The ePIC Streaming Computing Model

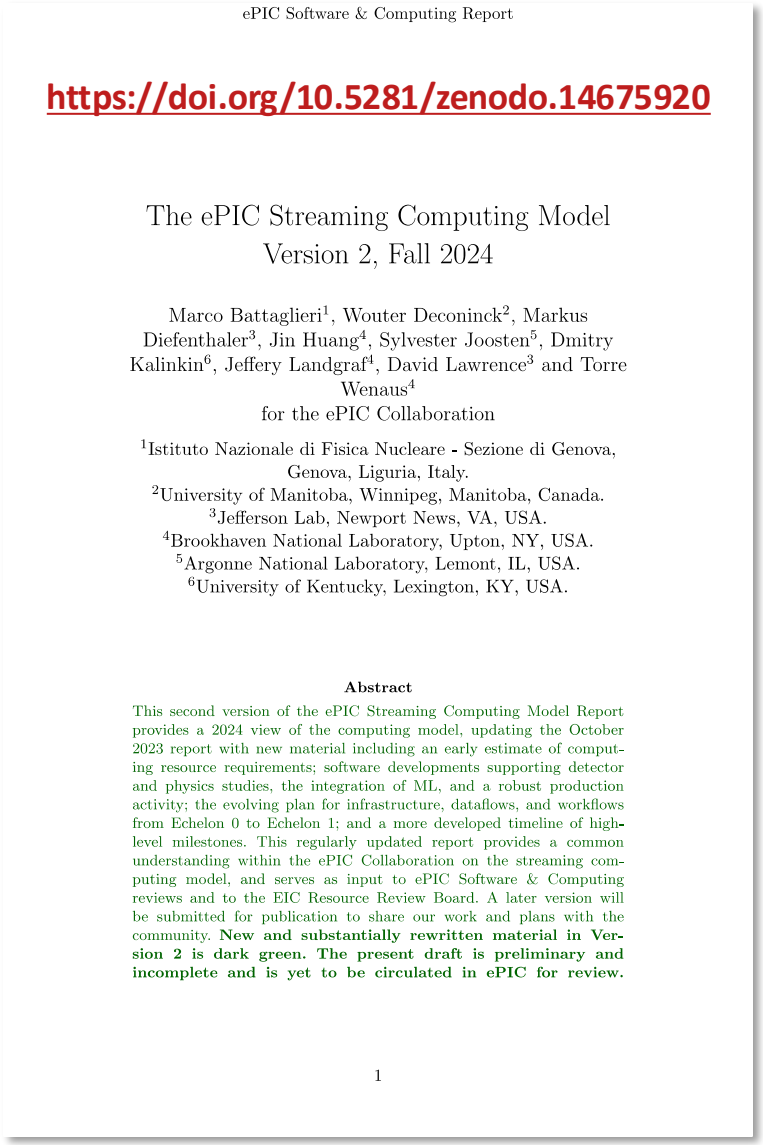

ePIC Software & Computing Report

https://doi.org/10.5281/zenodo.14675920

The ePIC Streaming Computing Model
Version 2, Fall 2024

Marco Battaglieri[1], Wouter Deconinck[2], Markus Diefenthaler[3], Jin Huang[4], Sylvester Joosten[5], Dmitry Kalinkin[6], Jeffery Landgraf[4], David Lawrence[3] and Torre Wenaus[4]
for the ePIC Collaboration

[1]Istituto Nazionale di Fisica Nucleare - Sezione di Genova, Genova, Liguria, Italy.
[2]University of Manitoba, Winnipeg, Manitoba, Canada.
[3]Jefferson Lab, Newport News, VA, USA.
[4]Brookhaven National Laboratory, Upton, NY, USA.
[5]Argonne National Laboratory, Lemont, IL, USA.
[6]University of Kentucky, Lexington, KY, USA.

Abstract

This second version of the ePIC Streaming Computing Model Report provides a 2024 view of the computing model, updating the October 2023 report with new material including an early estimate of computing resource requirements; software developments supporting detector and physics studies, the integration of ML, and a robust production activity; the evolving plan for infrastructure, dataflows, and workflows from Echelon 0 to Echelon 1; and a more developed timeline of high-level milestones. This regularly updated report provides a common understanding within the ePIC Collaboration on the streaming computing model, and serves as input to ePIC Software & Computing reviews and to the EIC Resource Review Board. A later version will be submitted for publication to share our work and plans with the community. **New and substantially rewritten material in Version 2 is dark green. The present draft is preliminary and incomplete and is yet to be circulated in ePIC for review.**

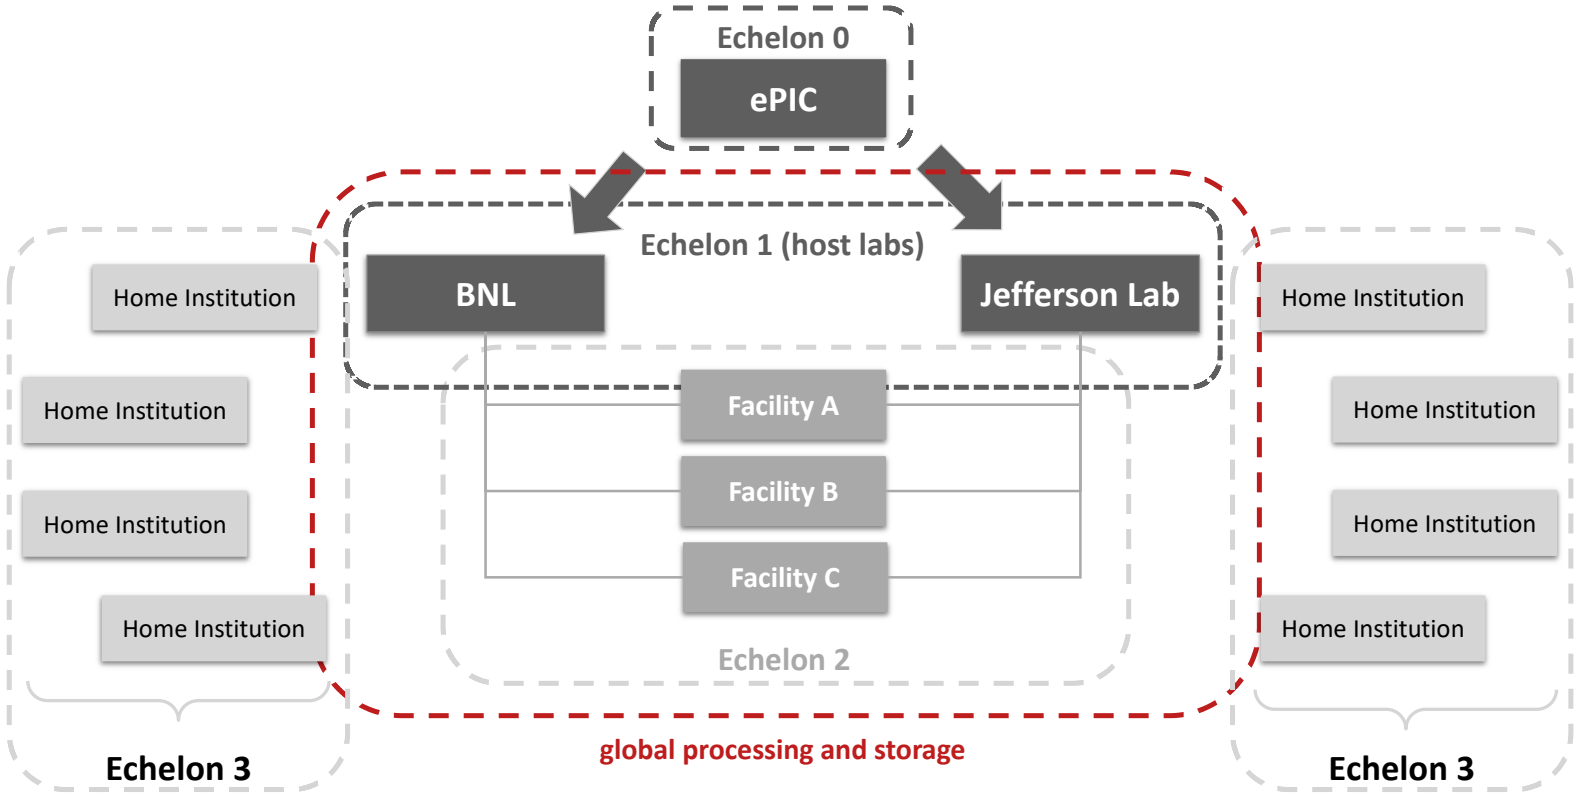

We developed the ePIC Streaming Computing Model to accelerate the pace of discovery and enhance scientific precision through improved management of systematic uncertainties. The model is documented in a detailed report and was reviewed by ESCAC in 2023 and 2024.

Computing Resource Needs (2034) and Their Implications

Processing by Use Case [cores]	Echelon 1	Echelon 2
Streaming Data Storage and Monitoring	-	-
Alignment and Calibration	6,004	6,004
Prompt Reconstruction	60,037	-
First Full Reconstruction	72,045	48,030
Reprocessing	144,089	216,134
Simulation	123,326	369,979
Total estimate processing	**405,501**	**640,147**

Storage Estimates by Use Case [PB]	Echelon 1	Echelon 2
Streaming Data Storage and Monitoring	71	35
Alignment and Calibration	1.8	1.8
Prompt Reconstruction	4.4	-
First Full Reconstruction	8.9	3.0
Reprocessing	9	9
Simulation	107	107
Total estimate storage	**201**	**156**

O(1M) core-years to process a year of data:

- Even with performance gains over the years, the required processing scale remains substantial.
- Highlights the need to leverage distributed and opportunistic resources from the outset.

~350 PB to store data of one year.

ePIC is a compute-intensive experiment. Its science must not be limited by computing constraints.

Computing Use Cases and Their Echelon Distribution

Use Case	Echelon 0	Echelon 1	Echelon 2	Echelon 3
Streaming Data Storage and Monitoring	✓	✓		
Alignment and Calibration		✓	✓	
Prompt Reconstruction		✓		
First Full Reconstruction		✓	✓	
Reprocessing		✓	✓	
Simulation		✓	✓	
Physics Analysis		✓	✓	✓
AI Modeling and Digital Twin		✓	✓	

Substantial role for Echelon 2 in preliminary resource requirements model

Assumed Fraction of Use Case Done Outside Echelon 1	
Alignment and Calibration	50%
First Full Reconstruction	40%
Reprocessing	60%
Simulation	75%

- **Echelon 1** sites uniquely perform the **low-latency streaming workflows** consuming the data stream from Echelon 0:
  - Archiving and monitoring of the streaming data, prompt reconstruction and rapid diagnostics.
- Apart from low-latency, **Echelon 2** sites fully participate in use cases and **accelerate** them:
  - Tentative resource requirements model assumes a **substantial role for Echelon 2**.
  - **Priority**: Capabilities and resource requirements for Echelon 2 resources developed jointly with the community.
  - **Priority**: Establishing EIC International Computing Organization (EICO):

Comment on Computing Modeling

- **Modeling** is fundamental and integral to science. My own experience is with physics and detector simulations.
- It is essential to understand the **accuracy**, **precision**, and, ideally, the **predictive power** of the tools and approach being used. Once this is established, modeling can be applied to **address concrete questions**.

- I listened with interest to your presentation on *"**Modeling the Data Transfer Between Echelon 0 and Echelon 1**."*
- In that case, the accuracy, precision, and predictive power of the tools and approach were not demonstrated; therefore, I cannot assess what was gained from the modeling effort.

- Alexei asked me what would be useful in terms of modeling. I have a **concrete proposal**:

  - To further develop the **ePIC Computing Model**, we need a better understanding of our use cases and how they can be distributed across Echelon 1 and 2 resources.
  - **Modeling** could play a vital role in achieving this by a) helping to define the use cases in more detail, and b) providing guidance for the requirements discussion for the Echelon 1–2 sites.
  - To enable tangible progress, I propose **limiting the scope of modeling to data flow**, i.e, distributed data management using Rucio / file transfer. This also aligns with the recommendation from the ECSAC review to begin developing a data management and lifecycle plan.

- However, we must **ensure that the modeling approach and tools are appropriate for the task**. This must be demonstrated.