

# **PERFORMANCE TESTING FRAMEWORK FOR THE ATLAS EVENTINDEX**

**E. Cherepanova<sup>1,a</sup>, F. Prokoshin<sup>1</sup>**

**On behalf of the Software and Computing Activity**

<sup>1</sup> *Joint Institute for Nuclear Research, Joliot-Curie 6, RU-141980 Dubna, Russia*

E-mail: <sup>a</sup> elizaveta.cherepanova@cern.ch

The ATLAS EventIndex is going to be upgraded in advance of LHC Run 3. A framework for testing the performance of both the existing system and the new system has been developed. It generates various queries (event lookup, trigger searches, etc.) on sets of the EventIndex data and measures the response times. Studies of the response time dependence on the amount of requested data, and data sample type and size, can be performed. Performance tests run regularly on the existing EventIndex and will run on the new system when ready. The results of the regular tests are displayed on the monitoring dashboards, and they can raise alarms in case (part of) the system misbehaves or becomes unresponsive.

Keywords: Scientific computing, BigData, Hadoop, EventIndex

Elizaveta Cherepanova, Fedor Prokoshin

Copyright © 2021 for this paper by its authors.  
Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

## 1. Introduction

The ATLAS EventIndex [1] is the tool that collects, checks and stores information about the main properties of all real or simulated events that were collected, processed or generated by the ATLAS experiment [2], and points to the files that contain them. The current storage implementation is based on having full data on the Hadoop [3] system and reduced information in Oracle [4] tables. The Hadoop system runs a variety of tasks, such as importing and cataloguing data, running consistency checks, establishing links between related datasets and responding to users' queries. Interactions with users are high-priority tasks. The system should be available and accessible under various conditions and give a response in the appropriate time, depending on the request. The daily access statistics are represented in Figure 1. To track the current system state and evaluate its performance, a variety of tests were developed.

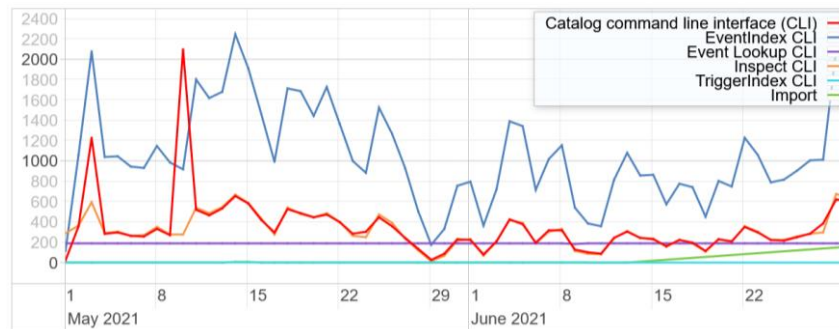


Figure 1. Daily access statistics of the Hadoop system by different EventIndex services between May and June 2021

## 2. Data for tests

The tests are carried out using two types of input data:

- a key, defined as a runNumber-eventNumber pair, e.g., “278880-251772208”.
- a full dataset name, e.g., “data16\_13TeV.00299584.physics\_Main.deriv.DAOD\_HIGG1D1.r9264\_p3083\_p4096”

The list of about 50 000 keys prepared for a physics analysis was used for the tests. This includes real data recorded in 2015-2018. For the tests several samples were made with keys from all years and for several selected datasets with a total size of 1 million events.

## 3. Types of queries

The EventIndex allows searching data using different types of input information and different commands for the search. Two of them are used:

- EventLookup (el) for fast search of the physical datasets corresponding to an event (specified as a pair of run number and event number or a key).
- EventIndex (ei) to search all datasets using either direct searches or complex Map/Reduce jobs. This can take both a key and a dataset name as input data.

All performed queries can be split into two groups:

- Event picking (performed using “el” and “ei” commands)
  - Fast search of events using key pairs.
  - Return GUIDs (Global Unique Identifier) of events, full dataset names and data types (RAW, AOD, DAOD).
- Search for events in a specified dataset (performed using “ei” command)
  - Return full information about an event (GUID, data type, production step, luminosity block, time, trigger chains, etc.)
  - Different filters can be applied.

## 4. Performance tests

### 4.1 EventLookup

The event lookup is performed through the "el" client command. The search is running every hour in a cron scheduler. Eight lists of keys are retrieved randomly each time to avoid using cached results. The summary information about data used and the performance tests results is listed in Table 1.

Table 1. Data used for EventLookup tests and tests result. 2021/06/18 – 2021/06/25

Data	Source of keys	Number of keys	Average execution time, s
data 2015	1 million events dataset	1 000	73.8
data 2015	several datasets	1 000	37.7
data 2017	several datasets	1 000	21.3
data 2018	1 million events dataset	10	4.2
data 2018	1 million events dataset	100	7.5
data 2018	1 million events dataset	1 000	30.5
data 2018	1 million events dataset	10 000	399.0
data 2018	1 million events dataset	50 000	964.2

The results of the tests are displayed in the Grafana Monitor [5]. Figure 2 shows the response times of the Hadoop server to event lookup queries selecting 1000 events out of a dataset with one million records or a mixture of several datasets as a function of time. The occasional glitches are due to other activities on the servers at the time of the queries. The response times are dominated by the query time for low numbers of events, and by the transmission time of the output for large numbers of events.




Figure 2. Response times of the EventIndex Hadoop server to event lookup queries selecting 1000 events out of a dataset with 1 million records or a mixture of datasets as a function of time, recorded between 2021/06/18 and 2021/06/25

### 4.2 EventIndex queries. No selection

The event picking and full dataset information retrieval are running once a day in a cron. These queries are performed through the "ei" client command. The key search is performed using one list with 10 keys for each year. The full dataset information retrieval is using datasets with sizes 10k (50k), 100k, 1M, 10M events for 2015-2018 data.

Figure 3 shows results of the tests for 2015 and 2017 datasets. The response times are dominated by the setup time of the Map/Reduce job for low numbers of events, and by the transmission time of the output record for large numbers of events.




Figure 3. Response times of the EventIndex Hadoop server to queries using Map/Reduce jobs to retrieve information on all events from datasets containing 10k, 100k, 1M and 10M events as a function of time, recorded between 2021/06/11 and 2021/06/25. Left: 2015 data, right: 2017 data

### 4.3 EventIndex queries. Trigger selection

The search of full dataset information for events that passed trigger selection is running in a cron. These queries are performed through the “ei” client command. The summary information about data used is listed in Table 2.

Table 2. Data used for EventIndex tests with trigger selection

Data	Total number of events	% of events passed trigger	Running
data 2015	500k, 1M, 10M, 50M	<1	twice a day
data 2016	500k, 1M, 10M, 50M	<1	twice a day
data 2017	500k, 1M, 10M, 50M	<1	twice a day
data 2018	500k, 1M, 10M, 50M	<1	twice a day
data 2018 fast	500k	<1	hourly
data 2018 long	50M	<40	once a day

“Data 2018 fast” is the data used to check the system availability. A dataset of 2018 with 500k events with <1% fraction of events passing a specified trigger was chosen because of the short execution time – about 2 minutes (see Figure 4).

“Data 2018 long” is the data used to check how the system can handle heavy task. A dataset of 2018 with 50M of events with <40% fraction of events passed a specified trigger was chosen because of long execution time – about 4 hours (see Figure 4).

The response times of the Hadoop server to queries searching events that passed specified triggers from datasets containing 500k and 50M events are shown in Figure 4 as a function of time.




Figure 4. Response times of the EventIndex Hadoop server to queries using Map/Reduce jobs searching events that passed specified triggers from datasets containing 500k and 50M events as a function of time, recorded between 2021/06/18 and 2021/06/25. The right axis applies to the long search, the left axis applies to the other queries

#### 4.4 Performance time dependence on the way of running queries

There are two ways of running queries: locally on the Hadoop cluster or remotely through the Internet via a Tomcat server [6]. The test described above are running locally, however a regular user usually runs queries remotely.

For comparison some of the tests were launched remotely. The response times for event lookup queries with search for data 2015 and for EventIndex queries with trigger selection for data 2018 are shown in Figure 5.

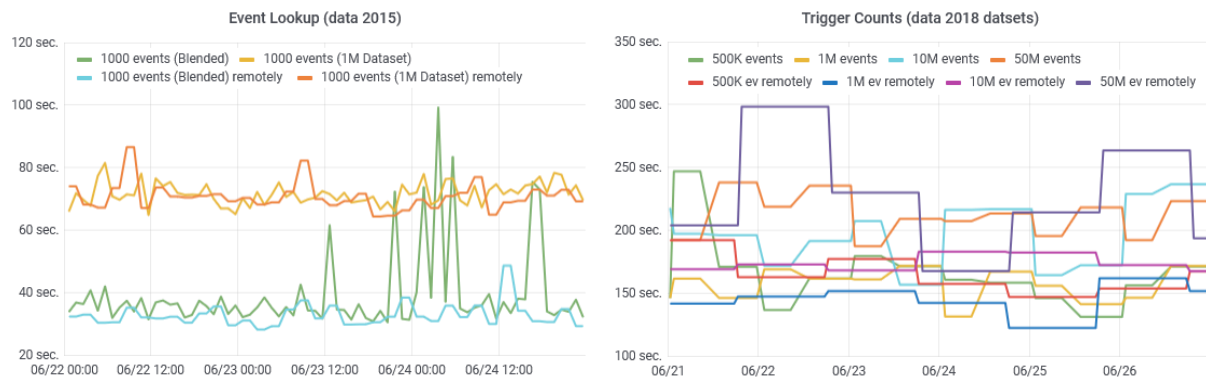


Figure 5. The dependence of response time on the way of running queries. Left: Event lookup queries for 2015 data: Blended – keys retrieved from several datasets, 1M Dataset – keys retrieved from one dataset with a million events. Right: EventIndex queries with trigger search for 2018 data

Remote queries run faster for almost all event lookup queries. The larger number of keys for search, the closer are performance times of queries launched locally and remotely. The EventIndex queries with trigger selection show similar performance time.

## 5. Conclusions

Performance tests for the existing ATLAS EventIndex system has been developed. Three types of jobs are running in a cron:

- Event lookup key search (once per hour)
- EventIndex key search and EventIndex full dataset information retrieval (once per day)
- EventIndex dataset search with trigger selection

The first and third are displayed using Grafana dashboards, where it is easy to track the current state of the system. Most of the locally running tests are also run remotely. Performance times show that for event picking tasks the local launch can be slower. The results of the tests are stored in Hadoop and can be used for comparison with performance of the newly-developed system implementation.

## References

- [1] Barberis D et al 2015 The ATLAS EventIndex: architecture, design choices, deployment and first operation experience, J. Phys.: Conf. Ser. 664 042003, doi: [10.1088/1742-6596/664/4/042003](https://doi.org/10.1088/1742-6596/664/4/042003)
- [2] ATLAS Collaboration 2008 The ATLAS Experiment at the CERN Large Hadron Collider, JINST 3 S08003 doi: [10.1088/1748-0221/3/08/S08003](https://doi.org/10.1088/1748-0221/3/08/S08003)
- [3] Hadoop and associated tools: <http://hadoop.apache.org>
- [4] Oracle: <https://www.oracle.com>
- [5] Grafana: <https://grafana.com>
- [6] Tomcat: <https://tomcat.apache.org>