

A High Performance Hierarchical Storage Management System For the Canadian Tier-1 Centre at TRIUMF

Monday 23 March 2009 17:30 (20 minutes)

We describe in this paper the design and implementation of TapeGuy, a high performance non-proprietary Hierarchical Storage Management System (HSM) which is interfaced to dCache for efficient tertiary storage operations. The system has been successfully implemented at the Canadian Tier-1 Centre at TRIUMF. The ATLAS experiment will collect a very large amount of data (approximately 3.5 Petabytes each year). An efficient HSM system will play a crucial role in the success of the ATLAS Computing Model which is driven by intensive large-scale data analysis activities that will be performed on the Worldwide LHC Computing Grid infrastructure around the clock.

TapeGuy is perl-based. It controls and manages data and tape libraries. Its architecture is scalable and includes Dataset Writing control, a Readback Queuing mechanism and I/O tape drive load balancing as well as on-demand allocation of resources. A central MySQL database records metadata information for every file and transaction (for audit and performance evaluation), as well as an inventory of library elements. TapeGuy Dataset Writing was implemented to group files which are close in time and of similar type. Optional dataset path control dynamically allocates tape families and assigns tapes to it. Tape flushing is based on various strategies: time, threshold or external callbacks mechanisms. TapeGuy Readback Queuing reorders all read requests by using a 'scan algorithm', avoiding unnecessary tape loading and unloading. Implementation of priorities will guarantee file delivery to all clients in a timely manner.

Presentation type (oral | poster)

oral

Authors: Ms DEATRICH, Denice (TRIUMF); Dr TAFIROUT, Reda (TRIUMF); Mr LIU, Simon (TRIUMF)

Presenter: Mr LIU, Simon (TRIUMF)

Session Classification: Hardware and Computing Fabrics

Track Classification: Hardware and Computing Fabrics