



Contribution ID: 230

Type: oral presentation

CASTOR: Operational issues and new Developments

Wednesday 29 September 2004 16:30 (20 minutes)

The Cern Advanced STORage (CASTOR) system is a scalable high throughput hierarchical storage system developed at CERN. CASTOR was first deployed for full production use in 2001 and has expanded to now manage around two PetaBytes and almost 20 million files. CASTOR is a modular system, providing a distributed disk cache, a stager, and a back end tape archive, accessible via a global logical name-space.

This paper focuses on the operational issues of the system currently in production, and first experiences with the new CASTOR stager which has undergone a significant redesign in order to cope with the data handling challenges posed by the LHC, which will be commissioned in 2007.

The design target for the new stager was to scale to another order of magnitude above the current CASTOR, namely to be able to sustain peak rates of the order of 1000 file open requests per second for a PetaByte disk pool. The new developments have been inspired by the problems which arose managing massive installations of commodity storage hardware. The farming of disk servers poses new challenges to the disk cache management: request scheduling; resource sharing and partitioning; automated configuration and monitoring; and fault tolerance of unreliable hardware

Management of the distributed component based CASTOR system across a large farm, provides an ideal example of the driving forces for the development of automated management suites. Quattor and Lemon frameworks naturally address CASTOR's operational requirements, and we will conclude by describing their deployment on the masstorage systems at CERN.

Authors: COUTURIER, B. (CERN); DURAND, J-D. (CERN); BARRING, O. (CERN); PONCE, S. (CERN)

Presenter: DURAND, J-D. (CERN)

Session Classification: Computer Fabrics

Track Classification: Track 6 - Computer Fabrics