



ePIC Streaming Computing Model

The background of the slide is a dynamic, abstract composition. It features a central point from which numerous bright, multi-colored light rays (red, orange, yellow, and blue) radiate outwards, creating a sense of depth and speed. Overlaid on this are streams of binary code (0s and 1s) in various colors, appearing to flow and converge towards the same central point. The overall effect is one of high-tech, digital energy.

Markus Diefenthaler (Jefferson Lab)


ePIC in a Nutshell



Captured from the [ePIC Event Display](#)

- **ePIC is the primary experiment at the Electron-Ion Collider at Brookhaven National Laboratory**, a project jointly realized by Brookhaven National Laboratory and Jefferson Lab.
- It will use electron-ion collisions to **image the quarks and gluons inside protons and nuclei** and study how the strong force shapes nuclear matter.
- ePIC is a **highly integrated, multipurpose detector** designed to reconstruct a wide range of particle signatures.
- The **international ePIC Collaboration** includes more than 1,100 collaborators from over 200 institutions in 25 countries.
- This scale, together with the detector's data rates and globally distributed collaboration, drives the need for a **distributed computing model**.

The ePIC Streaming Computing Model



Successfully reviewed in 2023–2025 and documented

The ePIC Streaming Computing Model
Version 2, Fall 2024

Marco Battaglieri¹, Wouter Deconinck², Markus Diefenthaler³, Jin Huang⁴, Sylvester Joosten⁵, Dmitry Kalinkin⁶, Jeffery Landgraf⁴, David Lawrence³ and Torre Wenaus⁴
for the ePIC Collaboration

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

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

The ePIC Streaming Computing Model is **physics-driven**. It has been developed to **accelerate the pace of discovery** and **enhance scientific precision** through streaming readout, AI, and distributed computing.




What the ePIC Streaming Computing Model Enables

The ePIC Streaming Computing Model integrates streaming readout, AI, and distributed computing.


- 1 Streaming Readout**
Maximizes scientific reach through unbiased event selection and continuous measurement of backgrounds and detector noise.
- 2 AI**
Accelerates science through autonomous alignment, calibration, and validation, shortening the path to physics results and publications.
- 3 Distributed Computing**
Enables flexible access to dedicated and opportunistic computing resources where they are needed.

Closed-loop experimental control



Investing now in the ePIC Streaming Computing Model ensures that the collaboration will have the infrastructure, resources, and workflows needed to **commission and operate the experiment**.

Streaming DAQ and Computing Milestones



- **Compute-Detector Integration:**

- Joint deliverables between **DAQ** and **computing** to develop an integrated system.
- We define and test DAQ-computing interfaces to reduce integration risk and enable early system testing.

- **Integrating AI with ePIC Detector:**

- AI has been built into the ePIC computing model from the outset.
- We are developing AI-driven workflows for alignment, calibration, validation, control, and diagnostics.
- This enables seamless processing from detector readout to physics analysis.

Computing Resource Needs (2034) and Their Implications

Based on a luminosity of $L = 1 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1} = 1 \text{ kHz}/\mu\text{b}$

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.

Computing resource needs at a scale of ATLAS and CMS today.

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

Distributed Computing for ePIC

- **Echelon 1** sites uniquely perform the **low-latency streaming workflows**.
- Apart from low-latency, **Echelon 2** sites fully participate in use cases and **accelerate** them.

Strong international interest and engagement:



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%

Distributed computing provides **flexibility**, influencing the fraction of use cases handled outside Echelon 1.

ECSJI Role and Echelon 1 Status

The **EIC Computing and Software Joint Institute (ECSJI)**, established by BNL and JLab, is the host laboratories' joint organization for EIC computing and software. It:

- provides a single interface to ePIC and the EIC Project, as well as DOE and international computing partners;
- coordinates responsibilities, resources, services, and personnel at the two host laboratories;
- manages agreements, service expectations, reviews, and resource planning;
- supports the development and operation of the distributed EIC computing infrastructure.



Alexei Klimentov (BNL) and David Lawrence (JLab) are the ECSJI directors.


Echelon 1 comprises the computing facilities operated by BNL and Jlab:

- Early deployment should begin in 2027.
- Initial resources will support detector studies, simulations, streaming tests, and joint Echelon 0 to Echelon 1 data challenges.

EIC International Computing Organization

Discussion Focus

- International partner engagement:



- **ePIC Streaming Computing Model:**
 - Coordination of data challenges with compute and data centers.
 - Echelon 2 roles and participation.
 - **Networking:** ESnet and LHCONE
- **WLCG:** Technical exchange and associated membership path.



ECSJI and EICO Annual Meeting (May 18–19, 2026, CERN)

ePIC is working closely with ECSJI and EICO to define international participation in the ePIC Streaming Computing Model and advance its implementation.

Use-Case-Driven Echelon 2 Participation

Participation being defined by use-case-specific capabilities, not by fixed site categories:

Capability Envelope

- Describes the **smallest set of resources and operational capabilities required** for a site to participate for a given use case.
- It is not a single number, but a combination of **requirements across multiple dimensions**:
compute · storage · network throughput · I/O performance · operational reliability
- A site **must satisfy all dimensions simultaneously**; otherwise, it cannot effectively execute the workload.
- Envelope will be defined pragmatically as an appropriate balance between added value and operational load.

Approach separates two distinct questions:

- **Qualification**: whether a site has sufficient capabilities to participate in a given use case (defined by the envelope).
- **Contribution**: what fraction of the global workload the site takes on (expressed as a percentage and coordinated across sites to match total resource needs).

Capability envelopes keep the computing model realistic and use-case driven, while allowing Echelon 2 contributions to meet global compute and storage targets.

Echelon 2 Resource Needs

Based on a luminosity of $L = 1 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1} = 1 \text{ kHz}/\mu\text{b}$

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.
- Distributed computing is an integral part of our computing model.


~350 PB to store data of one year.

Echelon 2 qualification may require a **defined fraction of the estimated resource needs**, e.g., 10% to 20% of the values shown, with the appropriate threshold under discussion within ePIC and EICO.


Distributed Computing and Storage: Status

Monthly simulation campaigns for preTDR and CD-2 preparations:

Compute Capacity



Storage Capacity



ePIC is building its distributed compute and storage fabric now, with ongoing discussions about additional international contributions.

Flags: France, Japan, United Kingdom

Distributed Computing Interfaces: **PanDA**, **Rucio**, and **XRootD**

- **PanDA integrates with existing site infrastructure:** Sites can keep their local batch systems and operational model rather than adopting a new architecture.
- **Central workload management, local resource control:** PanDA coordinates, dispatches, and monitors workloads across distributed resources, while the site retains control over local resource scheduling, configuration, allocation and operational policies.
- **Common interfaces are key to scaling across sites:** Authentication, authorization, job routing, software access, and monitoring should be standardized as much as possible across the collaboration. The initial model will be agreed upon between BNL and JLab, validated at both sites, and then rolled out to additional sites.
- **Monitoring should combine collaboration-wide and site-local views:** PanDA provides global workflow visibility, while sites retain their local monitoring tools.
- **Rucio provides the common data-management layer** across participating sites.
- **Data access and transfer:**
 - **XRootD:** unauthenticated read; authenticated read/write
 - **HTTPS:** third-party copy between participating sites
- **Authentication is moving to tokens** over the next few months:
 - Token-based authentication for storage write/TPC and Rucio production access
 - X.509 authentication is expected to be retired once token deployment is complete.

Coordinated Operations Across Computing Sites

- **Federation across independent sites requires coordination:** Experience from large distributed HENP computing shows that common services become difficult to operate reliably when configurations, deployments, and monitoring diverge across sites.
- **Collaboration-managed services and configuration:** ePIC aims to centrally maintain common configurations, deployments, and monitoring rather than independently configuring these components at each site.
- **Close partnership with hosting sites:** ePIC will work with each site on the underlying infrastructure and on integrating local resources into the common operational model.
- **Common operational visibility:** Collaboration-wide monitoring should provide a consistent view across sites while complementing site-local operational tools.
- **Preparing for AI-assisted operations:** We would like to work with sites toward:
 - AI-compatible logging and monitoring,
 - Appropriate access controls,
 - Well-defined interfaces for AI-based operational tools.

Summary

ePIC will require a **global distributed computing fabric** at substantial compute and storage capacity.

International sites are integral to the model, especially for reconstruction, reprocessing, simulation, analysis, and AI workloads.

Participation will be **use-case driven**, based on compute, storage, network, I/O, and operational capabilities.

ePIC will provide common workload-management, data-management, and operational interfaces.

We want to work with sites **now** on integration, data challenges, qualification criteria, and future AI-assisted operations.

Investing now in the ePIC Streaming Computing Model ensures that the collaboration will have the infrastructure, resources, and workflows needed to **commission and operate the experiment**.