

DEVELOPMENT OF THE EVENT METADATA SYSTEM FOR THE NICA EXPERIMENTS

E.I. Alexandrov¹, I.N. Alexandrov^{1,a}, A.G. Degtyarev², I.A. Filozova¹, K.V. Gertsenberger¹, P.A. Klimai², A.V. Yakovlev¹

¹*Joint Institute for Nuclear Research, 6 Joliot-Curie st., Dubna, 141980, Russia*

²*Moscow Institute of Physics and Technology, 9 Institutskiy per., Dolgoprudny, 141701, Russia*

E-mail: ^a alexand@jinr.ru

Particle collision experiments are known to generate a substantial amount of data that must be stored and then analyzed. Typically, only a small subset of all collected events is relevant when performing a particular physics analysis task. Although it is possible to obtain the required subset of event records directly, by iterating through the whole volume of the collected data, the process is very time and resource consuming. Instead, a more convenient approach widely used in modern particle collision experiments is to have an event metadata or indexing system that stores summary properties of all events and allows for the fast search for and retrieval of only necessary events based on given criteria. Such a system, called the Event Metadata System, has been developed for the current fixed-target and future collider experiments of the NICA project. The design of the system, its components, user interfaces and the REST API service, its integration with existing experiment systems and software, as well as associated challenges, are presented.

Keywords: particle collision experiment, data management, event metadata, event index, event database, NICA experiments, information system

Evgeny Alexandrov, Igor Alexandrov, Artyom Degtyarev, Irina Filozova, Konstantin Gertsenberger, Peter Klimai, Alexander Yakovlev

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 aim of the NICA (Nuclotron-based Ion Collider fAcility) project is to study hot and dense strongly interacting matter in heavy ion collisions at center-of-mass energies up to 11 GeV for gold ions, and up to 27 GeV for protons to study the spin structure and spin-related phenomena [1]. The NICA complex is currently under construction and will include, in its final form, both the fixed-target experiment, BM@N (Baryonic Matter at Nuclotron), and collider experiments, MPD (MultiPurpose Detector) and SPD (Spin Physics Detector). At present, experimental data from BM@N collected since 2015 (overall about 500 million events), as well as simulated data of both BM@N and the collider experiments, are available at distributed storages of the NICA computing system.

A large amount of data produced by particle collision experiments, including those of NICA, must be stored, then processed and analyzed. The data stream from the detector electronics up to the data storage of the NICA experiments is organized as a sequential data-driven pipeline [2], which includes data acquisition, online processing, and offline processing stages on distributed computing platforms. Physics analysis, which is a final stage of offline processing, is performed to obtain the desired physics results of an experiment. Physics analysis is typically carried out for a subset of events satisfying given criteria, such as events for a selected beam particle, or that lead to the production of certain particles, etc. To avoid a full scan of all files with the collected event data, an event metadata or indexing system is required.

2. Goals and requirements for the Event Metadata System

The main goal of the developed system was to provide a searchable event catalogue with basic information (metadata) on all reconstructed experimental collision events. Full experimental data are contained in ROOT files of the experiments, while the Event Metadata System (EMS) stores only the necessary event metadata, such as the number of primary and all tracks, track number of positively (negatively) charged particles, primary and secondary particles found in the event, number of hits by the detectors, total input and output charge in the event, etc. The exact set of stored summary event attributes (metadata) varies per experiment, so the system should be configurable to store a basic and arbitrary set of parameters per event being specific to the experiment. In addition, the software version used to obtain event data and a reference to the storage location are recorded for each event in the Event Catalogue.

Additional features of the EMS are providing statistics on the collected data and checking the quality of the catalogue of physics events. The requirements for the Event Metadata System include scalability (several billion events per year are expected for the NICA experiments), performance (an important characteristic is the processing time of one request), availability and fail safety, role-based access control. Access to metadata should be provided via several application programming and user interfaces, including Web service, REST API (application programming interface), and ROOT C++ methods.

3. Event Metadata System architecture and its developed components

The overall architecture of the Event Metadata System is shown in Fig. 1. The information system is based on the Event Database called Event Catalogue, which contains specific information on particle collision events, allowing the user to quickly search for a set of events required for particular physics analysis by various criteria. The Event Database identifies stored events with a file pointer defining a reference to the data file and event number (a unique integer within a run) in the file, which uniquely matches each event and allows finding it in the distributed file systems of the experiments. The Catalogue has been built on top of the PostgreSQL database management system (DBMS). The choice of the DBMS has been made for both performance and manageability reasons [3].

Run metadata for the NICA experiments, such as a period and run numbers, beam and target particles, and beam energy, are stored in the developed Condition Database [4]. To improve data query performance, run metadata are used to preselect experimental run numbers matching given criteria, before performing a search in the Event Catalogue.

The web user interface (UI) and REST API service, discussed in detail in the next section, are used to perform queries for required events to the Event Metadata System using various search criteria. Not only event metadata can be obtained via the developed interfaces, but also full event data merged into ROOT files containing a tree with the necessary events can be requested. A special Gathering Service is used to solve this task. Another service, Dispatcher, is employed to transfer event metadata directly to the workload management system of the experiments for further data processing or physics analysis on the NICA computing platform.

Authentication and authorization in the Event Metadata System have been implemented on the basis of the FreeIPA system, which is also used for the other information systems of the experiment. Each request to the EMS must be authenticated, and the available operations depend on the assigned user role: Index Administrator (full access, including the deletion of possibly erroneous data), Index Writer (can write new events, but not delete existing ones), Event Consumer (read-only user).

It is imperative that the system is ready and available to handle user requests at any time, so a monitoring service for the Event Metadata System has been implemented. The service is based on the following software stack: Telegraf, InfluxDB, Grafana, as well as custom Python scripts.

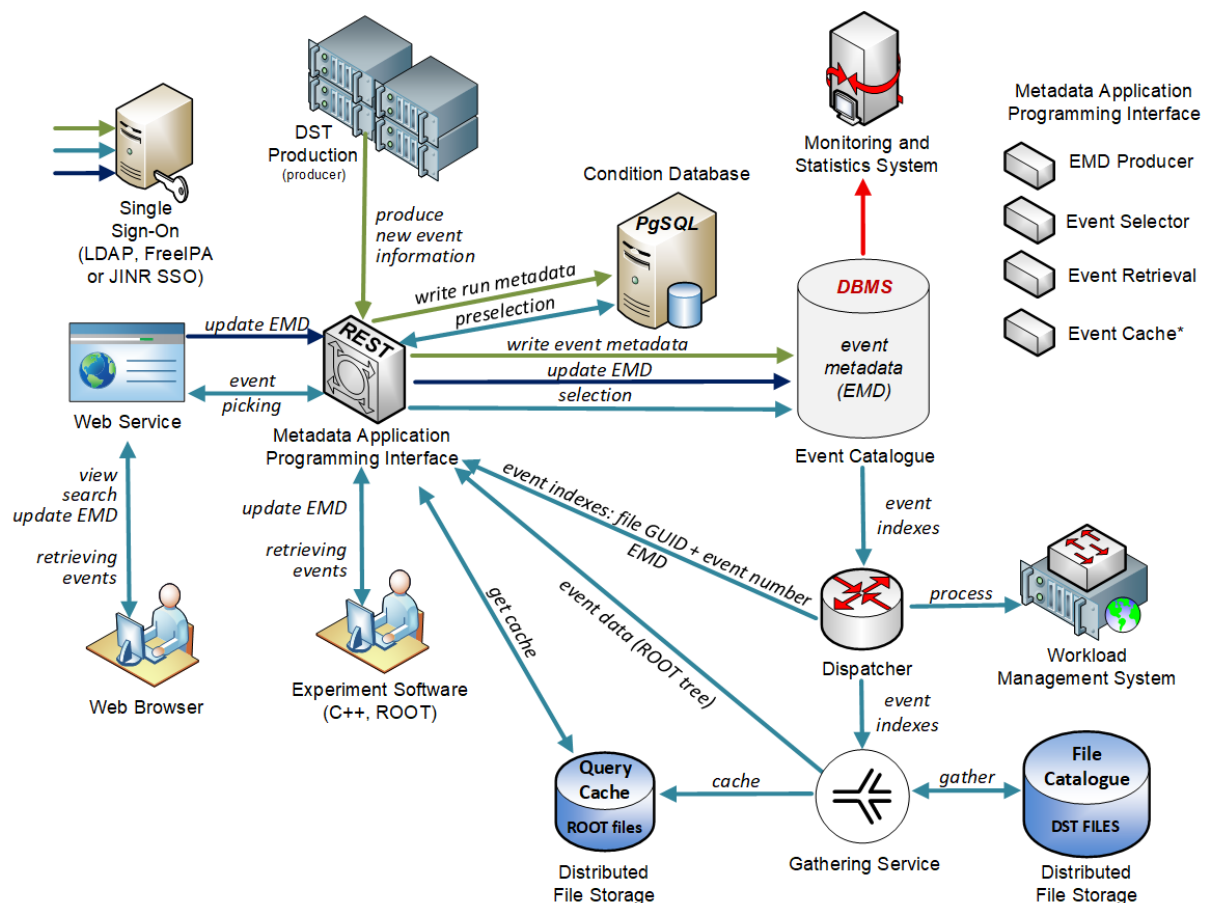


Figure 1. Common architecture of the Event Metadata System

4. Implementation of the REST API and Web UI for the EMS

Both the Web UI and REST API for the Event Metadata System have been implemented using the Kotlin programming language, ktor framework [5], and Java Virtual Machine runtime. For ease of provisioning, the components of the EMS architecture are packed in Docker containers.

The configuration for the REST API and Web UI services is provided in YAML format, specifying database connection parameters (for both the Event Catalogue and Condition database), base URLs for the Web UI and REST API, and custom metadata fields (defined by their name and type for each field).

The first version of the Web UI view (deployed for the BM@N experiment) is given in Fig. 2. It can be seen in the figure that base metadata (period number, run number, software version), Condition Database fields (beam and target particles, beam energy), and custom event attributes being specific to the experiment (for instance, positively charged track number) can be used to select the physics events required for particular analysis or data processing.

BM@N Events

Enter search criteria for events

Period Number

Run Number

Software Version

} Selection based on
standard parameters

Beam Particle

Target Particle

Energy, GeV

} Preselection based on
Condition DB

Total track number

} Selection based on
configured parameters

Events found:

storage_name	file_path	event_number	software_version	period_number	run_number	track_number
data1	/tmp/file1	100	19.1	7	5000	20
data1	/tmp/file1	101	19.1	7	5000	20
data1	/tmp/file1	102	19.1	7	5000	21

Figure 2. First version of the Web UI for the Event Metadata System

The developed REST API service employs JSON formatting illustrated in Fig. 3 to transfer event metadata. The *parameters* object includes the configurable part of event metadata, while the other attributes are the same for events of all the NICA experiments, regardless of configuration.

```
{
- events: [
- {
- reference: {
storage_name: "data1",
file_path: "/tmp/file1",
event_number: 1
},
software_version: "19.1",
period_number: 7,
run_number: 5000,
- parameters: {
track_number: 20
}
},
- {
- reference: {
storage_name: "data1",
file_path: "/tmp/file1",
event_number: 2
}
```

Figure 3. Example of an event metadata record in JSON format

The implemented API methods allow performing the following operations with event metadata:

- GET /count/ to obtain event count only, matching given criteria (available for any user role).
- GET /emd/ to obtain event metadata, matching given criteria (for any user role).
- POST /emd/ to add event metadata to the Event Catalogue (for Index Writer and Index Administrator roles).
- DELETE /emd/ to remove event metadata (for the Index Administrator role only).
- GET /eventFile to obtain a ROOT file with physics events satisfying given criteria (synchronous operation).
- GET /eventFileRef to obtain a reference to a ROOT file with physics events satisfying given criteria (asynchronous operation).

The above REST API methods support the same filtering criteria that are available in the Web UI, for example, to get event metadata for the given period number (7), software version (19.1), range of track numbers (from 10 to 15), and given beam particle (Kr), the user queries the following URL via the HTTP GET method:

```
/emd?period_number=7&software_version=19.1&track_number=10-15&beam_particle=Kr
```

In addition, a specialized C++ interface has been developed to write new metadata to the Event Catalogue while processing input data files with physics events in CERN ROOT format.

5. Conclusion

The first version of the Event Metadata System for the NICA experiments has been developed, and the information system is actively evolving to simplify user access and event metadata management. The Event Catalogue has been built on PostgreSQL and stores event metadata of the experiments to select only the events needed for particular physics analysis. The Web UI and REST API interfaces have been implemented to interact with NICA collaboration members, as well as with the other software systems of the experiments. All the services are available, and the Event Catalogue has been deployed for experimental physics events of the BM@N experiment, the first experiment of the NICA project. In addition, the Event Metadata System has been integrated with the Condition Database developed for the NICA experiments.

6. Acknowledgement

The work was supported by the Russian Foundation for Basic Research (RFBR) grant according to the research project No.18-02-40125.

References

- [1] M. Kapishin, Heavy Ion BM@N and MPD experiments at NICA // JPS Conf. Proc. 32, 010093, 2020.
- [2] A. Baskakov, S. Bazylev, A. Fediunin and I. Filippov, MPD Data Acquisition System: Technical Design Report // JINR, Dubna, 74 p., 2018.
- [3] E. Alexandrov, I. Alexandrov, A. Degtyarev, et al., Design of the event metadata system for the experiments at NICA // Physics of Particles and Nuclei Letters, Vol. 18, 603-616, 2021.
- [4] K. Gertsenberger, I. Alexandrov, I. Filozova, et al., Development of Information Systems for Online and Offline Data Processing in the NICA Experiments // Physics of Particles and Nuclei, Vol. 52, 801-807, 2021.
- [5] Ktor: <https://ktor.io/> (accessed 09.09.2021).