The ATLAS Event Index: The Architecture of the Core Engine

Dario Barberis[1], Zbigniew Baranowski[2], Andrea Favareto[1], Julius Hřivnáč[3], Fedor Prokoshin[4], Grigori Rybkin[3], Rainer Toebbicke[2] and Ruijun Yuan[3] on behalf of the ATLAS Collaboration

[1] Universita di Genova and INFN Sezione di Genova, Via Dodecaneso 33, 16146 Genova, Italy
[2] CERN, Route de Meyrin 385, 1217 Meyrin, Switzerland
[3] Linear Accelerator Laboratory, Batiment 200, 91440 Orsay, France
[4] Federico Santa Maria Technical University Valparaiso, Avenida de, Av. Espana 1680, Valparaiso, Region de Valparaiso, Chile

E-mail: Julius.Hrivnac@cern.ch

Abstract.

The global view of the ATLAS Event Index system has been presented in the 17th ACAT Conference. This article concentrates on the architecture of the system core component. This component handles the final stage of the event metadata import. It organizes its storage and provides a fast and feature-rich access to all information. A user is able to interrogate metadata in various ways, including by executing user-provided code on the data to make selections and to interpret the results. A wide spectrum of clients is available, from a set of Linux-like commands to an interactive graphical Web Service. The stored event metadata contain the basic description of the related events, the references to the experiment event storage and the full trigger record and can be extended with other event characteristics. Derived collections of events can be created. Such collections can be annotated and tagged with further information.

1. ATLAS Event Index

The ATLAS [1] Event Index [2] (presented in the 17th ACAT Conference [3]) is a system which provides experts and users with event-level metadata services. A user can find information about specific events from the Event Index data rather than from event dataset files. Event Index stores data GUIDs (Global Unique IDentifiers) as well as a variety of other event-wise information. All data are organized in TagFiles, which are currently implemented as Hadoop [4] MapFiles. A set of TagFiles is a TagSet, which can be accessed in the same way as a TagFile. All TagFiles are registered in a Catalog. Each operation on the Event Index database is implemented as a transformation of TagFiles, its result is again a standard TagFile subject to common interface. TagFiles can be processed in the most general way by applying any valid (Java) code to them. All operations are also registered in a Journal.

The system is integrated in the ATLAS offline environment and is available from all its distributions and via an Apache/Tomcat Web Service, which runs on two load-balanced servers. It runs on a Hadoop cluster provided by CERN IT. Currently, it handles over 80 billion events, both data and Monte Carlo.

2. Interfaces

External access to data is available via several interfaces:

- *Linux-like command line* which requires a small local client.
- *Interactive Graphical Interface* is just a simple graphical version of the command line interface for various commands (like Event Lookup command `el` (Figure 2)). Results are represented in a graphical view (like dataset overlaps table (Figure 4) or trigger statistics table and histogram (Figure 3)). The data oriented interface (Figure 1) presents all existing datasets with actions available to each of them.
- *Stateless Web Service* serves as an interface to other ATLAS services.

All interfaces implement the same functionality. External access is protected by the CERN Shibboleth [5] system.

3. Commands

The most generic access is possible by using two general commands available as Linux commands and as Web Service):

- `catalog` command allows searching and modification of the TagFile Catalog.
- `ei` command gives full searching access to all data. Result of each request is a new TagFile. User can specify search and assembly criteria using any valid (Java) code. This command is very rich, but more complex to use. There are several special auxiliary processors available to users, allowing more complex operations, like generating statistics histograms for selected TagFile variables.

Some frequently-used operations are available as special wrapper commands. The most important are:

- `el` is Event Lookup command to locate of events based on their event numbers.
- `ti` is Trigger Interface command to create special statistics tables about event triggers.
- `inspect` gives more direct access to TagFile content. It is useful for detailed study and for TagFiles with non-standard schema.

4. Auxiliary HBase Tables

Some data are assembled in auxiliary HBase [6] tables.

Catalog table keeps all information about TagFiles, their content, format, genealogy and basic characteristics. Catalog search is in most cases the first stage of general search.

Journal keeps track of all non-trivial operations. It allows to search for already delivered results, helping to investigate possible problems and to find usage patterns and usage statistics.

Event Lookup tables (one for real data and one for Monte Carlo data) give very fast mapping of event numbers to containing files. Tables are generated from imported HDFS files and keep references to them to allow more detailed searches and reporting. They are used by the `el` command.

5. Future improvements

5.1. Query Spaces

All TagFiles can be considered as entities in the TagFile space. All operations (derivations, merging, etc.) then create relations between those entities. We can develop a measure allowing to assess distances between entities. A user then can get an immediate approximate (or sometimes exact) result to the request based on distance to nearest already available TagFiles because results themselves are represented by TagFiles.

5.2. Universal Command Line Client

A downloadable application running on all usual platforms (Linux, MS Windows, Mac, Android) is under development. It will give command-like interface, so far only available on Linux, to all users.

5.3. Conclusion

The ATLAS Event Index system is in production and provides a succesfull service to the community. Its modular architecture allows rapid adaptation to new user requirements. New features and improvements are under active development.

References

[1] The ATLAS Collaboration 2008 The ATLAS Experiment at the CERN Large Hadron Collider JINST 3 S08003
[2] Barberis D et al 2015 The ATLAS EventIndex: architecture, design choices, deployment and first operation experience J. Phys.: Conf. Ser. 664 042003
[3] D Barberis et al 2016 J. Phys.: Conf. Ser. 762 012028
[4] Hadoop http://hadoop.apache.org
[5] Ormancey E 2008 CERN single sign on solution J. Phys.: Conference Ser. 119 (8) 082008
[6] HBase http://hbase.apache.org

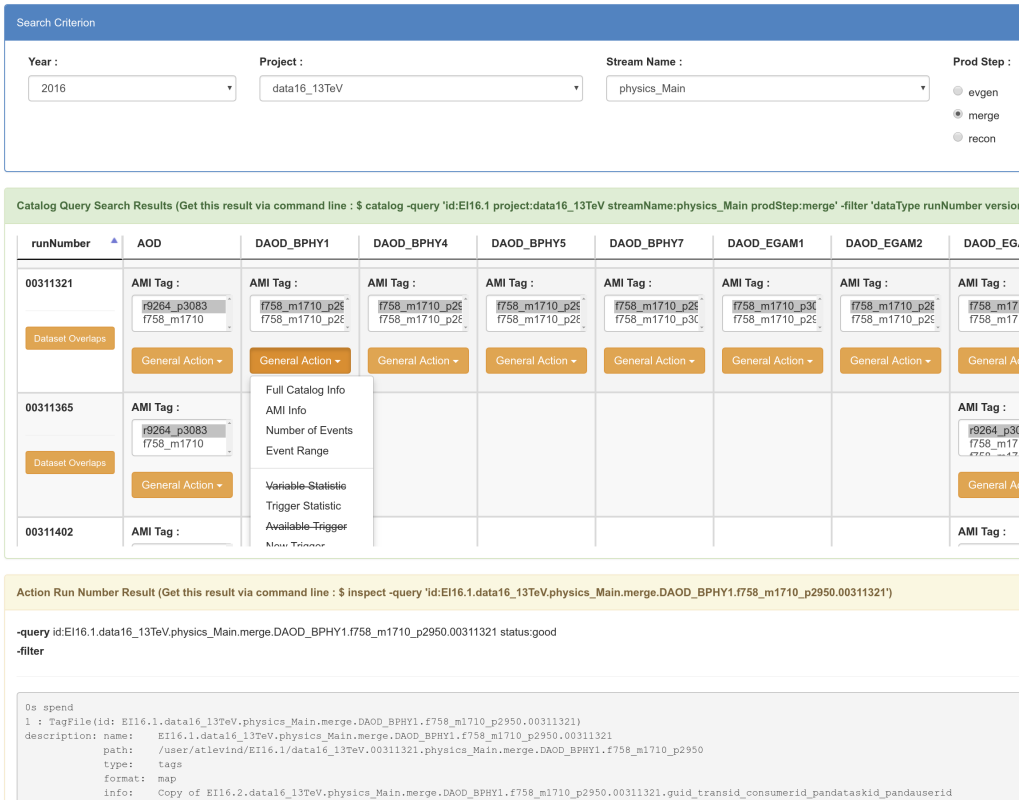

Figure 1. Data oriented Web Service proposes all actions available for a particular dataset.

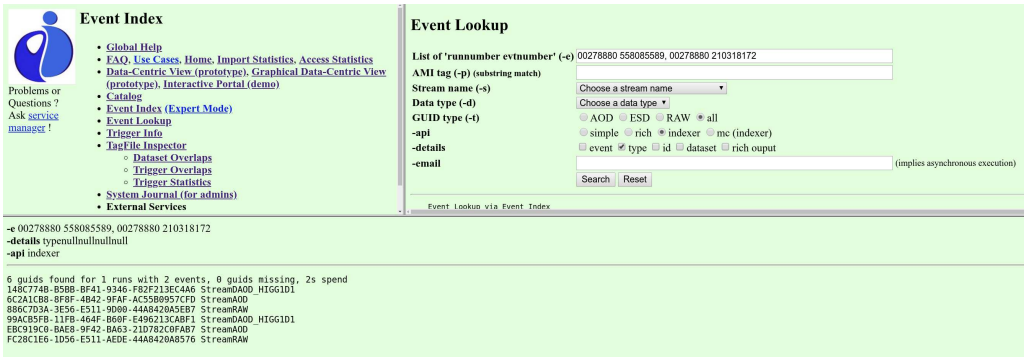

Figure 2. Generic Event Lookup Web Service.

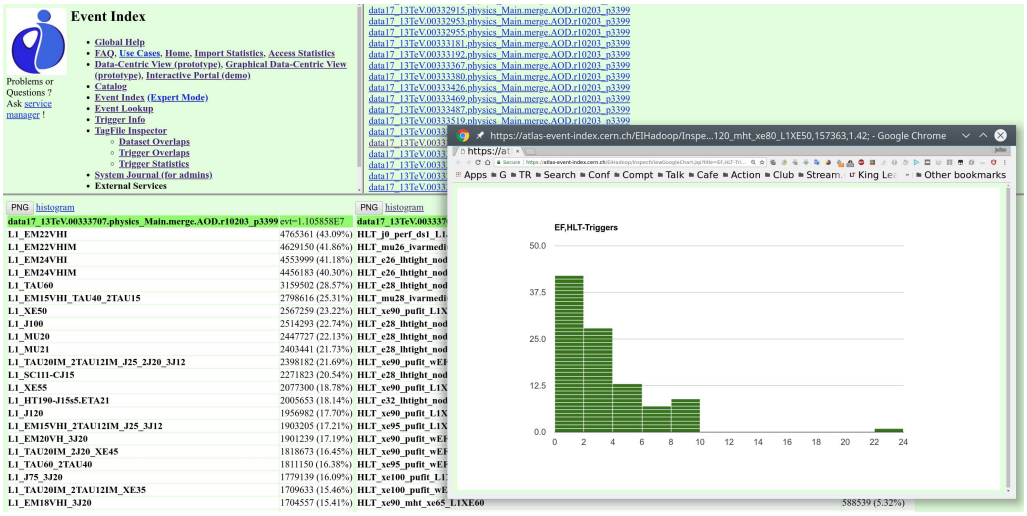

Figure 3. Trigger Statistics Web Service.

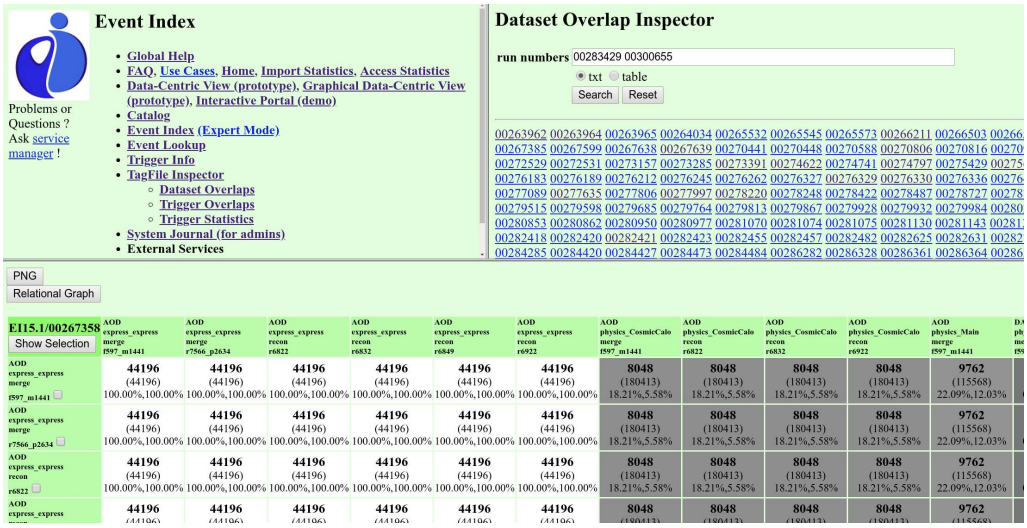

Figure 4. Web Service showing overlap between datases.