider range of analysis activity.

Adopting NFS v4.1 also brings advantages to site administrators. Adoption of standard protocols allows a mix-and-match approach, where storage from different vendors may be combined, based on their relative benefits, to produce a coherent storage solution. The standard also provides better toolage for monitoring.

Description of the work

The EMI project includes a programme where the storage elements are adopting industry standards. Within this programme dCache has adopted a trailblazer role: demonstrating the feasibility of adopting these industry standard protocols and demonstrating their functionality in production environments.

WebDAV is a industry standard protocol that allows end-users to read and write data. dCache has been adopting WebDAV, allow end-users to read and write data over the WAN. Within dCache, access via WebDAV is no different from other protocols: end-users experience a coherent namespace and are subject to the same set of access controls.

For local access dCache has implemented the NFS v4.1 protocol with pNFS, one of the first storage solution to do so. Through the use of realistic tests, dCache has demonstrated deploying NFS v4.1 for the high-throughput analysis work typical of HEP communities. There are sites in the process of rolling out NFS v4.1 support.

Conclusions

By adopting standard protocols, such as WebDAV and NFS v4.1, dCache is allowing end-users to be agnostic to the nature of the storage. This lowers the barrier for new communities that need access to large storage, allowing them to use grid resources without modifying their software. This work is already bringing benefits to end-users.

Primary authors: PETERSEN, Antje (DESY); BERNARDT, Christian (DESY); LITVINTSEV, Dmitry (Fermi-Lab); BEHRMANN, Gerd (NDGF); SCHWANK, Karsten (DESY); FUHRMANN, Patrick (DESY); MILLAR, Paul (DESY); LUNDQVIST, Per (National Supercomputer Centre in Sweden); BARANOVA, Tanja (DESY); MKRTCHYAN, Tigran (DESY)

Presenter: MILLAR, Paul (DESY)

Session Classification: EMI: Software for Distributed Computing Infrastructures

Track Classification: DCI - Standards