



Contribution ID: 233

Type: poster

Production Experience Using Resilient dCache

Wednesday, 5 September 2007 08:00 (20 minutes)

dCache is a distributed storage system which today stores and serves petabytes of data in several large HEP experiments. Resilient dCache is a top level service within dCache, created to address reliability and file availability issues when storing data for extended periods of time on disk-only storage systems. The Resilience Manager automatically keeps the number of copies within specified bounds by adjusting the number of replicas of each logical file on different units of disk hardware when files disk pool nodes are found to have crashed, been removed from, or added to the system.

We presented design of the dCache Resilience Manager in the CHEP2006 report "Resilient dCache: Replicating Files for Integrity and Availability". The present paper provides an update on further development of Resilient Manager and experience in the production deployment and operations in US-CMS T1 and T2 centers. The US-CMS T1 center substantially increased the size of their Resilient dCache and added second group of resilient pools for merging short files with production job output before storing files on tape. Two resilient pool groups operate independently of each other and other pool groups (tape-backed or volatile). A few more US-CMS T2 centers started to use Resilient Manager to increase the integrity and size of their systems. Based on experience with the Resilient Manager in US-CMS centers we added new features to drain files from the pools for hardware retirement and to avoid replication of files to the same pool host, while improving the Resilient Manager's performance and manageability.

Submitted on behalf of Collaboration (ex, BaBar, ATLAS)

dCache

Primary author: Mr KULYAVTSEV, Alexander (FNAL)

Presenter: Mr KULYAVTSEV, Alexander (FNAL)

Session Classification: Poster 2

Track Classification: Distributed data analysis and information management