

# The ATLAS EventIndex: possible evolution towards an Event Whiteboard

*Tuesday 30 May 2017 09:00 (20 minutes)*

The ATLAS EventIndex was designed during LS2 to satisfy a small but important number of use cases, primarily event picking. Its contents and storage architecture were tailored to the primary use cases, favouring performance and robustness over the possibility to expand its scope. The EventIndex is in operation since the start of Run2 and shows satisfactory performance for event picking and also for additional use cases like production consistency checks and trigger and derivation dataset overlap counts. Discussions are starting about the possibility to increment its contents and extend its scope for Run3. The general idea is to have an “Event Whiteboard” where the initial metadata can be supplemented with information produced during event processing tasks, like references to the algorithms that were run and their results, especially concerning event selections. With this tool one could construct “virtual datasets”, consisting only of the lists of events that satisfy some group of selection criteria, avoid creating many copies of the same events on storage, and run jobs to extract the main parameters for histogramming and fitting only at the end of the analysis tasks. Of course a thorough study of the available technologies has to be launched first, to make sure that this system can sustain the expected i/o rates.

**Author:** BARBERIS, Dario (Università e INFN Genova (IT))

**Presenter:** BARBERIS, Dario (Università e INFN Genova (IT))

**Session Classification:** Going beyond relational

**Track Classification:** Going beyond relational