

An Oracle-based Event Index for ATLAS

Monday 10 October 2016 15:15 (15 minutes)

The ATLAS EventIndex System has amassed a set of key quantities for a large number of ATLAS events into a Hadoop based infrastructure for the purpose of providing the experiment with a number of event-wise services. Collecting this data in one place provides the opportunity to investigate various storage formats and technologies and assess which best serve the various use cases as well as consider what other benefits alternative storage systems provide.

In this presentation we describe how the data are imported into an Oracle RDBMS, the services we have built based on this architecture, and our experience with it. We've indexed about 26 billion real data events thus far and have designed the system to accommodate future data which has expected rates of 5 and 20 billion events per year. We have found this system offers outstanding performance for some fundamental use cases. In addition, profiting from the co-location of this data with other complementary metadata in ATLAS, the system has been easily extended to perform essential assessments of data integrity and completeness and to identify event duplication, including at what step in processing the duplication occurred.

Secondary Keyword (Optional)

Storage systems

Primary Keyword (Mandatory)

Databases

Tertiary Keyword (Optional)

Primary author: GALLAS, Elizabeth (University of Oxford (GB))

Co-authors: FORMICA, Andrea (CEA/IRFU,Centre d'etude de Saclay Gif-sur-Yvette (FR)); DUMITRU, Andrei (CERN); DIMITROV, Gancho (CERN); CANALI, Luca (CERN); PETROVA, Petya Tsvetanova (University of Texas at Arlington (US)); BARANOWSKI, Zbigniew (CERN)

Presenter: GALLAS, Elizabeth (University of Oxford (GB))

Session Classification: Track 2: Offline Computing

Track Classification: Track 2: Offline Computing