

Automated load balancing in the ATLAS high-performance storage software

Thursday, May 25, 2017 4:48 PM (18 minutes)

The ATLAS experiment collects proton-proton collision events delivered by the LHC accelerator at CERN. The ATLAS Trigger and Data Acquisition (TDAQ) system selects, transports and eventually records event data from the detector at several gigabytes per second. The data are recorded on transient storage before being delivered to permanent storage. The transient storage consists of high-performance direct-attached storage servers accounting for about 500 hard drives. The transient storage operates dedicated software in the form of a distributed multi-threaded application. The workload includes both CPU-demanding and IO-oriented tasks. This paper presents the original application threading model for this particular workload, discussing the load-sharing strategy among the available CPU cores. The limitations of this strategy were reached in 2016 due to changes in the trigger configuration involving a new data distribution pattern. We then describe a novel data-driven load-sharing strategy, designed to automatically adapt to evolving operational conditions, as driven by the detector configuration or the physics research goals. The improved efficiency and adaptability of the solution were measured with dedicated studies on both test and production systems. This paper reports on the results of those tests which demonstrate the capability of operating in a large variety of conditions with minimal user intervention.

Summary

See attached file

Primary author: FABRICE, Le Goff (Rutherford Appleton Laboratory)

Presenter: FABRICE, Le Goff (Rutherford Appleton Laboratory)

Session Classification: R3-Trigger and data acquisition systems(5)

Track Classification: Trigger and data acquisition systems