

Increasing Performance and Scalability of dCache Storage Resource Manager

Thursday, 26 March 2009 08:00 (20 minutes)

The dCache disk caching file system has been chosen by a majority of LHC Experiments' Tier 1 centers for their data storage needs. It is also deployed at many Tier 2 centers. In preparation for the LHC startup, very large installations of dCache - up to 3 Petabytes of disk - have already been deployed, and the systems have operated at transfer rates exceeding 2000 MB/s over the WAN. As the LHC experiments go into production, it is expected that data storage capacity requirements and data transfer rates will continue to grow beyond currently tested limits. It is estimated that Tier-1 center serving just CMS experiment needs to support a sustained data throughput of 800 MB/sec.

As any other software being in production for years, dCache faced a change in access profile and required performance. In order to cope with evolving requirements and with access patterns requiring better performance, the dCache team regularly investigate improving components which might no longer be state of the art. As we did with other dCache components, we are now evaluating a possible redesign of The Storage Resource Manager (SRM) for scalability. SRM is a main Grid Storage Interface and a single point of entry into dCache, is one of the most critical components. SRM needs to be able to scale with increased load and to remain resilient against changing usage patterns.

We will present an analysis of the dCache architecture its performance bottlenecks with emphasis on the SRM and the current and future effort to improve scalability and stability in order to satisfy the ever increasing LHC experiments' requirements.

Author: PERELMUTOV, Timur (FERMI NATIONAL ACCELERATOR LABORATORY)

Presenter: PERELMUTOV, Timur (FERMI NATIONAL ACCELERATOR LABORATORY)

Session Classification: Poster session

Track Classification: Grid Middleware and Networking Technologies