TURNING LOGBOOK DATA INTO OPERATIONAL INSIGHT WITH A NEW STATISTICS SERVICE AT GSI AND FAIR

H. Stemmler*, A. Bloch-Späth, M. Stein, O. Geithner, S. Reimann
GSI Helmholtz Centre for Heavy Ion Research, Darmstadt, Germany

Abstract

It is common practice during accelerator operation to continuously document all relevant machine states and activities. At GSI Helmholtz Centre for Heavy Ion Research (GSI), this is done using an electronic logbook developed in-house and adapted to the specific requirements of accelerator operations, allowing shift crews to record failure events, setup times, and other important information in high detail. A new service, currently under development, aims to simplify the on-demand access to relevant statistics and is being prepared to include data from the Facility of Antiproton and Ion Research (FAIR) as its operations gradually commence in the upcoming beamtimes. This service is built as two layers: a backend layer using the Micronaut framework, which communicates with the database to retrieve logbook entries and calculate statistics, and a frontend layer using the Angular framework, which handles user requests and provides graphical representations of the statistics and their raw data. This paper introduces the new statistics service as an assessment tool based on operational data at GSI and FAIR.

INTRODUCTION

The operational success of large-scale facilities like GSI Helmholtz Centre for Heavy Ion Research (GSI) [1] requires close coordination between multiple departments and technical groups. Precise and continuous documentation ensures that operational data remains accessible and reliable across all teams, enabling efficient communication between interacting areas such as beam diagnostics, physics, technical infrastructure, and control systems.

At GSI, shift crew operators use OLOG [2], an in-house electronic logbook to continuously and manually record events that may influence beam delivery or machine performance. These include failure events, machine setups, system adjustments, and unusual operational conditions. Detailed records of such events allow operators, engineers, analysts and coordinators to later revisit operational periods, evaluate performance trends, and identify recurring issues that might otherwise go unnoticed. A complementary Key Performance Indicator approach for quantitative beam-based performance assessment at GSI was recently introduced in Ref. [3]. Results are presented in Ref. [4].

Documentation becomes increasingly important as facilities expand. With the gradual development of the Facility for Antiproton and Ion Research (FAIR) [5], the complexity and scale of accelerator operations will grow significantly. This is also reflected in the FAIR Control Centre concept, developed to support the larger and more complex operational scope of the FAIR/GSI accelerator complex [6, 7].

MOTIVATION

Previous Solution for Generating Statistics

At present, operational documentation at GSI is managed through the OLOG web application (see Fig. 2), which allows operators to record events in an electronic logbook. Until its latest version, OLOG included a feature for generating statistical summaries. Users could export tables containing statistics calculated from logbook entries stored in the database. These tables provide information about events documented by the operators. However, combining statistics over longer time periods had to be done manually, introducing potential errors and misleading conclusions, as also reported at other institutes [8]. Furthermore, although the system provided a full event export, it lacked the ability to export only the data relevant to a specific statistic.

Concept for New Statistics Service

The new statistics service seeks to provide fast and convenient access to accelerator performance data through a unified, automated interface for statistical generation. Users can request specific statistics for a selected time range and immediately receive both the calculated results and the underlying data.

After a statistic is selected, the service retrieves the required entries from the OLOG database. It then processes the data and calculates the requested statistics. The results are then displayed graphically together with the raw data used in the calculation. A schematic overview of the service and its functionalities is shown in Fig. 1.

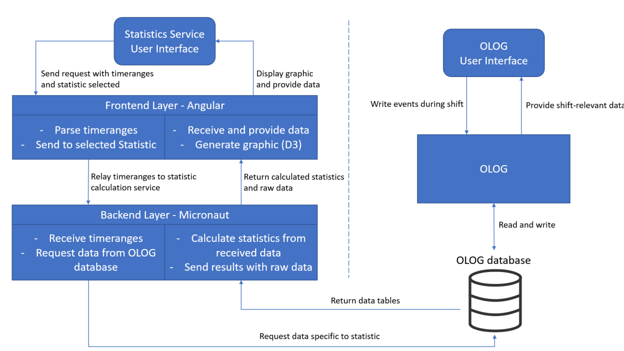

Figure 1: Schematic representation of the statistic service's layout and communication between its layers and with the OLOG database.

* h.stemmler@gsi.de

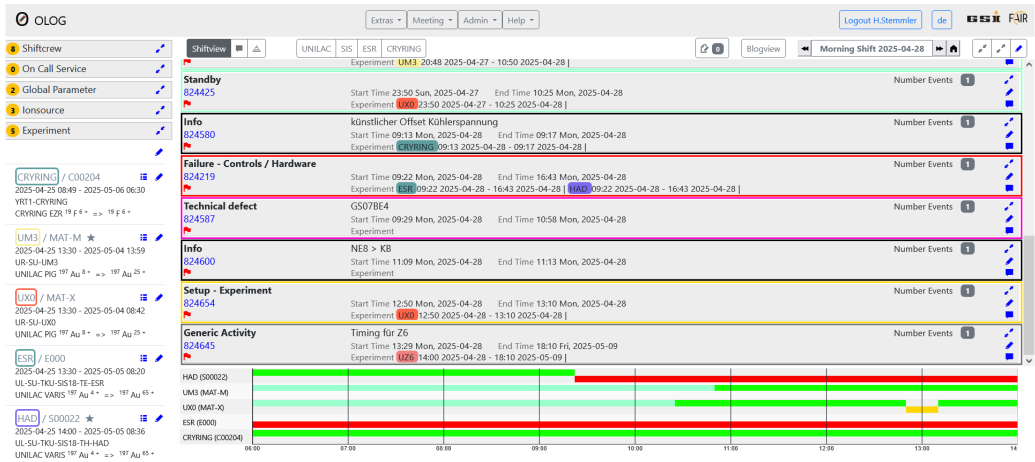

Figure 2: User interface of *OLOG*, the electronic logbook system used at GSI and FAIR to document operational data.

Providing both visualizations and raw data offers several advantages. Automatic graphical representations allow users to quickly use analyses in meetings or presentations, visualize the collected data, and identify patterns or anomalies, while access to the raw data enables more detailed analyses when needed.

Requirements In addition to ensuring accurate and reliable statistical calculations, the design of the new service is guided by three main objectives.

First, the system must be easy to use. The web interface must remain intuitive and require minimal training, enabling both technical and non-technical users to extract statistics with minimal effort.

Second, loading and computation times must remain low despite large data volumes, ensuring results are generated within seconds.

Third, the system should produce high-quality graphical outputs. Clear and customizable visualizations improve statistical interpretation and allow the results to be used in reports, presentations, or documentation.

Restrictions Despite its advantages, the new system has limitations. As the current OLOG database has been in use since 2016, earlier accelerator events are not retrievable by the service. Data from 2006 onwards is available on a different system, however the entry types do not necessarily correspond to the current OLOG data structures. This limits the historical range of analyses.

Additionally, much of the data in the OLOG system is documented manually by operators. While this allows for detailed descriptions of events, it also introduces potential inconsistencies, as naming conventions or classification methods may have changed over the years, making automated analysis more challenging. A strict standard for event documentation is therefore crucial.

IMPLEMENTATION

The new statistics service is implemented using a two-layer architecture consisting of a frontend layer using the Angular framework [9] and a backend layer based on the Micronaut framework [10]. This separation improves scalability, maintainability, and flexibility.

Technologies

Front-End Layer The frontend is implemented using the TypeScript-based framework Angular. Angular provides a robust environment for building responsive web applications and is already used in the development of many other web applications at GSI.

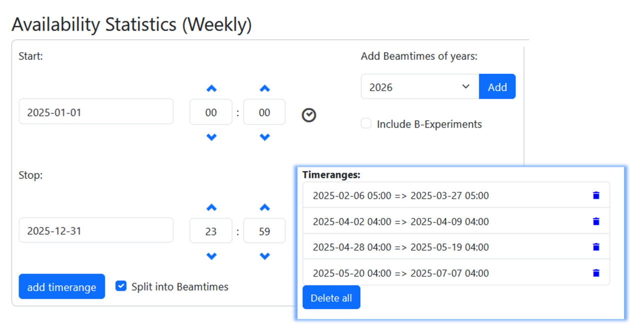

Figure 3: Static Angular component handling statistic and time range selection with the display of selected time ranges stamped in.

In the statistics service, the frontend is responsible for managing user requests. Users specify a time range on a static Angular component as shown in Fig. 3, that carries over to all pages by storing the selection in the top-level app component. Beam time periods can be automatically specified by the system, while users can select multiple custom time periods. After selecting the desired statistic component, the specified time range is sent to the backend layer by calling a service method with an HTTP GET wrapper corresponding to the statistic.

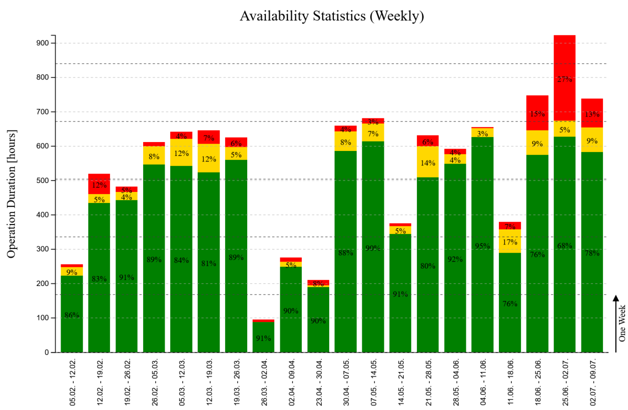

Figure 4: D3 generated SVG graphic of chosen statistic and time range as specified in Fig. 3.

Upon receiving a response from the backend, the frontend generates graphical representations with the *D3* [11] package. Figure 4 shows a graph generated by D3 for the selected statistic and time range specified in Fig. 3. Each graph is configured to display the corresponding statistic as accurate, comprehensive, and insightful as possible.

Finally, users can download each graph as an SVG file, all graphs in a ZIP file (if multiple graphs exist), and the raw data in a JSON format.

Back-End Layer The backend layer was developed as a Java microservice using the Micronaut framework, which is well suited for lightweight, high-performance services that interact with databases and external systems.

As the backend receives requests from the frontend specifying the selected statistic and the desired time range, it retrieves the corresponding logbook entries from the OLOG database. The backend then processes the data and calculates the requested statistics. These calculations include aggregating event times, handling overlapping entries, grouping and sorting entries into appropriate categories or time frames, counting event frequencies and occurrences, and analysing operational durations.

Once the calculations are complete, the backend sends both the results and the raw data within response objects specifically configured for presentation. The response objects are converted into JSON structures, with relevant statistical information at the top levels and the underlying raw data provided in deeper levels.

Data

All calculations are based on event entries manually documented in OLOG by GSI operators. Each entry typically contains information such as the time of the event, its category, a short description, fault-causing devices, affected machines and experiments, as displayed in Fig. 2. These reports are stored in a database used by the statistics service as the basis for its calculations.

Because these entries are recorded in real time during accelerator operation, they provide a detailed record of machine activity; however they are also susceptible to human error as some entries are made during hectic in-shift moments. During weekly meetings, events documented since the previous meeting are reviewed and potentially revised to better reflect the reported incident. This reduces errors and inaccuracies but does not guarantee complete correctness.

RESULTS

Initial testing shows that the statistics service performs well and fulfils its intended purpose. During development, the system helped identify inconsistencies within the database, similarly to what was experienced in Ref. [12]. For example, certain naming conventions had changed over the years, leading to variations in how events were recorded. Correcting these irregularities enhanced both the accuracy and consistency of future analyses, validating the service's usefulness as a diagnostic resource.

Another notable result is the system's performance. Despite retrieving thousands of logbook entries spanning over ten years, the service achieves loading times of only a few seconds. This gives confidence that future database growth and FAIR-related service expansion will have little impact on the system performance.

The automatically generated SVG graphs proved particularly practical, offering high-quality graphics that can be easily customised, scaled, or adapted to various formats without the loss of fidelity. These features make the tool valuable not only for internal performance tracking but also for inclusion in official reports or experimental run summaries.

CONCLUSION

The newly developed statistics service provides a powerful and efficient tool for analysing accelerator operation data at GSI. By combining automated data retrieval, statistical analysis, and graphical visualization, the system significantly improves the accessibility and usability of operational data. The service also demonstrates strong performance and scalability, enabling the analysis of large datasets with minimal waiting time.

Future plans include linking the service to the upcoming archiving system under development at GSI, which aims to automatically record and store machine data across the GSI and FAIR accelerator complexes. The inclusion of this data will further improve the quality and completeness of the operational data, thereby increasing the accuracy of the statistical analyses and expand the possible statistics. This expansion may also require the use of machine learning methods to support the analysis of increasingly large machine datasets.

Overall, the system represents an important step toward more efficient data-driven evaluation of accelerator operations and will play a valuable role in supporting both current GSI activities and future operations at FAIR.

REFERENCES

[1] GSI Helmholtz Centre for Heavy Ion Research, `https://www.gsi.de/en/start/news`

[2] P. Schuett *et al.*, "The GSI Operation Logbook OLog", presented at the 7th Int. Workshop on Accel. Oper. (WAO'10), Daejeon, Korea, Apr. 2010, https://www-conf.kek.jp/wao/WAO2010/PDF/O-17.pdf

[3] O. Geithner *et al.*, "Introduction of key performance indicators for the GSI accelerator facility", in *Proc. 16th Int. Part. Accel. Conf. (IPAC'25)*, Taipei, Taiwan, Jun. 2025, pp. 604–606. doi:10.18429/JACoW-IPAC2025-MOPS008

[4] O. Geithner *et al.*, "GSI Operation statistics in the FAIR Construction Phase (2012-2025): Trends, Failures and Lessons Learned", presented at the 17th Int. Part. Accel. Conf. (IPAC'26), Deauville, France, May 2026, paper WEP6127, this conference.

[5] S. Reimann *et al.*, "FAIR commissioning - Towards first science", in *Proc. 16th Int. Part. Accel. Conf. (IPAC'25)*, Taipei, Taiwan, Jun. 2025, pp. 2495–2498. doi:10.18429/JACoW-IPAC2025-THBN2

[6] M. Vossberg, K. Berkl, S. Reimann, P. Schuett, R. J. Steinhagen, and G. Stephan, "FAIR Control Centre (FCC) - Concepts and Interim Options for the Existing GSI Main Control Room", in *Proc. 8th Int. Part. Accel. Conf. (IPAC'17)*, Copenhagen, Denmark, May 2017, pp. 1791–1794. doi:10.18429/JACoW-IPAC2017-TUPIK047

[7] K. Grigoryev, R. Baer, S. Reimann, and V. Kamerdzhiev, "Commissioning of the FAIR Control Center", presented at the 17th Int. Part. Accel. Conf. (IPAC'26), Deauville, France, May 2026, paper WEP6100, this conference.

[8] P. Duval, H. Ehrlichmann, M. Lomperski, and J. Bobnar, "Automated Availability Statistics", in *Proc. 11th Int. Workshop Emerging Technol. Sci. Facil. Controls (PCaPAC'16)*, Campinas, Brazil, Oct. 2016, pp. 38–40. doi:10.18429/JACoW-PCAPAC2016-WEPOPRPO18

[9] Angular, https://angular.dev/

[10] Micronaut, https://micronaut.io/

[11] D3, https://d3js.org/

[12] P. Duval, R. Bacher, H. Ehrlichmann, D. Haupt, M. Lomperski, and J. Bobnar, "Automating Operation Statistics at PETRA-3", in *Proc. 16th Int. Conf. Accel. Large Exp. Phys. Control Syst. (ICALEPCS'17)*, Barcelona, Spain, Oct. 2017, pp. 876–879. doi:10.18429/JACoW-ICALEPCS2017-TUPHA189