CASTOR Tape Performance Optimisation at the UK LCG Tier-1

Tuesday 24 March 2009 08:00 (20 minutes)

The UK LCG Tier-1 computing centre located at the Rutherford Appleton Laboratory is responsible for the custodial storage and processing of the raw data from all four LHC experiments; CMS, ATLAS, LHCb and ALICE. The demands of data import, processing, export and custodial tape archival place unique requirements on the mass storage system used. The UK Tier-1 uses CASTOR as the storage technology of choice, which currently handles 2.3PB of disk across 320 disk servers. 18 Sun T10000 tape drives provide the custodial backend. This paper describes work undertaken to optimise the performance of the CASTOR infrastructure at RAL. Significant gains were achieved and the lessons learned have been deployed at other LHC CASTOR sites.

Problems were identified with the performance of tape migration when disk servers were under productionlevel load. An investigation was launched at two levels; hardware and operating system performance, and the impact of CASTOR tape algorithms and job scheduling. A test suite was written to quantify the low-level performance of disk servers with various tunings applied, and CMS test data coupled with the existing transfer infrastructure was used to verify the performance of the tape system with realistic experimental data transfer patterns. The improvements identified resulted in the instantaneous tape migration rate per drive reaching near line-speed of 100MB/s, a vast improvement on the previous attainable rate of around 16MB/s.

Presentation type (oral | poster)

oral

Summary

Performance optimisations of the CASTOR storage system used at the UK LCG Tier-1 at RAL were performed. Significant gains in the tape migration rate were achieved through hardware and operating system level optimisations, and by the development of new algorithms and job scheduling policies within CASTOR.

Authors: Dr SANSUM, Andrew (Rutherford Appleton Laboratory); Dr NEWBOLD, Dave (H.H. Wills Physics Laboratory - University of Bristol); Mr JACKSON, James (H.H. Wills Physics Laboratory - University of Bristol); Mr PONCE, Sebastien (CERN)

Presenter: Mr JACKSON, James (H.H. Wills Physics Laboratory - University of Bristol)

Session Classification: Poster session

Track Classification: Hardware and Computing Fabrics