



Contribution ID: 424

Type: **Poster**

Simulating the ATLAS Distributed Data Management System

Thursday 24 May 2012 13:30 (4h 45m)

The ATLAS Distributed Data Management system stores more than 75PB of physics data across 100 sites globally. Over 8 million files are transferred daily with strongly varying usage patterns. For performance and scalability reasons it is imperative to adapt and improve the data management system continuously. Therefore future system modifications in hardware, software as well as policy, need to be evaluated to accomplish good results and avoid unwanted side effects. Due to the complexity of large-scale distributed systems this evaluation process is primarily based on expert-knowledge, as conventional evaluation methods are inadequate. However, this error-prone process lacks quantitative estimations and leads to inaccuracy as well as incorrect evaluations.

In this work we present a novel, full-scale simulation framework. This flow-level based simulator is able to accurately model the ATLAS Distributed Data Management system. The design and architecture of the component-based software is presented and discussed. The evaluation concentrates on the accuracy and scalability of the simulation framework. Finally, selected use-cases where simulation could be hugely beneficial to distributed data management systems are presented and discussed.

Author: ATLAS, Collaboration (Atlas)

Co-authors: MOLFETAS, Angelos (CERN); LASSNIG, Mario (CERN); BARISITS, Martin (Vienna University of Technology (AT)); GARONNE, Vincent (CERN)

Presenter: BARISITS, Martin (Vienna University of Technology (AT))

Session Classification: Poster Session

Track Classification: Software Engineering, Data Stores and Databases (track 5)