

Query Service for new ATLAS EventIndex system

Grigory Rybkin
(on behalf of the ATLAS Collaboration)

Université Paris-Saclay, CNRS/IN2P3, IJCLab, 91405 Orsay, France



31 May 2023, Restitution CHEP



The ATLAS EventIndex system consists of the catalogue of all events collected, processed or generated by the ATLAS experiment at the CERN LHC accelerator, and all associated software tools. The new system, developed for LHC Run 3, makes use of Apache HBase — the Hadoop database — and Apache Phoenix — an SQL/relational database layer for HBase — to store and access all the events metadata. The Query Service allows searches for and retrieval of information of interest. Its design is based on a server-client architecture with the sides communicating via a RESTful API. The OpenAPI Specification is used to describe the API and generate parts of the source code of the server and client. Based on selection criteria received from the client, the server executes SQL queries against the Phoenix tables, retrieves the data, and returns it to the client. The client — a command line utility following the UNIX/POSIX-conforming program design style — can output results in various formats including plain text and JSON. The server and the client are developed in Java 8, with the use of the Jersey 2 JAX-RS API implementation and the Spring framework. The service performance is suitable for interactive use with queries for up to 1000 events at a time.



The ATLAS experiment [1] records several billion particle interactions every year of operation and produces even larger simulated data samples. The ATLAS EventIndex [2] is a BigData catalogue for all ATLAS events. Each EventIndex record includes summary information on the event itself and the pointers to the files containing the full event. The global catalogue allows to search for and retrieve specific events for in-depth investigations. EventIndex is an important tool, necessary to the good operations of the ATLAS collaboration.



A new major version of EventIndex was developed and put in production for Run3 of LHC.

The components of the EventIndex system are implemented using Big Data technologies and open-source and free software projects and tools. In particular, the projects

- Apache Hadoop — a framework that allows for the distributed processing of large data sets
- Apache HBase — the Hadoop database
- Apache Phoenix — an SQL/relational database layer for HBase

The **Query Service** is a key component of EventIndex.



server-client communicating via RESTful API
OpenAPI Specification [3] is used to

- describe the API
- generate parts of source code

client

- sends user specified selection criteria to server, returns server response to the user
- in Java 8, Jersey 2 JAX-RS API implementation [4]

server

- executes SQL queries against the Phoenix [5] tables — EventIndex' database
- retrieves the data, and returns it to the client
- in Java 8, JDBC (Java DataBase Connectivity), Spring framework [6]



Architecture and Technologies (cont'd)

- parallelism implemented
 - ▶ by adapting the SQL queries, then done by Phoenix
 - ▶ by means of Java Streams, then done by Java runtime
- configuration means provided by the Spring framework
 - ▶ configuration file, environment variables, command line options




Figure: Architecture and Technologies.



set of command line utilities

follows UNIX/POSIX-conforming program design style

- convenient to use with and by other programs

several output representation formats, e.g.,

- plain — lines of space separated fields, for processing with UNIX tools
- json — object-oriented representation, to use programmatically
- pretty — more human-readable json

possibility to read the input from the utility standard input

convenience features

- each option has two forms, short and long, giving a human readable name to the option [7], e.g., `--project`, `--data-type`
- positional arguments and options can appear in any order (be interspersed)



event-lookup utility for “event picking”

```
usage: event-lookup [-h] [-v] [-p PROJECT] [-s STREAMNAME] [-S PRODSTEP]
                  [-d DATATYPE] [-a VERSION] [-D GUIDDATATYPE] [-F FILE]
                  [-c plain,name,short,pretty,json] [-o OUTPUT]
                  [event [event ...]]
```

Search (indexed datasets) for GUIDs of files with events. The files with found GUIDs may be part of the searched datasets or provenance datasets.



- look up event 348895:3613825 (run number 348895, event number 3613825), output dataset name

```
$ event-lookup -c name 348895:3613825
348895 3613825 4aa369dc-fc28-3444-aab5-c2fdf27a9c69 DAOD_HDBS3\
  data18_13TeV.00348895.physics_Main.deriv.DAOD_HDBS3.f937_m1972_p4928
348895 3613825 78c6837f-f448-e811-ae3c-44a8420a7621 RAW\
  data18_13TeV.00348895.physics_Main.deriv.DAOD_HDBS3.f937_m1972_p4928
348895 3613825 0d1c27ff-c5b2-6b46-a0fa-edbdc51c6b77 AOD\
  data18_13TeV.00348895.physics_Main.deriv.DAOD_HDBS3.f937_m1972_p4928
```

- look up event 348895:3613825 (run number 348895, event number 3613825), select GUIDs of files with RAW data type, output format short — suitable for use with PanDA jobs that are given `--eventPickWithGUID` option

```
$ event-lookup -c short 348895:3613825 --guid-data-type RAW
348895 3613825 78c6837f-f448-e811-ae3c-44a8420a7621
```



event-lookup (JSON)

- look up event 348895:3613825, print the result in human readable JSON

```
$ event-lookup -c pretty 348895:3613825
[ {
  "eventNumber" : 3613825,
  "guid" : "4aa369dc-fc28-3444-aab5-c2fdf27a9c69",
  "prov" : [ {
    "dataType" : "RAW",
    "guid" : "78c6837f-f448-e811-ae3c-44a8420a7621"
  }, {
    "dataType" : "AOD",
    "guid" : "0d1c27ff-c5b2-6b46-a0fa-edbdc51c6b77"
  } ],
  "dataset" : {
    "runNumber" : 348895,
    "project" : "data18_13TeV",
    "streamName" : "physics_Main",
    "prodStep" : "deriv",
    "dataType" : "DAOD_HDBS3",
    "version" : "f937_m1972_p4928",
    "name" : "data18_13TeV.00348895.physics_Main.deriv.DAOD_HDBS3.f937_m1972_p4
  }
}
```



- look up events from FILE (one event per line), 2 events at a time, print each command line prior to invocation

```
$ <FILE xargs -r -d'\n' -n 2 -t event-lookup -s physics_Main
```

- support for wildcard characters in query parameters

```
$ event-lookup 348895:3613825 -a '*1972_?4928' -p 'da*'
348895 3613825 4aa369dc-fc28-3444-aab5-c2fdf27a9c69 DAOD_HDBS3 data18\
physics_Main deriv DAOD_HDBS3 f937_m1972_p4928
348895 3613825 78c6837f-f448-e811-ae3c-44a8420a7621 RAW data18_13TeV\
physics_Main deriv DAOD_HDBS3 f937_m1972_p4928
348895 3613825 0d1c27ff-c5b2-6b46-a0fa-edbdc51c6b77 AOD data18_13TeV\
physics_Main deriv DAOD_HDBS3 f937_m1972_p4928
```



dataset-list utility to query dataset information

```
usage: dataset-list [-h] [-v] [-p PROJECT] [-s STREAMNAME] [-S PRODSTEP]
                  [-d DATATYPE] [-a VERSION] [-F FILE]
                  [-c plain,name,short,pretty,json] [-o OUTPUT]
                  [run [run ...]]
```

List (indexed) datasets (metadata) for the runs (all by default).

event-list utility

```
usage: event-list [-h] [-v] [-B LB] [-C BCID] [--lpsk LPSK] [--id ID]
                 [--lb1 LB1] [--bcid1 BCID1] [--hpsk HPSK] [-l LIMIT]
                 [-S OFFSET] [-r] [-N] [-F FILE]
                 [-c {plain,long,short,pretty,json}] [-o OUTPUT]
                 [name [name ...]]
```

List metadata for or count events satisfying the selection criteria from the datasets.



References

- 1 Collaboration ATLAS (2008) The ATLAS experiment at the CERN Large Hadron Collider. *JINST* 3:S08003.
<https://doi.org/10.1088/1742-6596/331/7/072034>
- 2 Barberis, D., Alexandrov, I., Alexandrov, E. et al. The ATLAS EventIndex. *Comput Softw Big Sci* 7, 2 (2023).
<https://doi.org/10.1007/s41781-023-00096-8>
- 3 OpenAPI Specification: <https://swagger.io/specification>
- 4 Jersey 2 JAX-RS API implementation:
<https://eclipse-ee4j.github.io/jersey>
- 5 Phoenix: <https://phoenix.apache.org>
- 6 Spring framework: <https://spring.io>
- 7 GNU Standards for Command Line Interfaces:
https://www.gnu.org/prep/standards/html_node/Command_002dLine-Interfaces.html#Command_002dLine-Interfaces