

Kudu prototypes for the ATLAS EventIndex and the future Event WhiteBoard

Tuesday, July 10, 2018 4:45 PM (15 minutes)

The ATLAS EventIndex has been in operation since the beginning of LHC Run 2 in 2015. Like all software projects, its components have been constantly evolving and improving in performance. The main data store in Hadoop, based on MapFiles and HBase, can work for the rest of Run 2 but new solutions are explored for the future. Kudu offers an interesting environment, with a mixture of BigData and relational database features, which look promising at the design level and is used to build a prototype to measure the scaling capabilities as functions of data input rates, total data volumes and data query and retrieval rates.

The Event WhiteBoard will be an extension of the EventIndex functionalities to support the ATLAS Event Service and the concept of Virtual Datasets. These additional requirements are tested on the same Kudu prototype, in order to estimate the system performance and response times for different internal data organisations.

This talk will report on the selected data schemas and on the current performance measurements with the Kudu prototype.

Primary authors: BARBERIS, Dario (Università e INFN Genova (IT)); BARANOWSKI, Zbigniew (CERN); CANALI, Luca (CERN); FERNANDEZ CASANI, Alvaro (Univ. of Valencia and CSIC (ES)); GALLAS, Elizabeth (University of Oxford (GB)); GARCÍA MONTORO, Carlos (IFIC); GONZALEZ DE LA HOZ, Santiago (Univ. of Valencia and CSIC (ES)); HRIVNAC, Julius (Centre National de la Recherche Scientifique (FR)); PROKOSHIN, Fedor (Federico Santa Maria Technical University (CL)); RYBKIN, Grigori (Centre National de la Recherche Scientifique (FR)); SALT, Jose (Instituto de Fisica Corpuscular (IFIC) - Universidad de Valencia); SANCHEZ, Javier (Univ. of Valencia and CSIC (ES)); VILLAPLANA PEREZ, Miguel (IFIC)

Presenter: BARANOWSKI, Zbigniew (CERN)

Session Classification: Posters

Track Classification: Track 4 - Data Handling