


# Software & Computing Report



**Markus Diefenthaler (Jefferson Lab)**

# ePIC Software & Computing Organization



## Guiding Principles:

- Diversity, Equity, and Inclusion
- Statement of Software Principles
- Sustainability



**Software and Computing Coordinator**  
Markus Diefenthaler (Jefferson Lab)

## Cross-cutting Working Group:

- Data and Analysis Preservation



**Deputy Coordinator (Operations)**  
Wouter Deconinck (U. Manitoba)

## Operation Working Groups:

- Production
- User Learning
- Validation



**Deputy Coordinator (Development)**  
Dmitry Kalinkin (U Kentucky)

## Development Working Groups:

- Physics and Detector Simulation
- Reconstruction Framework and Algorithms
- Analysis Tools (not yet activated)



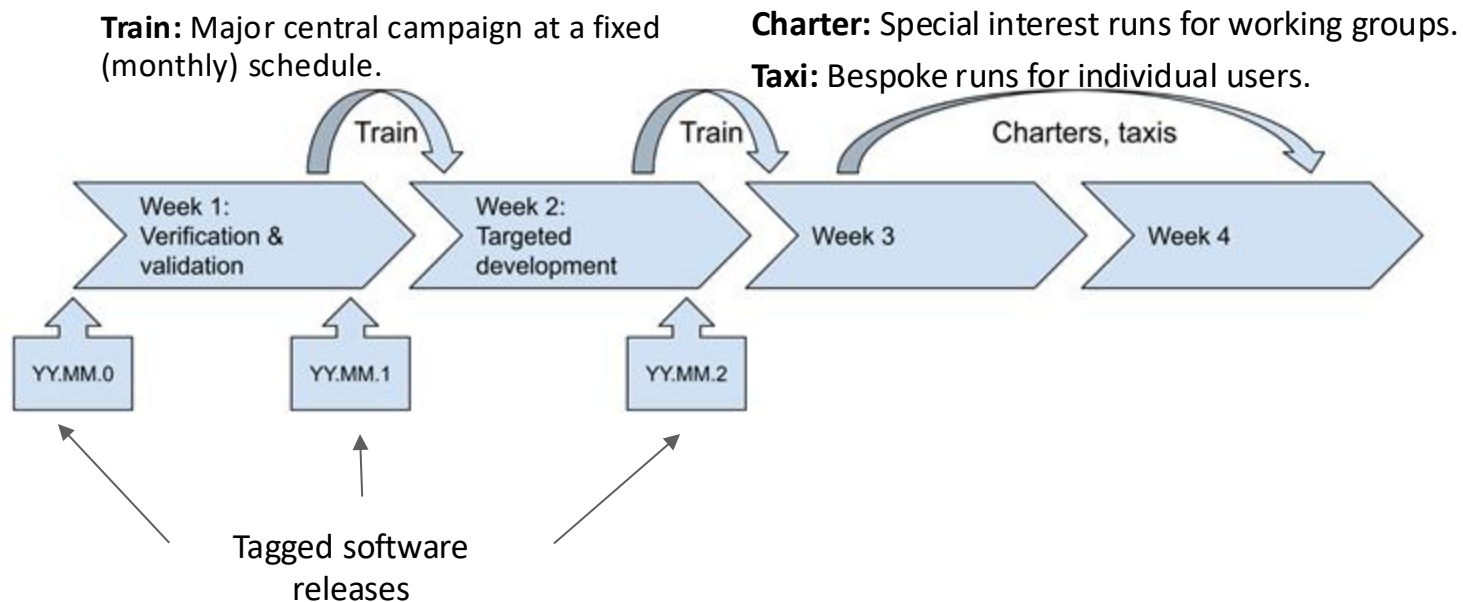
**Deputy Coordinator (Infrastructure)**  
Torre Wenaus (BNL)

## Infrastructure Working Groups:

- Streaming Computing Model
- Multi-Architecture Computing (not yet activated)
- Distributed Computing (not yet activated)

# Monthly Simulation Productions

1. **Regular updates** of simulation productions for detector and physics studies in preparation for the **preTDR**, subsequent CD milestones, and development of **early science program**.
2. **Timely validation and quality control** for simulation productions on datasets that require substantial time and resources. Focus on **benchmarks driven by Continuous Integration / Continuous Deployment**, a process that automates the testing and building of software.
3. **Continuous Deployment** of the software used for detector and physics simulations.
  - Latest merged features available the next day in nightly containers.



**Broad science program for the EIC:** The selection of physics processes and associated MC simulations for the preTDR has been finalized with the PWGs; however, it will continue to evolve over the coming years.

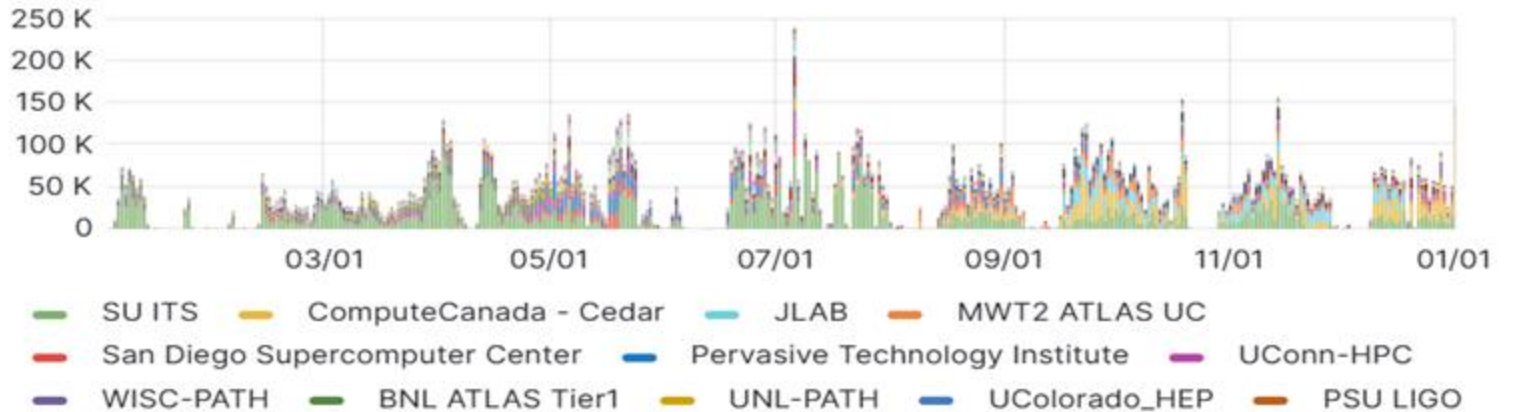
We can turn around new geometry and new algorithms within a month on the millions of events necessary to assess the impact. In 2024, monthly simulation campaigns have used **15 M core hours on the Open Science Grid** and produced over **~500 TB of simulation data**.

# Production WG (Sakib Rahman, Thomas Britton)

## Infrastructure improvements

- Improved throughput from Jefferson Lab
- Lustre disks made available at Jefferson Lab to improve performance
- Set up a cloud RSE for logging
- **Rucio instance** at Jefferson Lab
- Initiated and started testing new BNL submit host
- Added 200TB of XRootD storage at BNL with token-based read-write access
- **Successful integration of international resources:**  
Digital Research Alliance of Canada and INFN

Core Hours By Facility



## Rucio in Production

- Successful test jobs running on OSG and writing back to XRootD storage.
- January, 2025 campaign expected to be the first campaign with Rucio.

## Simulation Production Links

- [Live campaign updates](#)
- [Input Preprocessing Policy](#)
- [Default List of Datasets](#)
- [File access instructions](#)
- **NEW** [Dataset Request Form](#)

# Validation WG (Torri Jeske and Dmitry Kalinkin)

**Continuous Integration (CI):** Automated benchmarks are run with each code integration, ensuring that the code adheres to quality standards and functions as intended. These benchmarks provide immediate feedback to developers about the impact of their changes, enabling them to address issues promptly.

**Approach for CI:** capbara boosted code review productivity; successful adoption of Snakemake-based workflows.

**Data Quality Plots:** A standardized set of plots assists in visualizing and understanding the integrity of the data, identifying missing values or unexpected results. These plots facilitate the tracking of data quality changes over time, helping to pinpoint trends or shifts that may necessitate intervention. Additionally, making these plots accessible to everyone in the collaboration enhances communication and engagement.


**Approach for Data Quality:** Image browser with automated upload of CI detector benchmarks; underlying database is currently managing 250k+ images.

We benefit from excellent Continuous Integration workflows in the development of our software, and we are using the **TDR efforts to establish a standardized set of detector and physics performance plots.**

Here, the **collaboration-wide commitment** to ensuring that the **plots from the TDR are reproducible** makes a key difference.



# Physics and Detector Simulations WG (Chao Peng and Sakib Rahman)

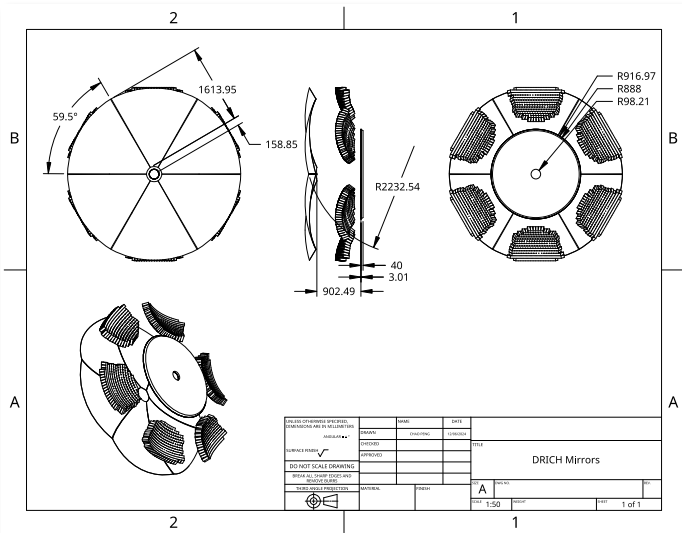


## Key achievements and milestones:

- Geometry update and detailed DD4hep implementation for all sub-detectors - in agreement with geometry database.
- Simulation to CAD model (step file) conversion.
- Begin of the evaluation process of simulation and engineering designs.

## Major impacts and successes:

- Maintenance of the interface for root geometry to step file conversion (an [npsim](#) tool).
- Discussions on the progress of CAD model to simulation conversion.
- Supports on the digitization improvements (simulation of detector readout electronics).
- Background merging revisited; simulation productions with and without background will be available in upcoming campaigns.



# Reconstruction Framework and Algorithms WG (Derek Anderson and Shujie Li)


## Key achievements and milestones:

- **Year One:**
  - Developed a full chain of ACTS-based track reconstruction with auto material map generation, realistic seed finder, combinatorial track finding/fitting, ambiguity solver, and primary vertexing.
  - Implemented initial prototype of electron-finder, which includes the implementation of track-calorimeter projections.
  - Made progress towards the implementation of particle flow in ePIC.
  - Made significant improvements across EICrecon, e.g., jet software.
- **Year Two:**
  - Led highly successful workfest at EIC UGM: provided critical input for priorities going into FY25.
  - Continued progress towards particle flow in ePIC (e.g. cluster merging/splitting)
  - Continued making significant improvements across EICrecon, e.g., refinement of truth-cluster associations.
  - Made progress towards implementation of secondary vertexing.

## Major impacts and successes:

- **Year One:**
  - Enabled tracking performance study with realistic event and beam background.
  - Supported internal review of tracking detectors
- **Year Two:**
  - Excellent progress on addressing pTDR priorities, both on Calorimeter and Tracking fronts.

# Visualization via Firebird: Web-Based Event Display Based on Phoenix



**Event Display for Reconstruction Development:**

**Available in eic-shell and at:**

**<https://eic.github.io/firebird/display>**

**Features:**

- Supports ePIC geometries.
- Integrates with simulation output for comprehensive studies and debugging.



# User Learning WG (**Stephen Kay** and Holly Szumila-Vance)

Landing Page on <https://eic.github.io/documentation/landingpage.html>:

- Full toolkit to get new users started: “**Getting Started**”, **software tutorials**, and **FAQs**.
- Lots of new documentation accessible to users in mostly uniform format.
- Any member of the collaboration can **directly contribute by submitting change requests**.
- Goal to fully integrate with new ePIC website.

Get started

ePIC Tutorials

HEP Software  
Training Center

FAQ

Welcome to the **ePIC Landing Page!**

Our mailing list: ✉ [eic-projdet-comp-sw-l@lists.bnl.gov](mailto:eic-projdet-comp-sw-l@lists.bnl.gov)

Subscribe here: <https://lists.bnl.gov/mailman/listinfo/eic-projdet-comp-sw-l>

**HelpDesk on Mattermost** consistently provides software support.

**New initiative for 2025** – Software manuals.

**Rolling programme of tutorials:**

**01/30** Dmitry Romanov – ePIC Event Display

**TBA** Shujie Li – Understanding the Simulation output”

**02/28** Stephen Kay – "Getting started with ePIC analysis”

# Streaming Computing (WG) (Marco Battaglieri, Jin Huang, **Jeff Landgraf**)

- July 11, 2023 – Kickoff Meeting, [detailed meeting notes](#) available since late 2024
  - WG essential in preparation for the streaming model computing reviews:
    - Oct. 2023 – ePIC Software and Computing Review
    - Dec. 2023 – EIC Resource Review Board
    - Sept. 2024 – ePIC Software and Computing Review
    - Nov. 2024 – EIC Resource Review Board
- 
- Map detector calibration logic, and dependencies
  - Collaboration with Electronics & DAQ WG towards defining interface, & handling of time frames
  - Iterating SRO data handling
    - Super Time Frames ~1000 Time Frames, non-overlapping, time ordered, unit of storage / data management
    - Discussing orchestration and storage e.g. Rucio / object storage tools
    - Continued discussion of infrastructure e.g. JANA2 updates for SRO
  - Opened discussions regarding requirements and plans
    - Housing part of Echelon 0 computing at SDCC
    - Network to building 1006 ~72 existing fibers shared with 1008, 1005 (need O(10) fibers)
    - Requirements for connection to Jefferson Lab
  - Strategy for demonstrating streaming reconstruction
    - TDR – separate analysis of streaming tracking based on “merged” timeframes
    - Post TDR – incorporate streaming digitization into simulation
  - SRO reconstruction from integrated time frames
    - Background mixer by Kolja Kauder
    - JANA2 updates for processing of timeframes, events and subevents by Nathan Brei
    - Demonstration that hits in a timeframe can be correlated to events in a timeframe by Barak Schmookler
    - Next steps are analysis with background/noise “blind” to event structure, & automated classification of background/physics
  - Integrating DAQ project milestones with SRO plans

# Compute-Detector Integration to Maximize and Accelerate Science

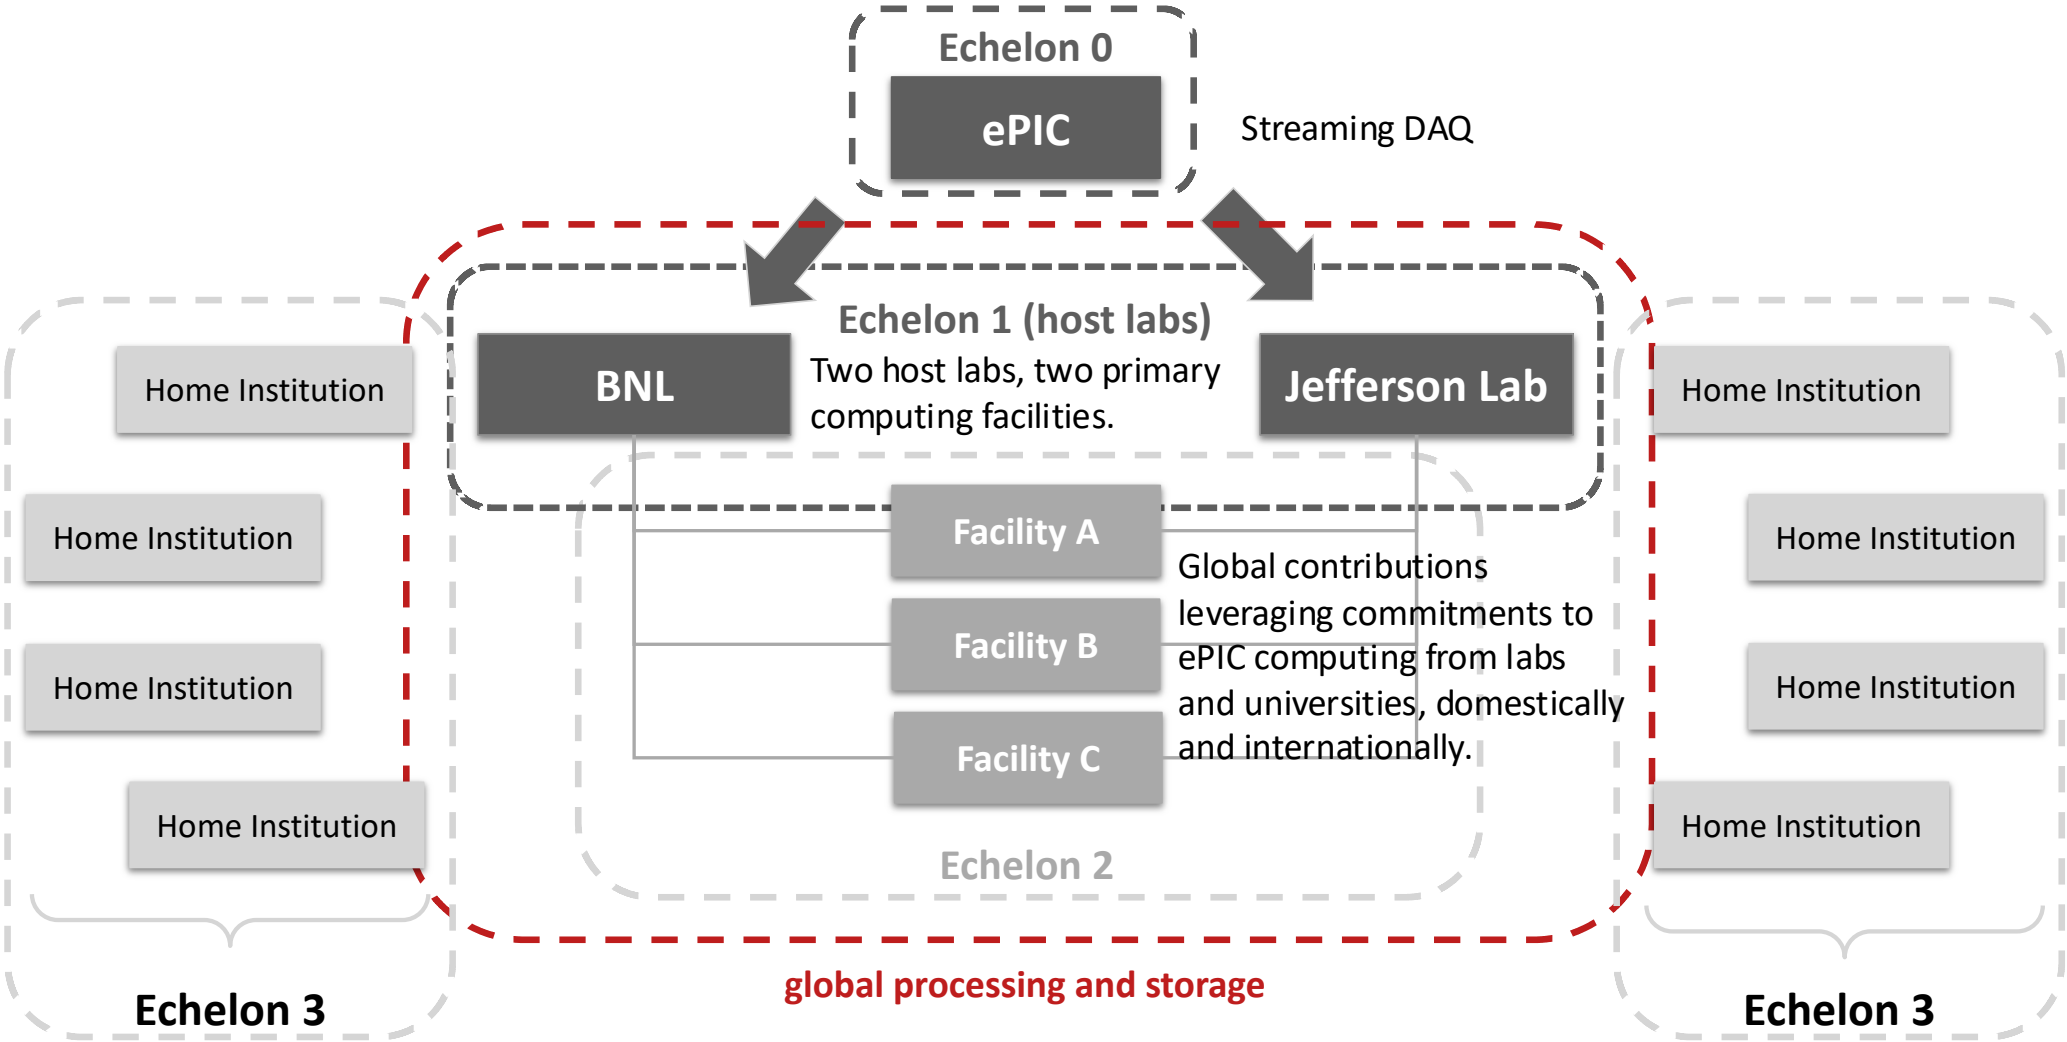
- **Maximize Science** Capture every collision signal, including background.
  - Event selection using all available detector data for **holistic reconstruction**:
    - **Eliminate trigger bias** and provide accurate estimation of uncertainties during event selection.
  - Streaming background estimates ideal to **reduce background** and related systematic uncertainties.
- **Accelerate Science** Rapid turnaround of 2-3 weeks for data for physics analyses.
  - Timeline driven by alignment and calibration.
  - Preliminary information from DSCs indicates that 2-3 weeks are realistic.
- **Technologies** Compute-detector integration using:

**Streaming readout** for continuous data flow of the full detector information.

**AI** for autonomous alignment and calibration as well as autonomous validation for rapid processing.

**Heterogeneous computing** for acceleration (CPU, GPU).

# The ePIC Streaming Computing Model



Supporting the analysis community *where they are* at their home institutes, primarily via services hosted at Echelon 1 and 2.









# 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 tentative 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 **accelerates** them:
  - Tentative resource requirements model assumes a **substantial role for Echelon 2**.
  - Capabilities and resource requirements for Echelon 2 resources developed jointly with the community.
  - Forming EIC International Computing Organization (EICO):      
  - The power of distributed computing lies in its flexibility to shift processing between facilities as needed.





# Computing Resource Needs (EIC Phase I) 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
<b>Total estimate processing</b>	<b>405,501</b>	<b>640,147</b>

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
<b>Total estimate storage</b>	<b>201</b>	<b>156</b>

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

- Optimistic scaling of constant-dollar performance gains would reduce the numbers about 5x:
  - Based on current WLCG measure of 15% per year.
  - But the trend is towards lower gains per year.
- Whatever the gains over time, processing scale is substantial!
- Motivates attention to leveraging distributed and opportunistic resources from the beginning.

~350 PB to store data of one year.

Computing resource needs comparable to LHC experiments ATLAS and CMS.

**ePIC is compute intensive experiment; must ensure ePIC is not compute-limited in its science.**

# Streaming DAQ and Computing Milestones (**work in progress**)

## Streaming DAQ Release Schedule:

### PicoDAQ

**FY26Q1**

- Readout test setups

### MicroDAQ:

**FY26Q4**

- Readout detector data in test stand using engineering articles

### MiniDAQ:

**FY28Q1**

- Readout detector data using full hardware and timing chain

### Full DAQ-v1:

**FY29Q2**

- Full functionality DAQ ready for full system integration & testing

### Production DAQ:

**FY31Q3**

- Ready for cosmics

## Streaming Computing Milestones:

**Start development of streaming orchestration**, including workflow and workload management system tool.

**Start streaming and processing streamed data between BNL, Jefferson, DRAC Canada, and other sites.**

**Support of test-beam measurements, using variety of electronics and DAQ setups:**

- Digitization developments will allow detailed comparisons between simulations and test-beam data.
- Track progress of the alignment and calibration software developed for detector prototypes.
- Various JANA2 plugins for reading test-beam data required. Work started on an example.

**Establish autonomous alignment and calibration workflows that allows for validation by experts.**

**Analysis challenges exercising end-to-end workflows** from (simulated) raw data.

**Streaming challenges** exercising the streaming workflows from DAQ through offline reconstruction, and the Echelon 0 and Echelon 1 computing and connectivity.

**Analysis challenges exercising autonomous alignment and calibrations.**

**Data challenges exercising scaling and capability tests** as distributed ePIC computing resources at substantial scale reach the floor, including exercising the functional roles of the Echelon tiers, particularly Echelon 2, the globally distributed resources essential to meeting computing requirements of ePIC.

# ePIC Computing & Software Reviews

## EIC Computing and Software Advisory Committee (ECSAC)

advise host laboratories on the progress and status of computing and software for the ePIC collaboration.



**Mohammad Al-Turany (GSI)**



**Simone Campana (CERN)**



**Christoph Pauss (MIT)**



**Verena Martinez Outschoorn (UMass Amherst)**



**Frank Würthwein (chair, UCSD)**

- **Annual reviews** with a charge reflective of the EIC schedule, the stage of the ePIC experiment, and impending deadlines.
- **Recent review** on September 26–27, 2024 included: **Assessment of the ePIC Computing Model.**

<b>Is there a comprehensive and cost-effective short and long-term plan for the software and computing of the experiment?</b>	<b>Yes</b>
• The pre detector technical design report (TDR) is scheduled to be delivered in 2025. Are the resources for software and computing sufficient to deliver the TDR?	<b>Yes</b>
• Is the design of the ePIC computing model and resource needs assessment adequate for this stage of the project?	<b>Yes</b>
• Is the ePIC computing and model flexible? Can it evolve and integrate new technologies in software and computing?	<b>Yes</b>
<b>Are the plans for software and computing consistent and integrated with standard practices across nuclear physics and particle physics communities, especially given technical evolution over the next decade?</b>	<b>Yes</b>
<b>Are the ECSJI plans to integrate into the software and computing plans of the experiment sufficient?</b>	<b>Yes</b>
<b>Are the plans for the integrating international partners' contributions flexible and adequate at this stage of the project?</b>	<b>Yes.</b>

### **Two Recommendation, to the Host Labs and ePIC:**

- Provide a detailed plan and timeline before the next ECSAC meeting for creating dedicated effort to ePIC Software & Computing team.
- Investigate how U.S. universities can contribute to the software and computing needs of the experiment, and present a plan at the next ECSAC review.

# Strategic Resource Allocation Plan

**ECSAC:** *“If you were given 4 dedicated experts, make a prioritized list of the tasks/areas where you would employ them for your short/medium/long term needs.”*

## Short/Medium-Term (next 3 years)

- Establish a dedicated effort in collaboration with Electronics & DAQ to develop integrated DAQ-computing workflows, working towards a full streaming DAQ chain test.
- Holistic full PID full reconstruction (lepton-hadron separation, lepton ID, hadron ID) implementation in the ePIC software stack utilizing the full capabilities of the integrated detector (PID, calo, tracking, etc.).
- Support AI/ML workflow integration in full simulation and reconstruction algorithms.
- ACTS expert for track seeding, track fitting, vertex finding algorithm development, tuning, and evaluation.

## Long-Term (4+ years)

- Continued support for streaming DAQ workflows in collaboration with Electronics & DAQ
- Expert in fast simulations to reduce the computational cost of the simulation campaigns to interpret data
- Expert in hardware accelerators to develop collaboration expertise to speed up simulation and reconstruction and leverage HPC platforms
- Distributed computing expert to develop operations between Echelon-1/2 and support progressively scaled up challenges



# Getting Involved

ePIC Software & Computing Report

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

The ePIC Streaming Computing Model  
Version 2, Fall 2024

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

<sup>1</sup>Istituto Nazionale di Fisica Nucleare - Sezione di Genova, Genova, Liguria, Italy.

<sup>2</sup>University of Manitoba, Winnipeg, Manitoba, Canada.

<sup>3</sup>Jefferson Lab, Newport News, VA, USA.

<sup>4</sup>Brookhaven National Laboratory, Upton, NY, USA.

<sup>5</sup>Argonne National Laboratory, Lemont, IL, USA.

<sup>6</sup>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.**

1


## Publication on “ePIC Streaming Computing Model”

- **Intended to be an ePIC publication.**
- **Version 1:** Prepared for ECSAC review in 2023. Feedback from the collaboration was solicited from October 11 to November 11, 2023, and incorporated into later versions.
- **Version 2:** Prepared for ECSAC review in 2024. **Added:**
  - 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.
  - A more developed timeline of high-level milestones.
- **We welcome your feedback now via email** and aim to submit the manuscript to *Computing and Software for Big Science* in Spring.

## Parallel Session on “ePIC Data: From Detector Readout to Analysis”:

- **Tomorrow** in the **morning** and **afternoon**.
- Foster a **broader dialogue** within the collaboration on our **data** and **data flow**, with the parallel session serving as a catalyst for that discussion.

# Outlook



*“Software is the soul of the detector”*  
(Ian Shipsey, University of Oxford)

## Great software for great science:

- **Design and Construction:** Integrated and validated simulations are essential for evaluating performance and determining physics reach.
- **Operation:** Autonomous experimentation and control using software/AI. Distributed and rapid processing of streamed data.
- **Research:** Software and data enable discovery.
  
- We **work together**, on a global scale and with other fields, on great software for great science.
- We focus on **modern scientific software & computing practices** to ensure the **long-term success of the EIC scientific program** throughout all CD milestones.

## Discussion Topics:

**Engagement:** Across the collaboration, we dedicate substantial time to ePIC. To ensure a sustainable effort, broader involvement is necessary. Within ePIC Software & Computing, we require additional engagement to sustain ongoing efforts—and even more to expand them. See also slides 18 and 19.


**Communication:** Effective communication among our various entities is a critical for our success. Regular joint meetings with PWGs have significantly improved communication and collaboration. For the DSCs, we propose defining software goals and using them as a foundation for both communication and collaboration.

# Backup

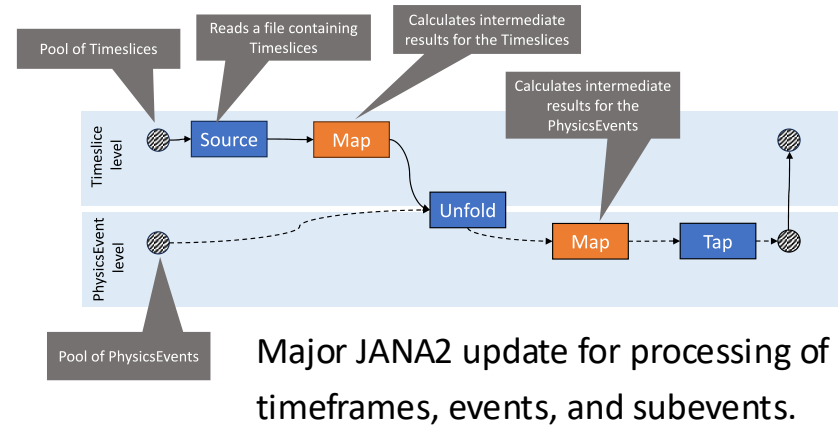
---

# Prototype of Event Reconstruction from Streaming Data

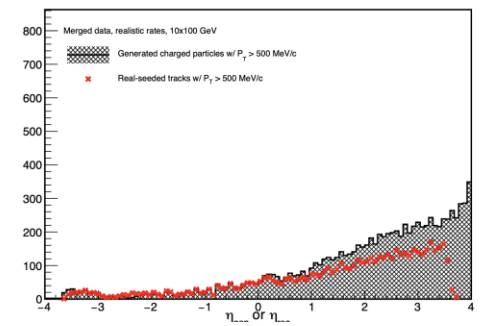
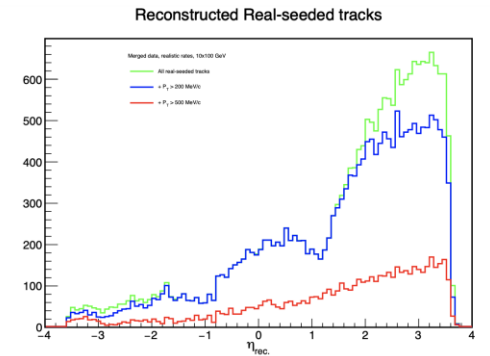
**Scope of the first prototype:** Track reconstruction only. Demonstrated that we can correlate hits in a realistic time frame to the various events in the time window of the MAPS of  $2\mu\text{s}$ .



- Data transferred in collections called *timeframes*.
- Each timeframe includes:
  - Data read from detectors over a time window of  $2^{16}$  cycles of the beam RF, equivalent to 0.6 ms.
  - Channel information and corresponding timing data



Major JANA2 update for processing of timeframes, events, and subevents.

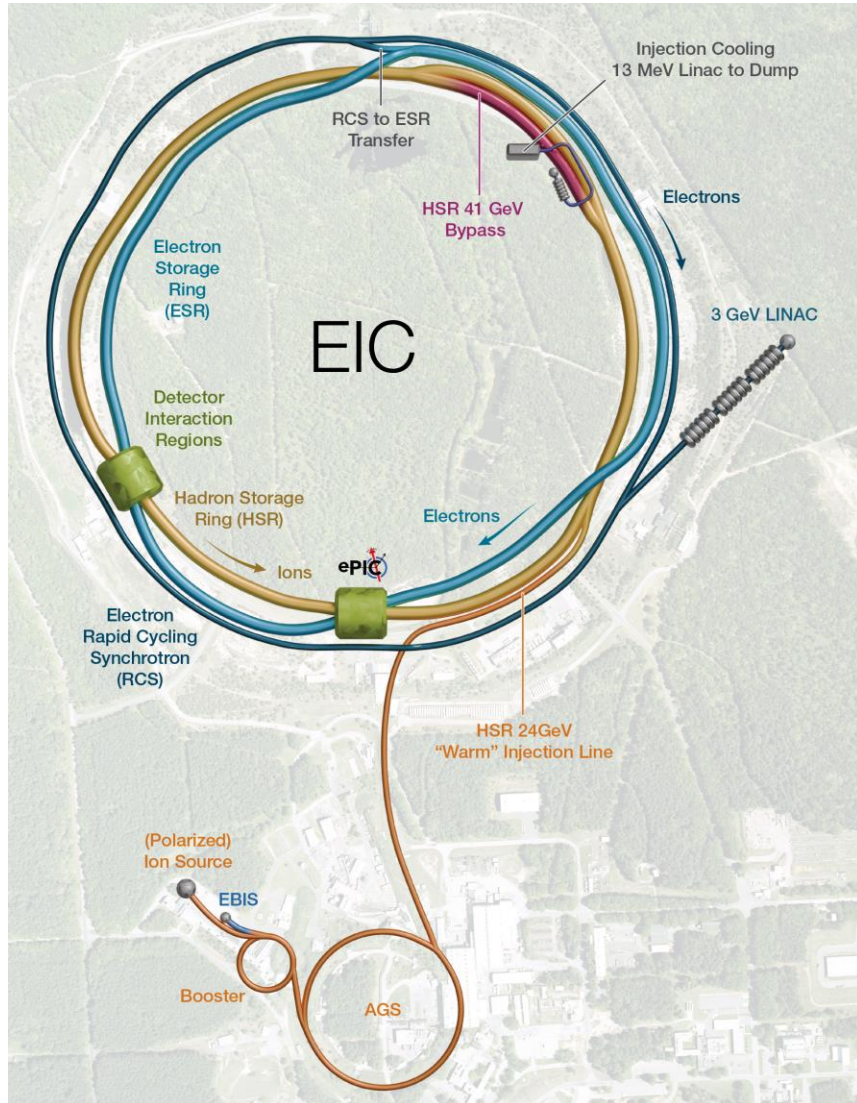


# Image Browser at [https://eic.jlab.org/epic/image\\_browser.html#](https://eic.jlab.org/epic/image_browser.html#)

The screenshot displays the ePIC image browser interface. On the left, there are filters for Campaign (24.03.1, 24.04.0, 24.05.0, 24.06.0), Plot Type (All, scan\_raw\_eta=0\_phi=0, scan\_raw\_eta=-2.5\_phi=0, scan\_raw\_eta=2.5\_phi=0), and Geometry (+). The main area shows 3 images for Campaign: 24.06.0 and Geometry: epic\_craterlake. Each image contains two plots:  $X/X_e$  (Cumulative) vs Path Length [mm] and  $A$  (Cumulative) vs Path Length [mm]. The top image is for  $\eta = 0, \phi = 0^\circ$  and the bottom image is for  $\eta = -2.5, \phi = 0^\circ$ . On the right, there is a search bar and a 'Sort by' dropdown. Below the search bar, there are two sections: 'scan\_raw\_eta=0\_phi=0' and 'scan\_raw\_eta=-2.5\_phi=0'. A text box explains that plots are generated for each simulation campaign. Another text box states that the user interface is user experience driven and provides feedback via the **Contact** section (or GitHub issues).



# Towards a Quantitative Computing Model: The EIC and Event Rates



- **Versatile machine:** versatile range of beam polarizations, beam species, center of mass energies.
- **High luminosity** up to  $L = 10^{34} \text{ cm}^{-2} \text{ s}^{-1} = 10 \text{ kHz}/\mu\text{b}$ .
  - The e-p cross section at peak luminosity is about  $50 \mu\text{b}$ . This corresponds to a signal event rate of about  $500 \text{ kHz}$ .
- The **bunch frequency** will be **98.5MHz**, which corresponds to a **bunch spacing** of about **10ns**.
  - For e-p collisions at peak luminosity, there will be in average 200 bunches or about  $2\mu\text{s}$  between collisions ( $98.5\text{MHz} / 500 \text{ kHz}$ ).
- The EIC Project and ePIC are currently discussing the early science program of the EIC:
  - For the computing resource estimate, we assume a luminosity scenario of  $L = 10^{33} \text{ cm}^{-2} \text{ s}^{-1} = 1 \text{ kHz}/\mu\text{b}$  in 2034.

# Towards a Quantitative Computing Model: Rate Estimates from Streaming DAQ

- **Event size of in average 400 kbit,**
  - Including signal and background apart from detector noise,
  - Assuming that detector noise can be substantially reduced in early stages of processing.
  - Event sizes will decrease in later stages of data taking as detector thresholds are raised.
- **Data rate of in average 30 Gbit/s,**
  - Estimate of upper limit: 10Gbit/s for detector noise + event rate \* event size.
  - Event rate = 50 KHz for EIC Phase 1 luminosity and maximum e-p cross section of  $50 \mu\text{b}$ .
- **Running 60% up-time for ½ year = 9,460,800 s:**
  - Data rate of 30 Gbit/s results in  $710 \times 10^9$  events per year.
  - The data volume of 35.5 PB per year will be replicated between Echelon 1 facilities (71 PB in total).

# Networking Estimates

---

**Echelon 0:** The raw data from the ePIC Streaming DAQ (Echelon 0) will be replicated across the host labs (Echelon 1). At the highest luminosity of  $1e34$ , the data stream from the ePIC Streaming DAQ is estimated at 100 Gbit/s. Consequently, Echelon 0 requires an outgoing network connection of at least 200 Gbit/s.

**Echelon 1:** Each Echelon 1 facility has similar requirements, as it will receive up to 100 Gbit/s of raw data and will share this data with Echelon 2. In addition, Echelon 1 will send a small amount of monitoring data, approximately 1 Gbit/s, back to Echelon 0. Echelon 1 will also receive calibration and analysis data from various Echelon 2 nodes at a comparable rate of about 1 Gbit/s.

**Echelon 2:** The network connection requirements for Echelon 2 facilities will depend on the proportion of raw data they intend to process. For the 10% of Echelon 1 scenario, a network connection of 20 Gbit/s would be required.

# Streaming Data Processing

## Traditional Workflow Characteristics in NP and HEP Experiments:

- Data is acquired in online workflows.
- Data is stored as large files in hierarchical storage.
- Offline workflows process the data, often with substantial latency.
- Batch queue-based resource provisioning is typical.
- Key features: discrete, coarse-grained processing units (files and datasets) and decoupling from real-time data acquisition.

## ePIC Streaming Data Processing Characteristics

- Quasi-continuous flow of fine-grained data.
- Dynamic flexibility to match real-time data inflow.
- Prompt processing is crucial for data quality and detector integrity.
- Processing full data set quickly to minimize time for detector calibration and deliver analysis-ready data.

## Challenging Characteristics of Streaming Data Processing:

- **Time critical**, proceeding in near real time.
- **Data driven**, consuming a fine-grained and quasi-continuous data flow across parallel streams.
- **Adaptive and highly automated**, in being flexible and robust against dynamic changes in data-taking patterns, resource availability and faults.
- **Inherently distributed** in its data sources and its processing resources.

## Assumptions for Infrastructure:

- Existing batch-style processing likely to remain.
- Dynamic processing, e.g. Kubernetes, may displace the batch model.
- Design the system for both batch and dynamic processing to ensure resilience against technology evolution.
- Accommodate but effectively hide these underlying infrastructure characteristics.

# Technology Choices and Evolution

- **Modularity is Key:** We will have a modular software design with structures robust against changes in the computing environment so that changes in underlying code can be handled without an entire overhaul of the structure.
- **Lessons Learned from the NHEP Community** informed the ePIC Streaming Computing Model:
  - Our software is deployed via containers.
  - Our containers are distributed via CernVM-FS.
  - We run large-scale simulation campaigns on the Open Science Grid.
  - Access to our simulations is facilitated through XRootD.
  - We are in the process of deploying Rucio for distributed data management, improving access for collaborators to specific simulation files.
- **Engagement in Advanced Scientific Computing Discussions**, including:
  - DOE SC Round Table on *“Transformational Science Enabled by Artificial Intelligence”* .
  - DOE SC ASCR's Integrated Research Infrastructure (IRI) program. Data-integration-intensive and time-sensitive patterns highly relevant for ePIC.
  - DOE SC ASCR's High Performance Data Facility (HPDF) project, not only enabling these patterns but also potential partnership on data and analysis preservation.
  - International HEP Software Foundation’s discussions on analysis facilities, analysis use cases, and analysis infrastructure.
- **Software Stewardship by the NHEP Community:**
  - **ECSAC Review:** *“ePIC Software & Computing uses many common tools and are active contributors to several.”*
  - Engaging in discussions about HEP Software Foundation Affiliated Projects and Software.

**ECSAC Review:** *“ePIC Software & Computing plans well integrated with standard practices across NHEP.”*

# The Role of AI

- **Compute-detector integration** using:

**Streaming readout** for continuous data flow of the full detector information.

**AI for autonomous alignment and calibration as well as reconstruction and validation for rapid processing.**

**Heterogeneous computing** for acceleration.

- AI will **empower the data processing** at the EIC.
  - Rapid turnaround of data relies on autonomous alignment and calibration as well as autonomous validation.
- AI will also **empower autonomous experimentation and control** beyond data processing:
  - Vision for a responsive, cognizant detector system, .e.g., adjusting thresholds according to background rates.
  - Enabled by access to full detector information via streaming readout.