Contribution ID: **287** Type: **Poster**

Multi-threaded checksum computation for the ATLAS high-performance storage software

ATLAS is one of the generic-purpose experiments observing hadron collisions at the LHC at CERN. Its trigger and data acquisition system (TDAQ) is responsible for selecting and transporting interesting physics events from the detector to permanent storage where the data are used for physics analysis. The transient storage of ATLAS TDAQ is the last component of the online data-flow system. It records selected events at several GB/s to non-volatile storage before transfer to the offline facilities. The transient storage is a distributed system consisting of high-performance direct-attached storage servers accounting for 480 hard drives. A distributed multi-threaded C++ application operates the hardware. Reliability and efficiency of this system are critical for the operations of ATLAS and TDAQ.
As part of the transient storage workflow, checksums of the recorded data files are calculated. The checksums are used throughout the offline data management and data distribution system to guarantee the integrity and correctness of the raw data. This paper presents the current multi-threading strategy of the transient storage software and the associated trade off between compute and storage hardware resources. We then introduce a novel multi-threaded checksum computation strategy. We discuss the key concepts of the implementation with a focus on the importance of overhead minimization. Finally the paper reports on the tests performed on the production system to demonstrate the validity of the implementation. A 30 % increase in the overall throughput performance will be demonstrated and discussed in the view of future LHC and ATLAS upgrades.

Authors: LE GOFF, Fabrice (CERN); AVOLIO, Giuseppe (CERN)

Presenter: AVOLIO, Giuseppe (CERN)

Session Classification: Poster Session

Track Classification: Track 1: Computing Technology for Physics Research