



Contribution ID: 288

Type: Poster

Evaluating InfluxDB and ClickHouse database technologies for improvements of the ATLAS operational monitoring data archiving

The Trigger and Data Acquisition (TDAQ) system of the ATLAS experiment the Large Hadron Collider (LHC) at CERN currently is composed of a large number of distributed hardware and software components (about 3000 machines and more than 25000 applications) which, in a coordinated manner, provide the data-taking functionality of the overall system.

During data taking runs, a huge flow of operational data is produced in order to constantly monitor the system and allow proper detection of anomalies or misbehaviors. The Persistent Back-End for the ATLAS Information System of TDAQ (P-BEAST) is a system based on a custom-built time-series database and it is used to archive and retrieve for applications any operational monitoring data. P-BEAST stores about 18 TB of highly compacted and compressed raw monitoring data per year acquired at 200 KHz average information update rate during ATLAS data taking periods.

Since P-BEAST has been put into production, 4 years ago, several promising database technologies for fast access to time-series and column-oriented data have become available. InfluxDB and ClickHouse were the most promising candidates for improving the performance and functionality of the current implementation of P-BEAST.

This paper presents a short description of main features of both technologies and a description of the synthetic tests ran on both database systems which try to leverage the best possible options for storage of the P-BEAST data using their respective data model capabilities. Then, the results of the performance testing that has been performed using a subset of archived ATLAS operational monitoring data are presented. Finally, a comparison of the results is presented.

Primary authors: VASILE, Matei (IFIN-HH (RO)); SOLOVIEV, Igor (University of California Irvine (US)); AVOLIO, Giuseppe (CERN)

Presenter: VASILE, Matei (IFIN-HH (RO))

Session Classification: Poster Session

Track Classification: Track 1: Computing Technology for Physics Research