Letter of intent for the HORIZON-INFRA-2025 call

ALLEN4EIC:

A Heterogeneous, Open-Access Framework for Real-Time High-Throughput Data Processing

Project leaders:

- Vladimir Gligorov, LPNHE, CNRS/IN2P3
- Carlos Muñoz, IJCLab, CNRS/IN2P3

(UPDATED on July 27, 2025)

Abstract

This project proposes to adapt the Allen real-time processing framework, developed at LHCb, for use at the ePIC experiment at the Electron-Ion Collider (EIC). By leveraging GPU-based and heterogeneous computing architectures, the project will enable scalable, low-latency event reconstruction and calibration tailored to ePIC's streaming readout. A core outcome is a modular, open-source software infrastructure offering Virtual Access to real-time data processing pipelines and validation tools. This service will empower researchers to test and deploy high-throughput workflows remotely, supporting FAIR data practices and advancing collaborative, reproducible research in hadron physics.

Contents

1 Research Objectives

This project will adapt the **Allen real-time processing framework**, developed by LHCb at CERN, for use at the **ePIC experiment** at the future Electron-Ion Collider (EIC). The goal is to build a modular, open-access software infrastructure for high-throughput, real-time event processing, enabling scalable workflows that will be shared with the broader research community through a **Virtual Access** model.

The ePIC detector, the flagship experiment at the Electron-Ion Collider (EIC), is being developed with a fully streaming data acquisition system, one of the first in nuclear and particle physics. This architecture enables continuous readout of the entire detector, with rapid processing supported by artificial intelligence to facilitate autonomous calibration, detector alignment, and validation of reconstructed physics events [1]. To meet the performance demands of such workflows, the collaboration is exploring the use of heterogeneous computing resources, particularly GPUs.

Allen is a proven, open-source framework that currently powers LHCb's first-level trigger, processing 30 million bunch crossings per second using GPU-based resources [2]. Its architecture supports cross-platform portability (x86, CUDA, and soon ARM), integrates Python configuration with high-performance C++/CUDA code, and includes built-in tools for performance monitoring. Unlike FPGA-based solutions, Allen achieves microsecond-scale processing using commodity hardware, making it ideal for the heterogeneous compute environments used at ePIC. Its energy efficiency and scalability have been studied in detail as part of LHCb's real-time computing evolution [3].

At ePIC, Allen's role would shift from its use at LHCb as a trigger-level filter to a downstream data processing tool, with a focus on autonomous calibration of detector systems. A first application will target the ePIC **backward electromagnetic calorimeter**, a detector currently under design and construction at IJCLab.

As part of the ODISSEE (HORIZON-INFRA-2024-TECH-01) project, the IN2P3 Allen team has received funding to modularize the generally usable components of Allen, particularly those related to code portability, high-throughput execution, and monitoring, separating them from the algorithms specific to the LHCb use case. A first version of the *AllenCore* framework is scheduled for release in Q1 of 2026. This is therefore the right time to begin planning for non-LHCb applications that can be built on top of the framework. Our project will thus extend Allen to a new experimental domain, supporting ePIC's specific data structures, geometry, and real-time requirements. A demonstrator pipeline will be deployed on a French computing cluster configured to emulate a **future EIC Echelon-2** computing site.

To implement this, we request support for one postdoctoral researcher (22 months) to:

- Integrate ePIC geometry and simulated data formats into AllenCore;
- Design a performant low-level data model informed by LHCb experience;
- Develop a real-time processing pipeline for ePIC's streaming architecture;
- Implement calibration and reconstruction algorithms;
- Extend validation and benchmarking tools to assess computational and physics performance.

The resulting infrastructure will be containerized, openly licensed, and supported by documentation and training materials to ensure wide reusability across the hadron physics community.

Scientific and Infrastructure Impact

This project will demonstrate the broader applicability of Allen as a sustainable, high-performance framework for real-time data processing in hadron physics. It will establish an open-access, containerized software suite, including real-time pipelines, validation tools, and documentation, integrated with the European Open Science Cloud (EOSC) ecosystem. Training workshops, both in-person and online, will equip researchers with skills in real-time computing, FAIR data stewardship, and EOSC-compatible workflows. In doing so, the project directly supports the Horizon call's goals of improving access to advanced research infrastructure services and preparing the next generation of researchers to exploit them effectively.

In summary, the project will deliver a sustainable, open-access processing service that integrates into the European research infrastructure landscape, enabling wide adoption of cutting-edge real-time analysis tools in hadron physics.

2 Connection to Transnational Access infrastructures (TAs) and / or Virtual Access projects (VAs)

This project directly supports and enhances the computing infrastructure of the EIC at Brookhaven National Laboratory (BNL), where the ePIC detector is a central experimental effort. By integrating the Allen framework into ePIC's data processing pipeline, we aim to strengthen BNL's capacity to manage the EIC's demanding data rates through advanced, heterogeneous computing techniques. The development of a real-time processing solution tailored for ePIC complements the existing EIC infrastructure goals by improving computational efficiency, reducing data storage demands, and enabling rapid physics analysis. Furthermore, this collaboration aligns with BNL's commitment to fostering open, scalable, and sustainable computing environments that empower the broader nuclear and particle physics communities engaged with the EIC.

A key objective of this project is the establishment of a Virtual Access facility offering open, reusable software services tailored for real-time processing in streaming experiments. All tools and documentation will be publicly released through community repositories, enabling researchers across Europe and beyond to adapt these workflows to their own needs. This supports the Horizon call's emphasis on integrated, user-ready access to high-level research infrastructure services.

3 Estimated Budget

The project requests funding to support the following key components:

- **Personnel:** One full-time **postdoctoral researcher for 22 months** (estimated CNRS contract cost of **128 196.00 €**, including an allowance for the projected increase in postdoctoral employment costs over the duration of the project), based at IJCLab within the EIC group, who will lead the development and integration of the AllenCore-based real-time processing pipeline for ePIC.
- **Travel and Collaboration:** We anticipate complementary travel support from the BNL Transnational Access program to facilitate collaboration among the participating institutions. These funds will cover project-related meetings, workshops, and joint development sessions held at BNL.

4 Participating and partner institutions

- CNRS/IN2P3: CPPM (lead: D. vom Bruch), IJCLab (lead: C. Muñoz), LPNHE (lead: V. Gligorov)
- INFN: Genova (lead: M. Battaglieri)
- CERN (lead: S. Mariani)
- Jefferson Lab (lead: M. Diefenthaler)

The project brings together a tightly integrated team with complementary expertise. LPNHE and CPPM contribute core knowledge as developers of the Allen framework at LHCb, while IJCLab leads its adaptation to ePIC, focusing on calibration of the backward electromagnetic calorimeter. INFN-Genova, experts in streaming readout (SRO), bridge the data acquisition chain up to storage, ensuring compatibility with real-time processing. Jefferson Lab, leading ePIC computing and SRO, will provide integration support and enable testing in experimental environments. CERN and the LHCb team will assist in adapting AllenCore to new use cases. This collaboration ensures a coherent development pipeline from detector to analysis, fully aligned with the goals of open, high-performance data services.

5 Integration with and Benefits to BNL, ePIC, and the EIC Infrastructure

This project delivers concrete, high-impact benefits to BNL, the ePIC collaboration, and the broader EIC infrastructure. By introducing a real-time, heterogeneous data processing framework tailored to the needs of the EIC, it supports faster detector calibration, streamlined data quality assurance, and accelerated scientific output. These capabilities will directly contribute to the success of BNL's flagship project.

5.1 Strategic Relevance to ePIC and the EIC Scientific Timeline

A central objective of Allen4EIC is to enable fast and reliable calibration and monitoring for the backward electromagnetic calorimeter in ePIC. This subsystem poses particular challenges due to its use of SiPMs, sensitivity to environmental conditions, and stringent energy resolution requirements. To meet these demands, Allen4EIC will provide:

- Low-latency, streaming-based calibration using live data;
- Online correction of environmental effects, such as temperature-dependent gain shifts;
- Real-time monitoring of detector health and stability during operation.

These features will lead to quicker commissioning, reduced downtime, and high-quality data from the start of operations. Most importantly, the ability to apply calibration and quality control in real time will accelerate the production and publication of scientific results, enabling the EIC to deliver early physics fast. By minimizing the dependence on offline calibration cycles, Allen4EIC helps shorten the time from data-taking to discovery.

5.2 Alignment with and Enhancement of BNL's EIC Computing Strategy

Allen4EIC is fully aligned with BNL's computing strategy for the EIC, which emphasizes heterogeneous architectures and real-time capabilities. The project builds on the proven Allen framework, originally developed at LHCb, adapting its GPU-accelerated, event-parallel processing model to the EIC context. Specifically, Allen4EIC will:

- Demonstrate efficient deployment of GPU-based calibration pipelines within BNL's computing environment;
- Serve as a test case for heterogeneous DAQ integration, interoperating with tools such as JANA2, EICrecon, and the streaming data acquisition systems under development at BNL.

Thanks to its open and modular design, the framework will be broadly applicable across multiple EIC subsystems, not limited to the backward calorimeter, making it easier to deploy real-time calibration and monitoring solutions across multiple subsystems and future detector upgrades.

5.3 Contribution to Shared Infrastructure and Long-Term EIC Support

Allen4EIC contributes to the long-term sustainability of the EIC computing environment by delivering tools, components, and workflows that are directly compatible with BNL-led software initiatives. The project will:

- Ensure full compatibility with EIC software standards for reconstruction, simulation, and quality assurance;
- Develop reusable modules for calibration, clustering, and monitoring, suitable for adaptation by other detector groups;
- Participate actively in integration efforts and working groups, including test beam support and collaborative development activities at BNL.

By contributing advanced European software developments to the core EIC infrastructure, Allen4EIC reinforces international collaboration while delivering tools that will benefit not only BNL and the EIC, but also the broader hadron physics community in Europe. The resulting modular and adaptable components can be reused at other experimental facilities where similar calibration and real-time processing challenges arise.

References

[1] M. Battaglieri, W. Deconinck, M. Diefenthaler, J. Huang, S. Joosten, D. Kalinkin, J. Landgraf, D. Lawrence, and T. Wenaus, "The ePIC Streaming Computing Model," 2024. [Online]. Available: https://doi.org/10.5281/zenodo.14675920

[2] R. Aaij *et al.*, "Allen: A high level trigger on GPUs for LHCb," *Comput. Softw. Big Sci.*, vol. 4, no. 1, p. 7, 2020.

[3] R. Aaij, D. H. Cámpora Pérez, T. Colombo, C. Fitzpatrick, V. V. Gligorov, A. Hennequin, N. Neufeld, N. Nolte, R. Schwemmer, and D. Vom Bruch, "Evolution of the energy efficiency of LHCb's real-time processing," *EPJ Web Conf.*, vol. 251, p. 04009, 2021.