

Scalable sync-and-share service with dCache

Monday 30 January 2017 16:30 (20 minutes)

Scientific exploration and exploitation of data is undergoing a revolution as communities explore new ways of analysing their data. One solution that is being used increasingly is sync-and-share, where data, presentations, graphs and code are shared in an ad hoc fashion. This allows communities to explore data in new and innovative ways.

Sites that have already invested in dCache to solve their large data storage requirements are keen that their services integrate seamlessly with sync-and-share systems. Such combined systems should function without any degradation of either system.

The combined system, involving ownCloud and dCache, is also attractive for sites that have not yet invested in dCache as they will want a solution that scales as well as dCache.

Naturally, dCache has many years of experience handling multiple petabytes of data: we know how to scale such data services. Although large data is problematic for (own|next)Cloud, the ownCloud+dCache deployment at DESY ("DESY Cloud") shows such a deployment is practical.

Although a simple deployment functions well enough, we see opportunities to improve the performance; in particular, there is the possibility to improve the interface between (own|next)Cloud and dCache; to avoid duplication of information and avoid potential bottlenecks.

We will present a summary of the ownCloud+dCache hybrid system, identifying problem areas and present our solutions to solve these issues.

Authors: FUHRMANN, Patrick (DESY); MILLAR, Paul (DESY); Mr MKRTCHYAN, Tigran (DESY)

Presenter: MILLAR, Paul (DESY)

Session Classification: Technology