



Contribution ID: 64

Type: poster presentation

The new CERN tape software - getting ready for total performance

CASTOR (the CERN Advanced STORAGE system) is used to store the custodial copy of all of the physics data collected from the CERN experiments, both past and present. CASTOR is a hierarchical storage management system that has a disk-based front-end and a tape-based back-end. The software responsible for controlling the tape back-end has been redesigned and redeveloped over the last year and shall be put into production at the beginning of 2015. This paper summarises the motives behind the redesign, describes in detail the redevelopment work and concludes with the short and long-term benefits.

Modern tape drives achieve 250 MB/s and speeds up to 1GB/s are on the roadmaps. To achieve this performance and to drop support for obsolete requirements, a new tape server software has been designed from the ground up, with a fully pipelined architecture. Disk and tape transfers, instruction fetches and result reporting are all running in independent threads communicating through queues to prevent external latencies from impacting tape drive performance.

The software has been developed in C++, facilitating modularity, unit testing, and faster development turnaround to face future challenges. The main developments included a novel methodology for SCSI development, polymorphic drive classes and factories to minimize extra coding, a tape file layer re-implementing the existing file format, polymorphic disk file access methods for easy extension to new protocols, thread-safe FIFO containers, easy to use threading building blocks and a memory management system for buffering data between disk and tape transfers.

Development was backed by a continuous integration system that included unit testing. Thanks to the use of mock objects, such as simulated tape drives, the unit tests covered complex use cases comparable with full production sessions. The unit tests were also validated by memory leak and race condition detectors.

Thanks to strict adherence to unit testing methodologies, the project went very smoothly. The development started on the 31st May 2013. The first read-only test version was delivered to tape operators mid-August 2014 for running against production stagers and the read-write version followed in mid-September. The project took the equivalent 2 FTEs for 5 people involved.

Starting from this solid base, the CERN tape infrastructure will evolve to address future needs such as long term data preservation, striped file systems access for higher throughput and session pre-emption to improve the granularity of drive allocations and to ensure for the full utilization of all drives at all times.

Primary authors: CANO, Eric (CERN); MURRAY, Steven (CERN)

Co-authors: KRUSE, Daniele Francesco (CERN); COME, David (CERN/ISAE-Supaéro); KOTLIAR, Viktor (Institute for High Energy Physics (RU))

Presenter: CANO, Eric (CERN)

Track Classification: Track3: Data store and access