



Contribution ID: 444

Type: **Poster presentation**

Performance evaluation of a dCache system with pools distributed over a wide area

Monday, 14 October 2013 15:00 (45 minutes)

Distributed storage systems have evolved from providing a simple means to store data remotely to offering advanced services like system federation and replica management. This evolution has been made possible due to the advancement of the underlying communication technology, that plays a vital role in determining the communication efficiency of the distributed systems. The dCache system, which has a wide installation base in the High Energy Physics community, is a distributed data caching environment built over a large number of heterogeneous storage servers usually located at a single physical site that is able to offer to the end user a unified filesystem view of the entire data repository by means of several protocol standards. One of the key features of the dCache system is the dissociation of the file namespace of the repository from the actual location of the data files in the storage servers. This allows files to reside anywhere in the distributed system, facilitating data replication and improving data access. Another interesting feature is load balancing by means of hot-spot detection, live migration, and caching mechanisms. The latest versions of dCache are capable of providing access to data through the new Network File System version 4.1 (NFSv4.1) protocol, which extends the capabilities of its predecessors by supporting parallel I/O capabilities, thus increasing scalability and overall performance.

In this work we propose to evaluate the performance of the new dCache system implemented with NFSv4.1, with the goal of analyzing its behavior when combining distributed server pools spread over a wide area and interconnected by a dedicated optical network system. To establish reference points we will begin by analyzing dCache deployed as a central storage system, providing data to local clients through NFSv4.1. A simple local dCache system will be built, in which the server pools are basically low cost JBODs. The centralization of different storage systems into a single structure allows the reduction in the number of storage systems deployed in an enterprise datacenter, which consequently leads to cost reduction. Another interesting advantage is the decrease of the setup time for new storage servers, which basically consists of adding more JBODs to the already deployed centralized storage system. The dCache I/O performance will then be tested with Iozone, a commonly used benchmarking tool, and the behavior of the system will be evaluated in response to increasing loads. During the tests we will evaluate system scalability (by adding new pools) and fault tolerance (by detaching some of the pool servers).

Primary authors: BACH, Eduardo (UNESP - Universidade Estadual Paulista (BR)); IOPE, Rogerio (UNESP - Universidade Estadual Paulista (BR))

Presenter: BACH, Eduardo (UNESP - Universidade Estadual Paulista (BR))

Session Classification: Poster presentations

Track Classification: Facilities, Production Infrastructures, Networking and Collaborative Tools