

OPEN ACCESS

The ALICE Electronic Logbook

To cite this article: V Altini *et al* 2010 *J. Phys.: Conf. Ser.* **219** 022027

View the [article online](#) for updates and enhancements.

You may also like

- [ALICE moves into warp drive](#)
F Carena, W Carena, S Chapeland et al.
- [Principles of Plasma Spectroscopy](#)
A.L. Osterheld
- [Relativistic Figures of Equilibrium](#)
M Mars

ECS
The
Electrochemical
Society
Advancing solid state &
electrochemical science & technology

DISCOVER
how sustainability
intersects with
electrochemistry & solid
state science research

The ALICE Electronic Logbook

V Altini², F Carena¹, W Carena¹, S Chapeland¹, V Chibante Barroso¹, F Costa¹,
R Divià¹, U Fuchs¹, I Makhlyueva^{1,3}, F Roukoutakis^{1,4}, K Schossmaier¹, C Soòs¹,
P Vande Vyvre¹ and B Von Haller¹ for the ALICE collaboration

¹ CERN, Physics Department, Geneva, Switzerland.

² INFN, Dipartimento di Fisica dell'Università and Sezione INFN Bary, Italy.

E-mail: Vasco.Chibante.Barroso@cern.ch

Abstract. All major experiments need tools that provide a way to keep a record of the events and activities, both during commissioning and operations. In ALICE (A Large Ion Collider Experiment) at CERN, this task is performed by the Alice Electronic Logbook (eLogbook), a custom-made application developed and maintained by the Data-Acquisition group (DAQ). Started as a statistics repository, the eLogbook has evolved to become not only a fully functional electronic logbook, but also a massive information repository used to store the conditions and statistics of the several online systems. It's currently used by more than 600 users in 30 different countries and it plays an important role in the daily ALICE collaboration activities. This paper will describe the LAMP (Linux, Apache, MySQL and PHP) based architecture of the eLogbook, the database schema and the relevance of the information stored in the eLogbook to the different ALICE actors, not only for near real time procedures but also for long term data-mining and analysis. It will also present the web interface, including the different used technologies, the implemented security measures and the current main features. Finally it will present the roadmap for the future, including a migration to the web 2.0 paradigm, the handling of the database ever-increasing data volume and the deployment of data-mining tools.

1. Introduction

1.1. The ALICE experiment

ALICE (A Large Ion Collider Experiment) [1] is an LHC (Large Hadron Collider) experiment designed to study the physics of strongly interacting matter at extreme energy densities, where the formation of a new phase of matter, the quark-gluon plasma, is expected. Consisting of 18 sub-detectors, it primarily targets heavy-ion collisions (in particular Pb-Pb), but it will also be able to cope with proton-proton (pp) and proton-ion (pA) collisions.

After 15 years of design and installation, a series of dedicated sessions were organized (starting from December 2007) not only to commission the sub-detectors, the online systems and the online-offline interfaces but also to exercise operational activities like 24/7 shift crews, allowing the ALICE experiment to be ready for the LHC startup in September 2008 [2].

³ Now at ITEP Institute for Theoretical and Experimental Physics, Moscow, Russia

⁴ Now at University of Athens, Physics Department, Athens, Greece

1.2. The ALICE Electronic Logbook

In large scientific experiments like ALICE, it's essential to have a bookkeeping facility to keep a record of the experiment's operational activities. As shifters come and go, a central information repository is needed to store reports of incidents, configuration changes, achievements or planned operations. Furthermore, data-taking conditions and statistics are needed to allow not only the selection of good run candidates for offline processing but also to detect trends, correlations and bottlenecks.

In ALICE, these requirements are implemented by the ALICE Electronic Logbook (eLogbook), a custom-made application developed by the ALICE Data Acquisition team (DAQ) [3] and in production since August 2007. This paper will present the architecture of the eLogbook, its main components and interfaces, the current status and the foreseen expansions and improvements.

2. Architecture and components

The eLogbook architecture is based on LAMP (Linux, Apache, MySQL and PHP) software stack, with the relational database (DB) serving as a data repository and the web-based Graphical User Interface (GUI) providing interactive access to members of the ALICE collaboration. A C Application Programming Interfaces (API) and a Command-line Interface (CLI) grant the different automatic actors read/write access to the DB.

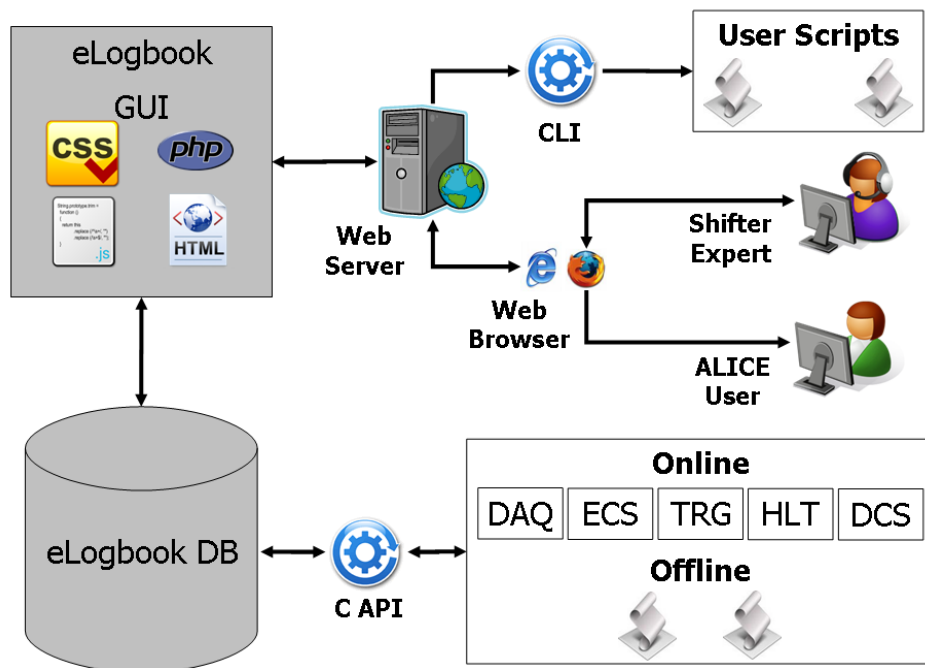


Figure 1. The Alice Electronic Logbook architecture

As shown in figure 1, the different ALICE online and offline subsystems access the DB via the *DAQlogbook* API and the shifters and members of the collaboration via a web browser. The *logbookGetRunsInfo* CLI is provided for users that want to develop applications reading from the DB via the web server in a non-interactive way.

2.1. Database

The DB, running on a MySQL 5.1 Community Server, is used to store heterogeneous data related with the experiment's activities. InnoDB is used as a storage engine for its support of both transactions and

foreign keys constraints. As shown in figure 2, the tables that compose this DB can be grouped into 3 different categories:

- RUN CENTERED: related to a specific run (data-taking period identified by a unique ID called 'Run Number')
- LOG ENTRY CENTERED: related to a specific human or automatic text report with optional file attachment
- USER CENTERED: related to the GUI users

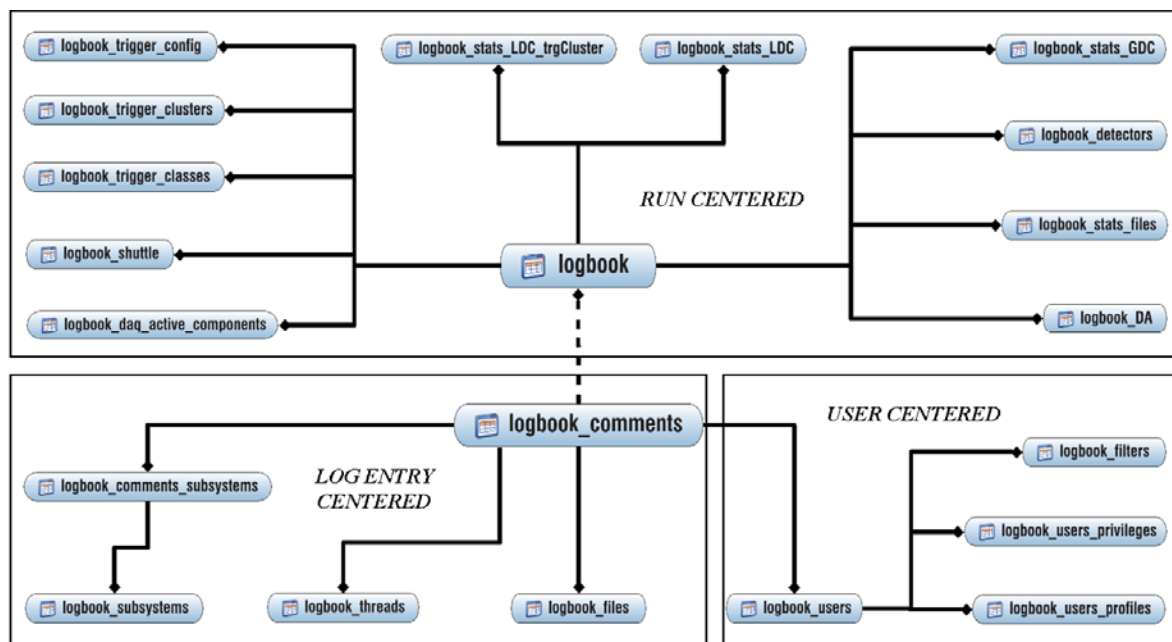


Figure 2. The ALICE Electronic Logbook DB schema

A stored procedure, called `update_logbook_counters`, is executed every X seconds (plus at the end of each run) to update the different global counters in the `logbook` table whose value depends on partial counters spread throughout several tables.

Daily backups are performed both to a RAID 6 disk array and the CERN Advanced STORAge manager (CASTOR) [4].

2.2. Graphical User Interface

The eLogbook web-based GUI (<https://cern.ch/alice-logbook>) was developed using modern web technologies, including PHP5, Javascript and Cascading Style Sheets (CSS). It is hosted on an Apache web server and it can be accessed from the experimental area (inside the experiment's technical network), the CERN General Purpose Network (GPN) and the internet.

2.2.1. Authentication and Authorization. Authentication is implemented via the CERN Authentication [5] central service, providing Single Sign On (SSO) and removing the effort of authenticating the users from the eLogbook software. This way, when a user tries to access the GUI, he/she is redirected to the CERN Login page where it has to provide his/her credentials. If successful, he/she is then redirected back to the GUI.

Authorization is implemented in the GUI with 5 different levels of privileges:

- NONE: no access the GUI

- READ: read-only access
- WRITE: read/write access (e.g. can write reports)
- ADMIN: same as WRITE + can change privileges of users (except ADMIN and SUPER)
- SUPER: same as ADMIN + can change privileges of users (except SUPER)

2.2.2. *Run Statistics*. The Run Statistics pages give users access to several parameters and statistics of each individual run. This information is grouped in 4 different tabs:

- Statistics: run conditions (e.g. magnetic field current) and data-taking metrics (e.g. event rate)
- Detectors: which sub-detectors participated in each run
- Trigger Clusters: which trigger clusters were defined for each run
- Overview: aggregated data displayed in tabular and graphical format

Users can order the runs, select the list of fields to be displayed or export the results in XML or ASCII format. Some fields have extra information made available through *onMouseOver* boxes.

2.2.3. *Run Overview*. The Run Overview page, accessible via the Run Statistics page, displays aggregated counters in a tabular and/or graphical format. This allows users to have a global view of the experiment's data-taking statistics, serving as a source for management reports, conference papers and presentations. There are 5 available tabs:

- Totals: aggregates counters globally and displays them in a table
- Per Detector: aggregates counters per sub-detector, displaying a summary table and several histograms
- Per # of Detectors: aggregates counters per number of sub-detectors, displaying a summary table and several histograms
- Per Partition: aggregates counters per partition (a logical group of 1 or more sub-detectors), displaying a summary table and several histograms
- Over Time: aggregates counters per time interval, displaying them as a function of time

The graphs and histograms are generated using the *JpGraph* PHP library, with users being able to choose between linear and logarithmic scales.

2.2.4. *Run Details*. The Run Details page provides a detailed view of the available data concerning a specific run. The following tabs are accessible:

- Run Conditions: global configuration parameters and conditions (e.g. list of participating sub-detectors)
- Run Statistics: data-taking statistics (e.g. data rates)
- Trigger Clusters Info: trigger system configuration
- Run Quality: global and per sub-detector quality flags
- LDCs Statistics: Local Data Concentrators (DAQ readout nodes) statistics
- GDCs Statistics: Global Data Collectors (DAQ event building nodes) statistics
- Shuttle Info: Offline system pre-processing status
- File Info: data migration status
- Log Entries: reports related to this run
- InfoLogger Messages: log messages created during data-taking by the different processes of the online systems

2.2.5. *Log Entries*. The Log Entries pages allow users to either read existing reports or (if they have the necessary privileges) write new reports. Any type of file can be attached to a report, with thumbnails being created for images. This information is grouped into 3 different tabs:

- Log Entries: all reports inserted by the members of the collaboration. Users can choose between ‘Compact’, ‘Extended’ or ‘Detailed’ view mode for different detail levels.
- Files: files attached to the reports displayed in the former tab
- Overview: aggregated data displayed in a tabular format

Users can optionally specify a run number to link the report to one existing run. If so, the report will be displayed in the Log Entries tab of the corresponding Run Details page.

2.2.6. *Search filters*. To allow the possibility to search for runs or log entries that match a given set of criteria, column based filters were implemented in the Run Statistics and Log Entries pages. Using either predefined or user-defined values, users can combine several filters to restrict the runs or log entries displayed. These filters are also active in the Run Overview page, allowing users to only aggregate runs that match the specified criteria.

2.3. Application Programming Interfaces

The *DAQlogbook C* API allows the different ALICE subsystems to access the DB in order to either store or read data. It is ported to other languages (like Tcl) using the SWIG tool and distributed with the DAQ DATE software package.

2.4. Command-line Interfaces

A CLI was developed to allow access, via the web-server, to specific parts of the DB for users who need to select runs matching specific criteria non-interactively. The *libcurl C* library was used to create and send the HTTP request and to process the HTTP response. In order to keep authentication, users must request beforehand a client certificate to the CERN Authentication service and provide it to the CLI as an argument.

3. Current status

Starting as a normal electronic replacement of the old paper logbooks, the eLogbook is now a central piece that the members of the collaboration access to follow and report the experiment’s activities. There are currently more than 600 registered users in the eLogbook GUI with identified accesses from 30 different countries. The DB stores more than 60 000 runs conditions and statistics and more than 5 000 human reports, with a total amount of data of 400 MB.

4. Future developments

In the near future we plan to expand the eLogbook to store data from the Data Quality Monitoring (DQM) [6] facility, the High Level Trigger (HLT) online system, the Detector Algorithms (DA) [7] facility and the different sub-detectors. This will allow users not only to have better information on the configuration of the data-taking sessions but also to access more criteria for run selection. In terms of DQM, it’s worth to mention that monitoring histograms will also be made available via the GUI, allowing remote access to this valuable information.

Concerning the DB, and with the recent migration to MySQL 5.1, we foresee to setup horizontal partitioning for better performance and simplified data management.

In terms of the GUI, apart from the display of the new data mentioned before, we plan to include a section for DAQ error analysis in order to easily identify misbehaving processes, sub-detectors or machines. In the mid-term we plan to move to the Model-View-Controller (MVC) architecture and add *Web 2.0* technologies such as AJAX and Adobe Flex.

5. Conclusions

The eLogbook facility is now widely used by the experiment's subsystems and members of the collaboration, making it an indispensable tool for the ALICE activities. Since 2007, several improvements have been done with the scope of the eLogbook growing steadily. As more subsystems enter into nominal operation mode and sub-detectors become fully operational, we can expect that a growing amount of information will become available via the eLogbook facility. With the LHC startup later this year, the eLogbook will continue to play an important role not only for data-taking operations but also as an access point where all the members of the collaboration will follow the experiment's activities.

References

- [1] ALICE 1995 Technical Proposal, CERN-LHCC-1995-71
ALICE 2008 The ALICE experiment at the CERN LHC 2008 JINST **3** S08002
- [2] Vande Vyvre P et al. Commissioning and initial experience with the ALICE on-line *Proc. CHEP 2009* (Prague, CZ: JPCS)
- [3] ALICE 2004 The Technical Design Report of the Trigger, Data-Acquisition, High Level Trigger, and Control System CERN-LHCC-2003-062
- [4] Cancio G, Duellmann D and Pace A 2009 Data Management Evolution and Strategy at CERN *Proc. CHEP 2009* (Prague, CZ: JPCS)
- [5] Ormancey E CERN Single Sign On solution *Proc. CHEP 2007* (Victoria, CA: JPCS)
- [6] von Haller B et al. The ALICE Online Monitoring *Proc. CHEP 2009* (Prague, CZ: JPCS)
- [7] Chapeland S et al. Online processing in the ALICE DAQ – The Detector Algorithms *Proc. CHEP 2009* (Prague, CZ: JPCS)