



Contribution ID: 415

Type: **Poster**

Tape write efficiency improvements in CASTOR

Tuesday 22 May 2012 13:30 (4h 45m)

The CERN Advanced STORage manager (CASTOR) is used to archive to tape the physics data of past and present physics experiments. Data is migrated (repacked) from older, lower density tapes to newer, high-density tapes approximately every two years to follow the evolution of tape technologies and to keep the volume occupied by the tape cartridges relatively stable. Improving the performance of writing files smaller than 2G to tape is essential in order to keep the time needed to repack all of the tape resident data within a period of no more than 1 year. Until now CASTOR has flushed the write buffers of the underlying tape-system 3 times per user-file, using up to 7 seconds. With current drive-writing speeds reaching over 240MB/s per second, 7 seconds of flush-time equates to an approximate loss of 1.5 GB of data transfer time per user-file. This paper reports on the solution to writing efficiently to tape that is currently in its early deployment phases at CERN. Write speeds have been increased whilst preserving the existing tape-format by using immediate (non-flushing) tape-marks to write multiple user-files before flushing the tape-system write-buffers. The solution has been realized as a set of incremental upgrades to minimize risk, maximize backwards compatibility and work safely with the legacy modules of CASTOR. Unit testing has been used to help reduce the risk of working with legacy code. This solution will enable CASTOR to continue to be a long-term and performant tool for archiving past and present experiment data to tape.

Summary

The CERN Advanced STORage manager (CASTOR) is used to archive to tape the physics data of past and present physics experiments. For reasons of physical storage space, all of the tape resident data in CASTOR are repacked onto higher density tapes approximately every two years. Improving the performance of writing files smaller than 2G to tape is essential in order to keep the time needed to repack all of the tape resident data within a period of no more than 1 year. This paper reports on the solution to writing efficiently to tape that is currently in its early deployment phases at CERN.

Primary author: MURRAY, Steven (CERN)

Co-authors: CANO, Eric (CERN); Mr CANCIO, German (CERN); Dr LO PRESTI, Giuseppe (CERN); LO RE, Giuseppe (CERN); PONCE, Sebastien (CERN); KOTLYAR, Victor (Institute for High Energy Physics (RU)); BAHYL, Vlado (CERN)

Presenter: MURRAY, Steven (CERN)

Session Classification: Poster Session

Track Classification: Computer Facilities, Production Grids and Networking (track 4)