ePIC Software & Computing Report

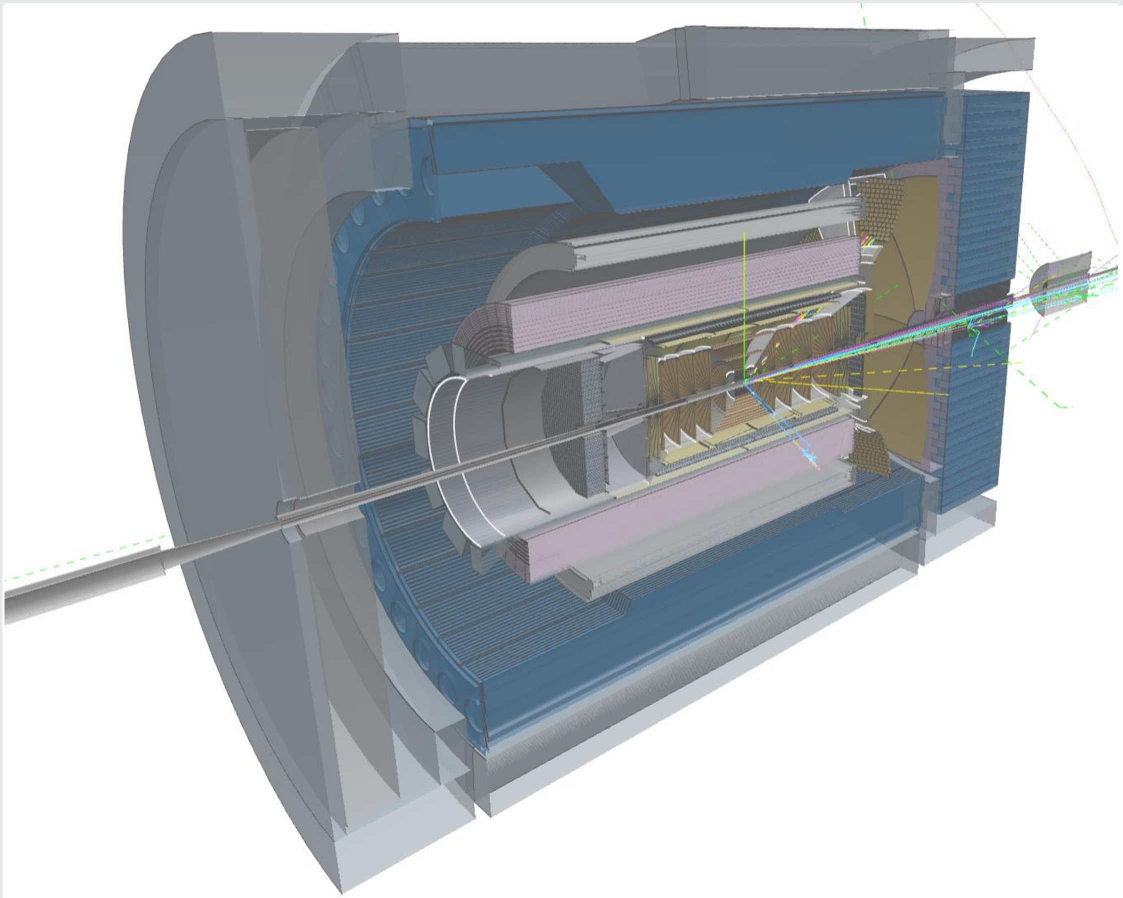

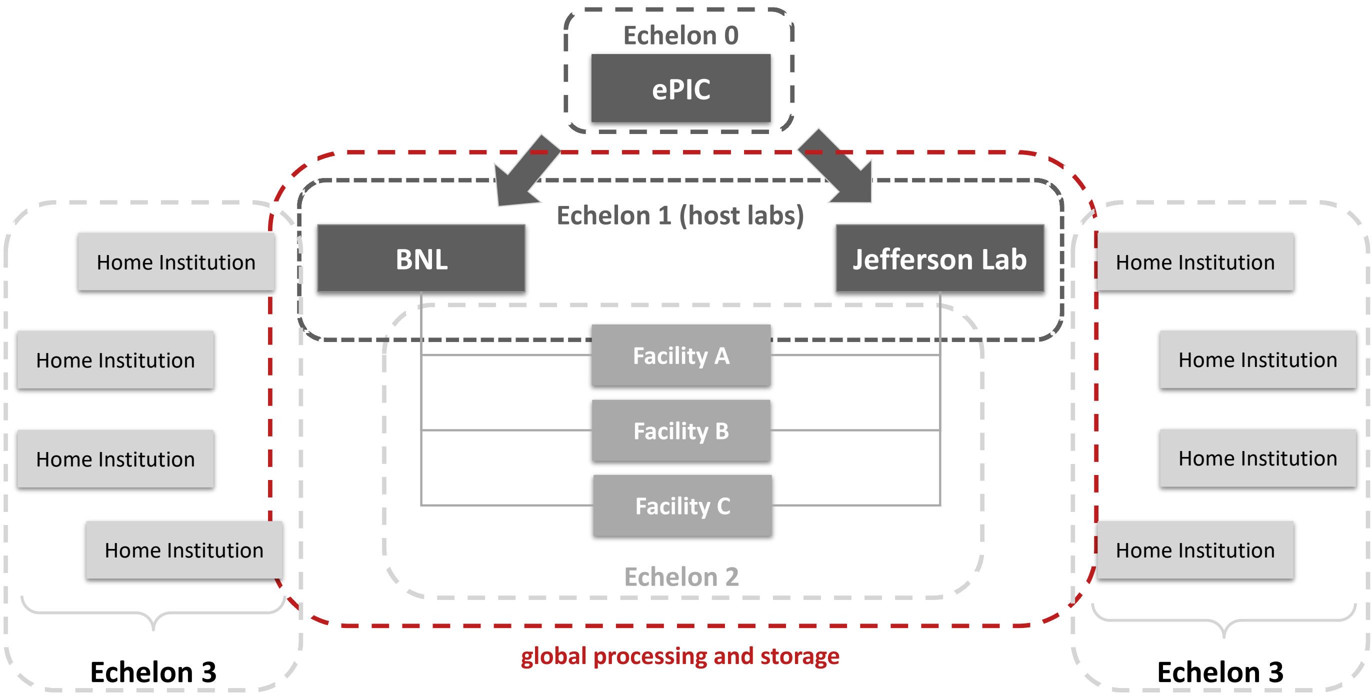

Markus Diefenthaler (Jefferson Lab)

"Software is the Soul of the Detector"

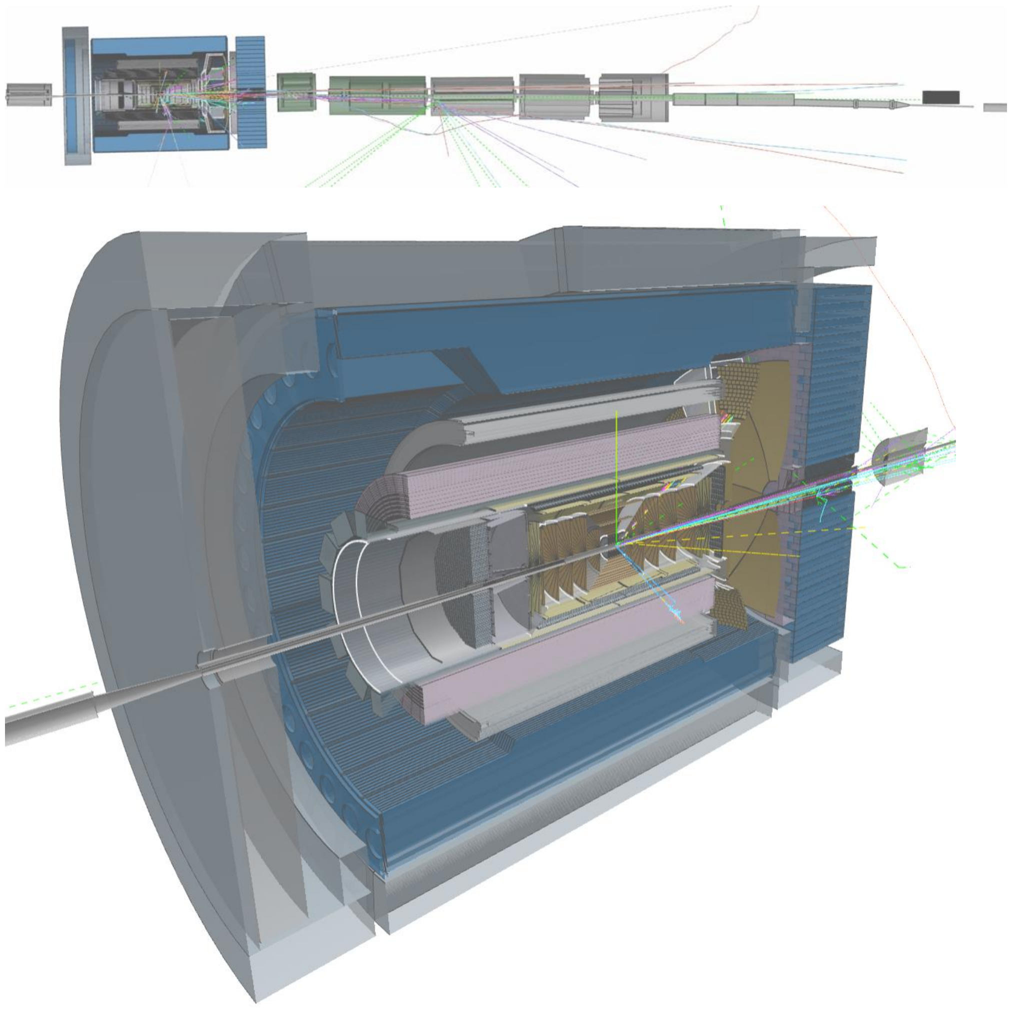

Captured from the ePIC Event Display

Great Software for Great Science:

- **Design and Construction:** Integrated and validated simulations are essential for evaluating detector performance and determining physics reach.
- **Operation:** Seamless processing of data from detector readout to analysis using streaming readout, AI, and distributed computing. Autonomous experimentation and control.
- **Physics Analysis:** 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**.

ePIC Software & Computing Report: Three Talks and One Discussion

MD – ePIC Computing Model: A co-design approach between the detector and the computing enables seamless data flow from detector readout to physics analysis, using streaming readout and AI. Current developments aim to define and test the interface between DAQ and computing, and to mitigate risks in the integrated DAQ-computing system.

Alexei Klimentov – ESCJI Update: We are working closely with the EIC Computing and Software Joint Institute (ECSJI). The joint institute by BNL and JLab supports the computing and software needs and activities of the EIC. Current priorities are the formation of the EIC international Computing Organization (EICO) and resource estimates for the Echelon 1 sites.

Dmitry Kalinkin – ePIC Software and Simulation Campaigns: Advanced software and large-scale simulations underpin both pre-TDR effort and the Early Science Program. Priorities focus on closing gaps in reconstruction and simulation production for an expanding set of physics processes, beam energies, and backgrounds.

Holly Szumila-Vance, Sasha Prozorov, Stephen Kay – Discoverable Software Discussion: Our goal is to make it easy to find, understand, and use the right software. A landing page for onboarding new collaboration members and tutorials aimed at detector experts and physics analyzers have been our first steps. We will discuss next steps toward discoverable software.

Organization

Organization

ECSAC Review

Computing
Model
and Testbeds

Thank You to Wouter!

- Wouter Deconinck has stepped down as Deputy Coordinator for Operations.

- We thank Wouter for his foundational contributions to ePIC software, simulation campaigns, and for his outstanding service to the collaboration.

- Wouter will remain active in ePIC.

Thank You to Chao!

- Chao Peng has stepped down as Convener of the Physics and Detector Simulation WG.
- We thank Chao for his integral contributions to the ePIC software and simulations.
- Chao will remain active in ePIC.

ePIC Software & Computing Organization

Coordinators and WG Conveners

Chandra has been approved as **Reconstruction Working Group convener**, and we are pleased to work with him on **PID software**.

Development

Dmitry Kalinkin





Simulation WG: Simon Gardner (nominated)S, akib Rahman. **Reconstruction WG**: Derek Anderson, Chandra Chatterjee, Shujie Li

Infrastructure

Torre Wenaus



Streaming Computing WG: Marco Battaglieri, Taku Gunji, Jeff Landgraf.

We have nominated **Simon** and **Sasha** as new conveners for the **Simulation** and **User Learning WGs** and are looking forward to working with them.

Operations

Holly Szumila-Vance





Production WG: Thomas Britton, Sakib Rahman. **User Learning WG**: Stephen Kay, Sasha Prozorov (nominated). **Validation WG**: Torri Jeske.

We are excited to work with Holly in her new role as Deputy Coordinator for Operations.

Further Changes in the ePIC Software & Computing Organization

ePIC Committee on Collaborative Software Development Guidelines and Policies

Findings and Recommendations

Version 1.0 (31 Oct 2025)

Derek Anderson (JLab), Nathan Brei (JLab), Johannes Elmsheuser (BNL, chair), Alexander Jentsch (BNL), Sylvester Johannes Joosten (ANL), David Lawrence (JLab), Holly Szumila-Vance (FIU)

Abstract:

The ePIC Committee on Collaborative Software Development Guidelines and Policies was set up in September 2025 to review the current software development process. The review has two primary goals: (a) Fostering the growth of a software developer community within ePIC by assessing our workflows and procedures from a user-centered approach, and (b) Facilitating the integration of new code into the ePIC software stack. This document provides a summary of the findings and recommendations of the committee.

- Charged a **committee to review our collaborative software development guidelines and policies**.
- The committee delivered detailed recommendations on code reviews and Continuous Integration.

- Based on these recommendations, we are revising how **code reviews** are conducted.
- We are developing **clear, practical collaboration guidelines** that define expectations before submitting a pull request.
- These guidelines will improve review quality and support efficient onboarding of new reviewers.
- They will be maintained as a living document as our collaborative practices mature.

preTDR Priorities for 2026

In 2024, we defined preTDR readiness for software and simulations with the collaboration—and successfully met those goals. In 2025, we renewed our priorities at the Frascati Collaboration Meeting, with many topics already reaching an advanced stage and positioning us well for the next phase. **We will discuss priorities for 2026 during this collaboration meeting.**

Production Provide the simulation campaigns with background to finalize the ePIC detector design and validate its physics performance, roll out Rucio as the default system for finding and accessing simulation data, and automate production workflows.

We have reached a level of complexity in simulation production that necessitates a **workflow management system (WFMS)**. We have decided to use **PanDA** for our current simulation productions. This decision will be revisited in **FY28**, when we evaluate WFMS options for **streaming orchestration** and select one.

Reconstruction Coordinate the effort to address gaps in reconstruction. Work toward a holistic reconstruction approach, such as particle identification that integrates information from calorimeters, Cherenkov detectors, and time-of-flight systems.

Simulation Implement and operate a workflow between detector and simulation experts to track the status of the comparison between the simulation design and the engineering design, and to resolve any discrepancies in a timely and systematic manner.

Streaming Orchestration Define and test the interface between DAQ and computing to mitigate risks in the integrated DAQ-computing system. As part of the integration, test develop and deliver a functional testbed that validates a workflow management system for the autonomous calibration of a detector subsystem.

Workforce Foster a developer community within ePIC and support the careers of scientists focused on software & computing.

Additional Priorities for 2026

Priority: With substantial progress in analysis prototyping, it is timely to discuss analysis requirements for Software & Computing.

Analysis Tools Review analysis workflows and requirements with the Physics WGs to enhance reconstruction algorithms and simulation output, and to identify priorities for the development of the analysis model and tools for ePIC.

Priority: We have many prototypes for AI approaches and need to integrate them into software and simulation productions.

AI at Scale Coordinate and support AI development within ePIC. Use the integration of AI into simulation, reconstruction, and analysis workflows as a primary success metric. Provide MLOps infrastructure and support.

Topic in AI Workfest

ECSAC Recommendation: We **recommend** the ePIC collaboration and the host labs to start developing a data management and lifecycle plan, agreed with the Funding Agencies.
Priority: We are finalizing the DAQ-computing interface need to determine whether there are requirements related to AI readiness.

Data Develop a data management and lifecycle plan that addresses the AI readiness of the ePIC experiment, supports multimodal data (including metadata), and ensures the reproducibility and reusability of both data and analysis workflows.

Topic in AI Workfest

ECSAC Review

Organization

ECSAC Review

Computing
Model
and Testbeds

EIC Computing & Software Advisory Committee (ECSAC) Review


Ilya Baldin (JLab)


Simone Campana (CERN, chair)


Mario Cromaz (LBNL)


Pere Mato (PIC)


Verena Martinez Outschoorn (UMass Amherst)

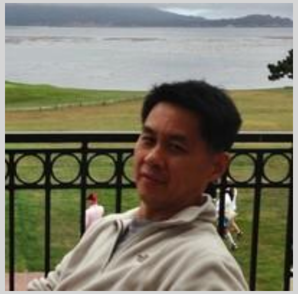
Wei Yang (SLAC)

Reviews occur annually, with a charge reflective of the EIC schedule, the stage of the ePIC experiment, and impending deadlines. The 2025 review was held at Jefferson Lab on October 6–7. 2025.

Charge of 2025 ECSAC Review:

- Progress toward TDR readiness and implementation of the computing model
- Short- and long-term resource planning
- ECSJI resource support and laboratory collaboration (see ECSJI Report later today)
- Engagement with the broader software and computing community
- ECSJI-ePIC collaboration effectiveness

ECSAC Review (Oct. 6–7): Highlights and Recommendations

From the ECSAC Closeout Report

“We commend the ePIC collaboration for the progress in Software & Computing. It is very positive to see the active engagement of additional partners from U.S. and international institutes. There has been substantial progress in software and computing development, notably in streaming data processing, simulation campaigns, and software readiness for the pre-TDR.”

ECSAC commended ePIC’s progress and made five recommendations:

1. We **recommend** the ePIC collaboration to establish a multi-year plan for SW&C, focusing on the deliverables for the next three years. Such a plan should be presented at the next ECSAC meeting in Fall 2026.
2. We **recommend** the ePIC collaboration and the host labs to start developing a data management and lifecycle plan, agreed with the Funding Agencies.
3. We **recommend** the ePIC collaboration to continue the commissioning plan through the ongoing demonstrators and testbeds. Such a commissioning process should include assessing the E0-E1 interfaces and workflows.
4. We **recommend** the ePIC to continue engaging the institutes (including the host labs) to identify possible additional contributions for SW&C developments and operations. We expect to hear about the level of engagement of the institutes and the success in addressing resource gaps at the next ECSAC meeting.
5. We **recommend** implementing the EICO as described during the review. We suggest a staged approach focusing initially on the aspects more related to the collaboration (Collaboration/Overview boards) and the links with WLCG.

ECSAC Review

Organization

ECSAC Review

Computing Model and Testbeds

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.**

1

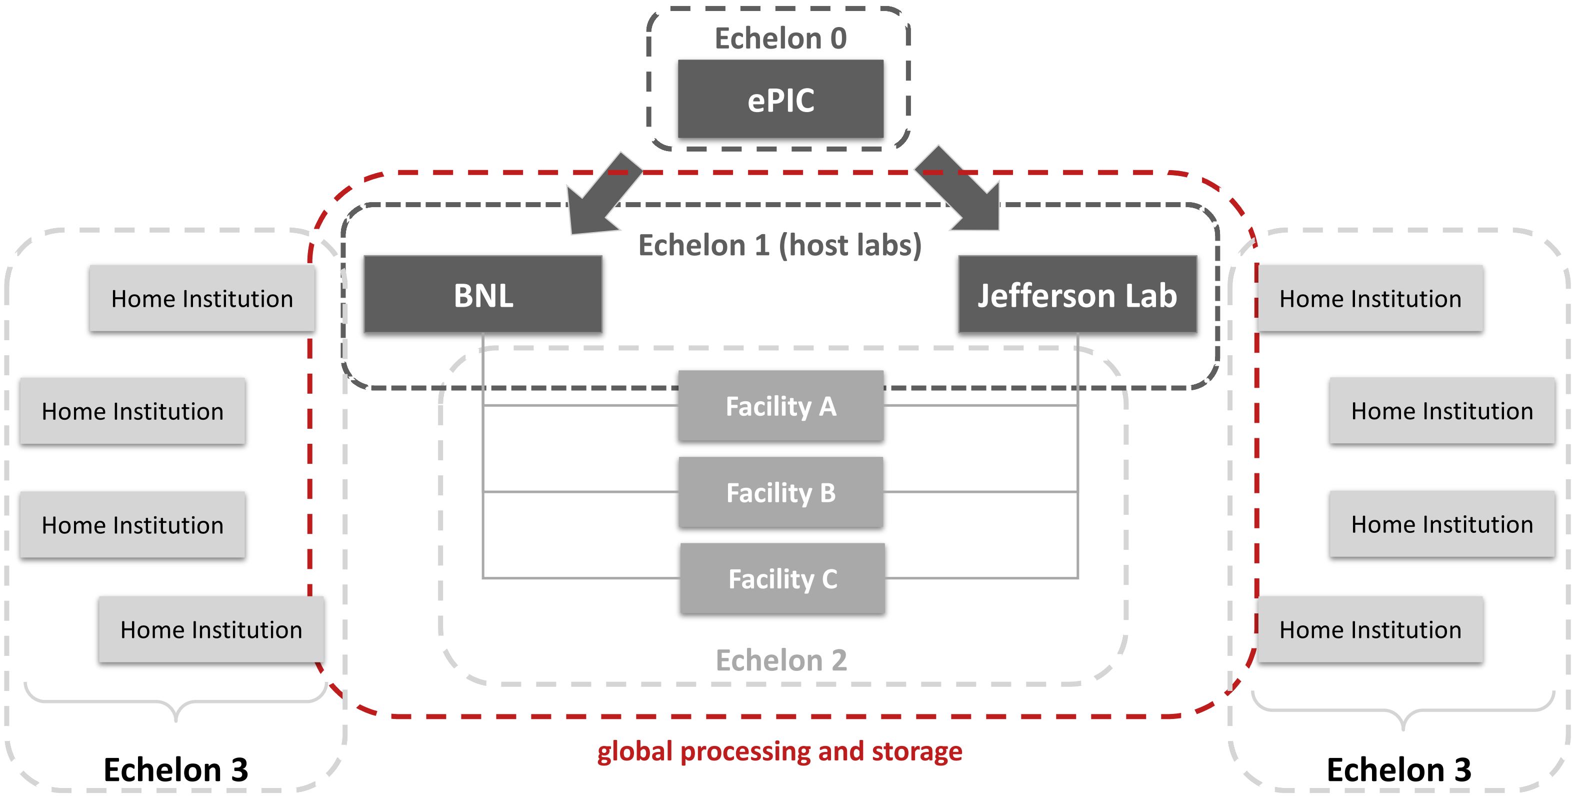

We developed the ePIC Streaming Computing Model to accelerate the pace of discovery and enhance scientific precision through improved management of systematic uncertainties. The model is documented in a detailed report and was reviewed during the 2023 and 2024 ECSAC reviews.

ePIC Within the Global Particle Physics Experiments Landscape

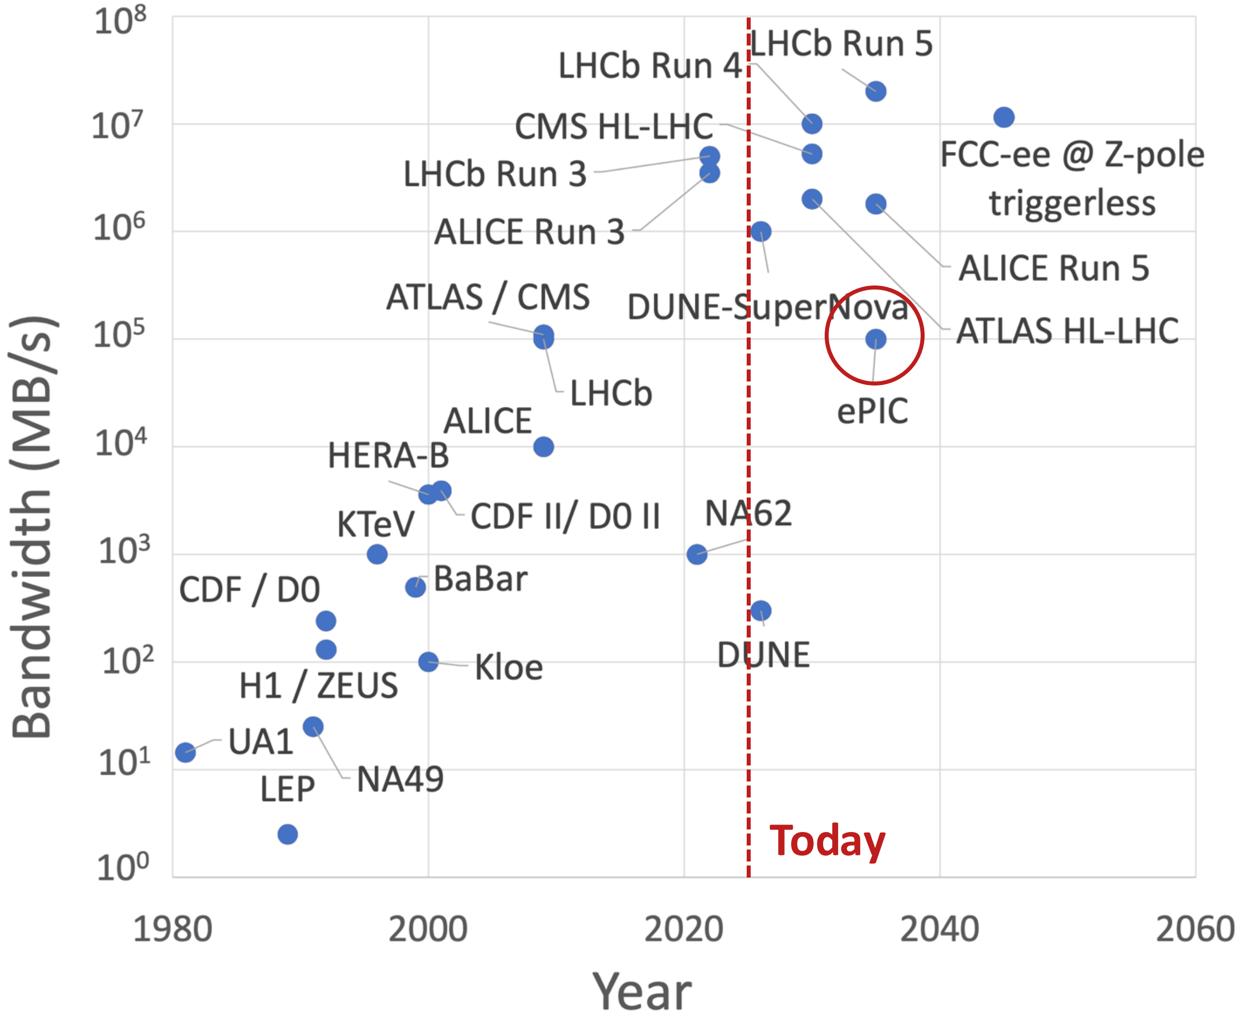

Streaming Readout

Data rate of up to 100 Gbit/s

after low-level data reduction in the Streaming DAQ

Aarrestad, Thea, and Dorothea vom Bruch. *Trigger and Data Acquisition: Challenges and Perspectives*. Presentation at the Open Symposium on the European Strategy for Particle Physics, Venice, Italy, June 23, 2025. https://agenda.infn.it/event/44943/contributions/265988/

Prototyping Ideas and Tools in Testbeds

With active testbeds and functional prototypes now in place, the effort is moving from design to implementation. These developments aim to define and test the interface between DAQ and computing, and to mitigate risks in the integrated DAQ-computing system.

- **Streaming orchestration**, i.e., a workflow and workload management system for streaming data—is essential for system testing. A requirements document has been developed and is now guiding testbed and prototype development.
  - **Testbed plans** are taking concrete shape:

Streaming reconstruction: Raw data stream to event identification, reconstruction, and analysis.

Streaming orchestration: Developing E0-E2 streaming workflows in the testbed, utilizing Rucio and PanDA.

Streaming processing: Developing E0-E2 streaming workflows using EJFAT.

  - Not covered in this presentation but starting efforts:

Streaming analysis: Demonstrate simulation data production streaming to E2 site.

Rapid data processing: Autonomous calibration workflow for one detector system.

JANA2 for Streaming Processing

- Multithreaded JANA2 framework provides a component-level hierarchical decomposition of data boundaries into **Run**, **Timeframe**, **PhysicsEvent**, and **Subevent** levels. This is essential for streaming processing.
- The **Folder** and **Unfolder** component interfaces enable traversal of this hierarchy by supporting operations such as splitting and merging data streams. This functionality has been tested and validated within **EICrecon**.

Introducing multilevel sources

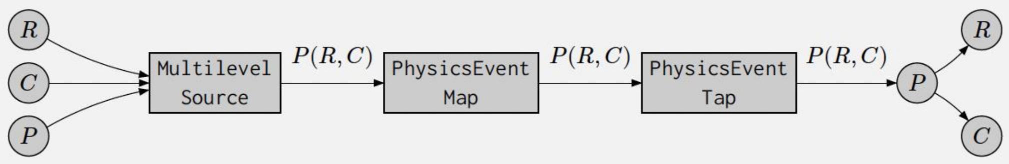

EICrecon timeframe splitting

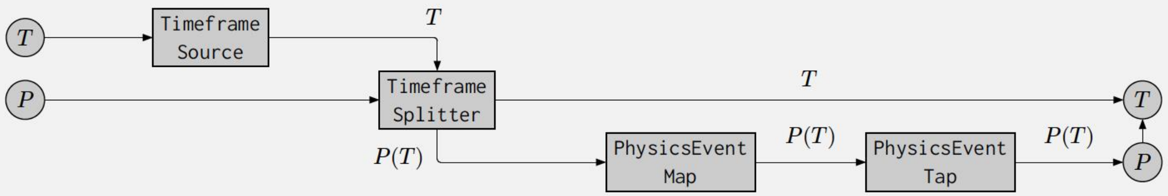

Multilevel sources with timeframe splitting

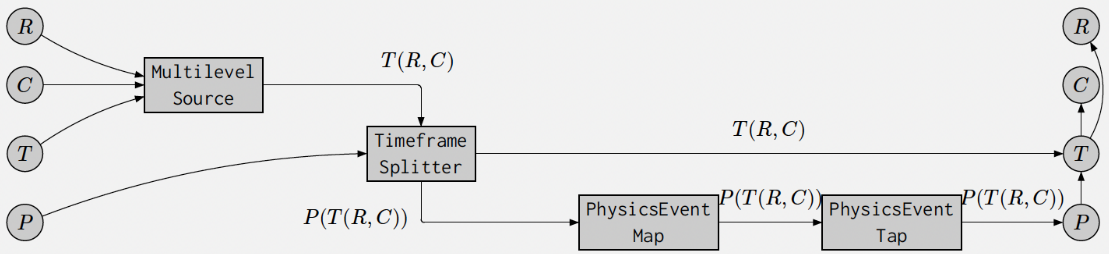

Timeframe sources with multilevel splitting

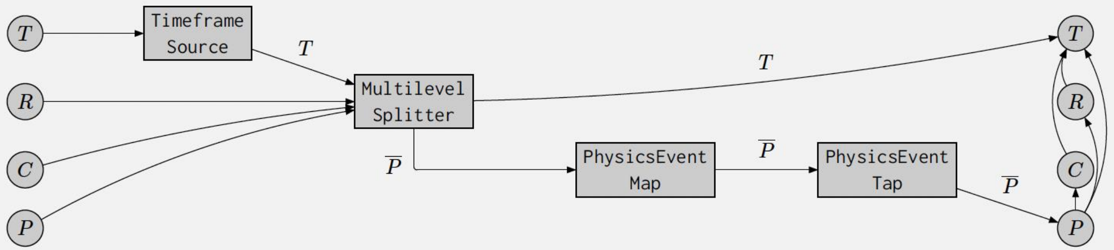

Streaming Reconstruction Prototype: Event-Building in JANA2

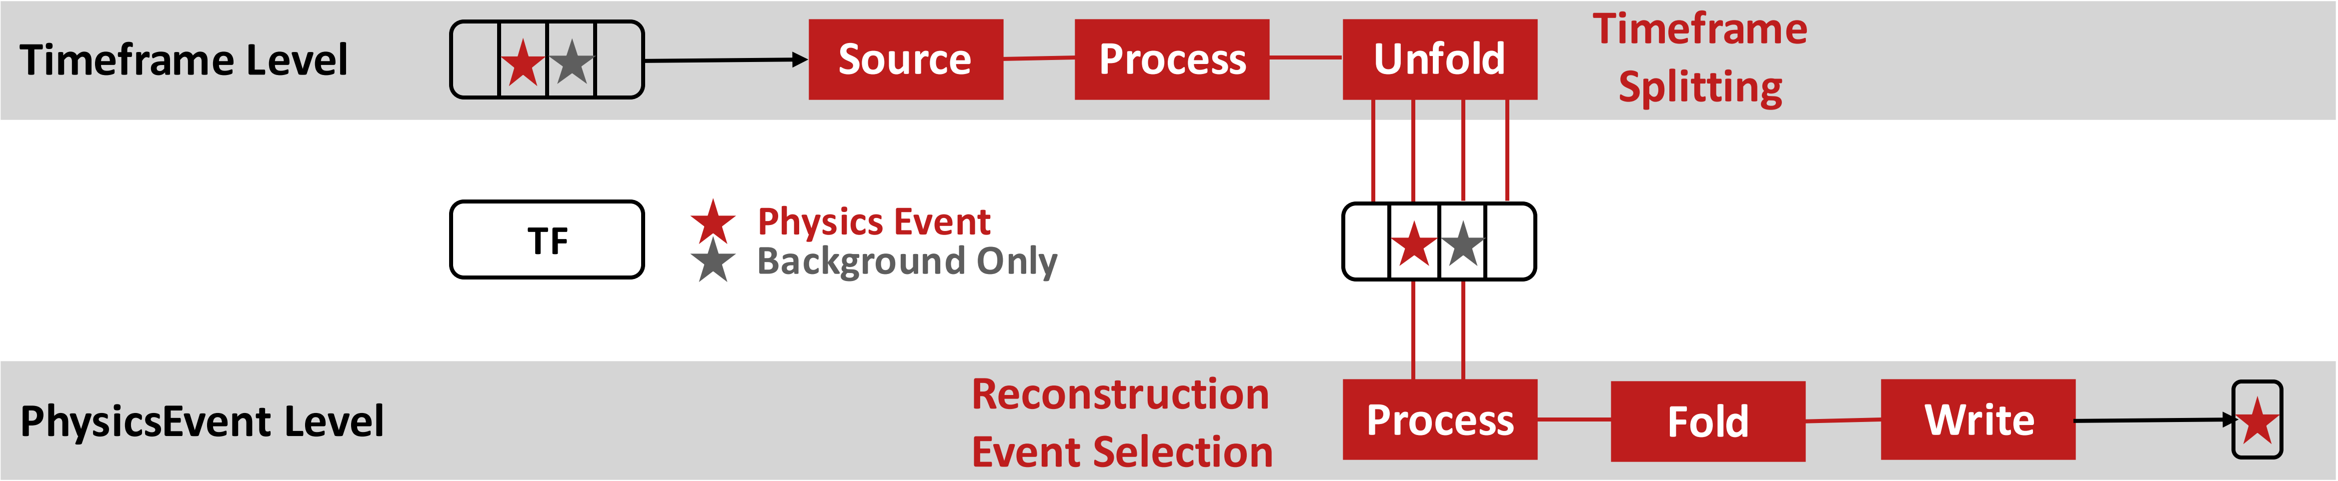

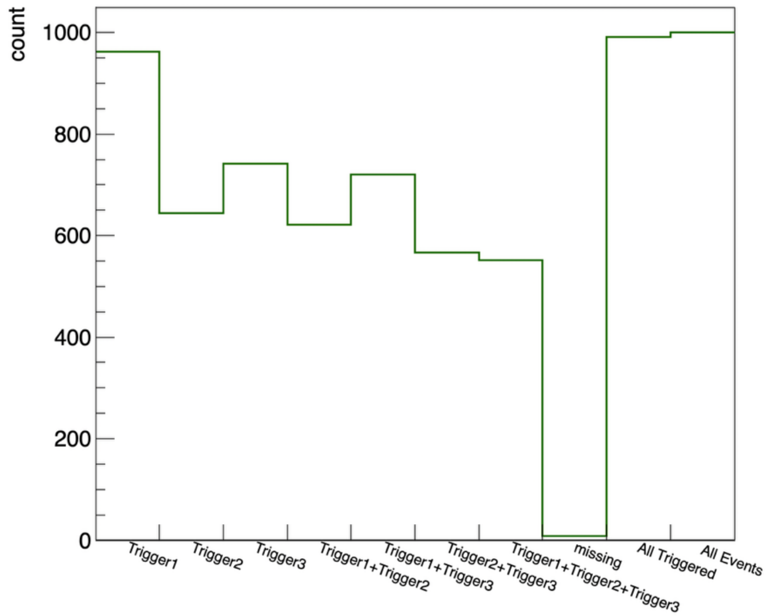

Trigger 1: Coincidence hits in the SVT Endcap and Forward MPGD Endcap.
Trigger 2: Coincidence hits in the TOF Barrel and MPGD Barrel.
Trigger 3: Coincidence hits in the SVT Endcap and Backward MPGD Endcap.

A straightforward event selection based on raw simulated hits in the SVT, MPGD, and TOF detector systems achieves greater than 99 % efficiency and less than 1 % background in identifying physics events in TFs.

Streaming Orchestration Testbed

Motivation:
- Evaluate how well existing distributed computing tools support streaming orchestration.
- Focus on practical deployment and performance in realistic environments.

- **Design Precepts**:
  - Robust geographical distribution across real-world networks.
  - Full automation of data processing workflows.
  - Complete exposure of system status and operational analytics.
- **Approach**:
  - PanDA and Rucio align with the stated design precepts and are deployed in live testbed instances at BNL.
  - Assume that data is delivered in STFs, each consisting of 1000 aggregated TFs, with a size of ~2 GB at a rate of ~1 Hz.
- **System Architecture:** Simulated DAQ with Distributed Agents and Monitoring
  - A **set of collaborating agents** communicate via ActiveMQ and form the core of the system:
    - **Data handling agent**: executes Rucio actions based on data and conditions.
    - **Processing agent**: executes PanDA actions based on available data and conditions.
    - **Monitoring agent**: skims STF data for fast feedback.
  - The **driver simulates the DAQ** and other external components to evaluate their impact.
  - A **monitoring backend** with a database aggregates and exposes system state.

End-to-End Workflow: From Run Start to Monitoring

Workflow driven by DAQ simulator and 3 agents communicating via ActiveMQ.
All system activity and state recorded in the database via REST and displayed in the monitor.

1. **Run Start** DAQ simulator broadcasts a message signaling the start of a new data-taking run (**working**).
2. **Dataset Creation** The data agent receives the message and instructs Rucio to create a dataset for the run (**working**).
3. **Processing Task** Processing agent sets up a PanDA task based on the run start message (**working**).
4. **STF Available** DAQ simulator broadcasts availability of a new STF file; this continues while the run is active.
5. **STF Transfer** The data agent triggers Rucio registration and transfers the STF to E1 storage (**not yet integrated**).
6. **STF Processing** PanDA detects the new STF at E1, transferred via Rucio, and launches jobs to process it (**working**).
7. **Fast Monitoring** The fastmon agent sees the STF broadcast, performs a partial read, and injects a data sample into the E1/E2 monitoring stream (**fast monitor emulation working based on STF metadata**).

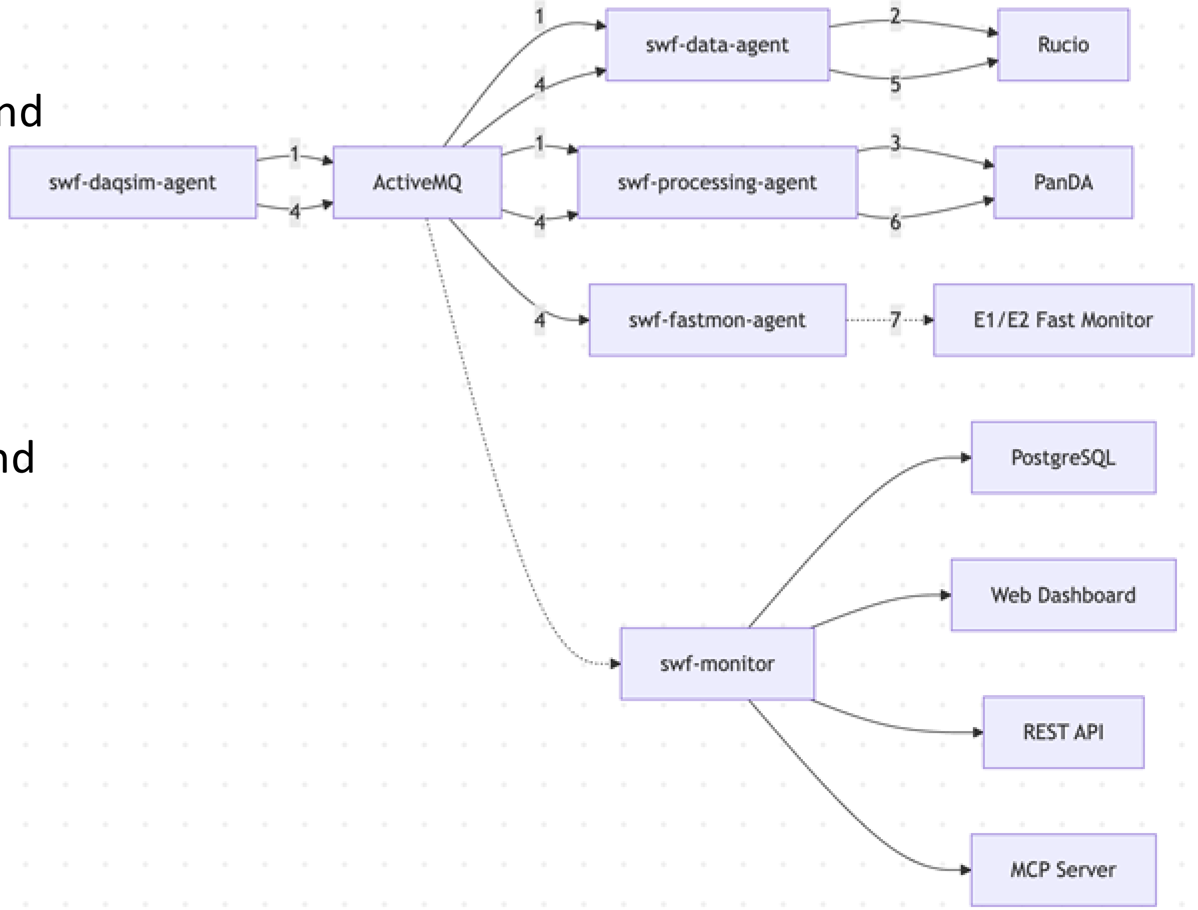

State Machine: Stream-Oriented Workflow States and Substates

States

- no_beam
  - Collider not operating
- beam
  - Collider operating
- run
  - Physics running
- calib
  - Dedicated calibration period
- test
  - Testing, debugging
  - Any substates can be present during test

Substates

- not_ready
  - detector not ready for physics datataking
  - occurs during states: no_beam, beam, calib
- ready
  - collider and detector ready for physics, but not declared as good for physics
  - when declared good for physics, transitions from beam/ready to run/physics
  - occurs during states: beam
- physics
  - collider and detector declared good for physics
  - if collider or detector drop out of good for physics, state transitions out of 'run' to 'beam' or 'off'
  - occurs during states: run
- standby
  - collider and detector still good for physics, but standing by
  - occurs during states: run
- lumi
  - detector, machine data that is input to luminosity calculations
  - occurs during states: beam, run
- eic
  - machine data, machine configuration
  - occurs during states: all
- epic
  - detector configuration, data
  - occurs during states: all
- daq
  - info, config transmitted from DAQ
  - occurs during states: all
- calib
  - a catch-all for a great many calib data types, we can start small
  - occurs during states:

The baseline workflow exercises nearly all states, with the exception of detector and machine configurations.

Streaming Data Processing Demonstration

Secretary Wright's Visit to Jefferson Lab on August 21

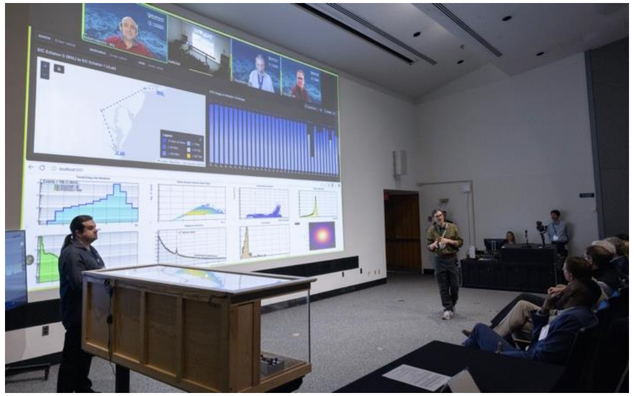

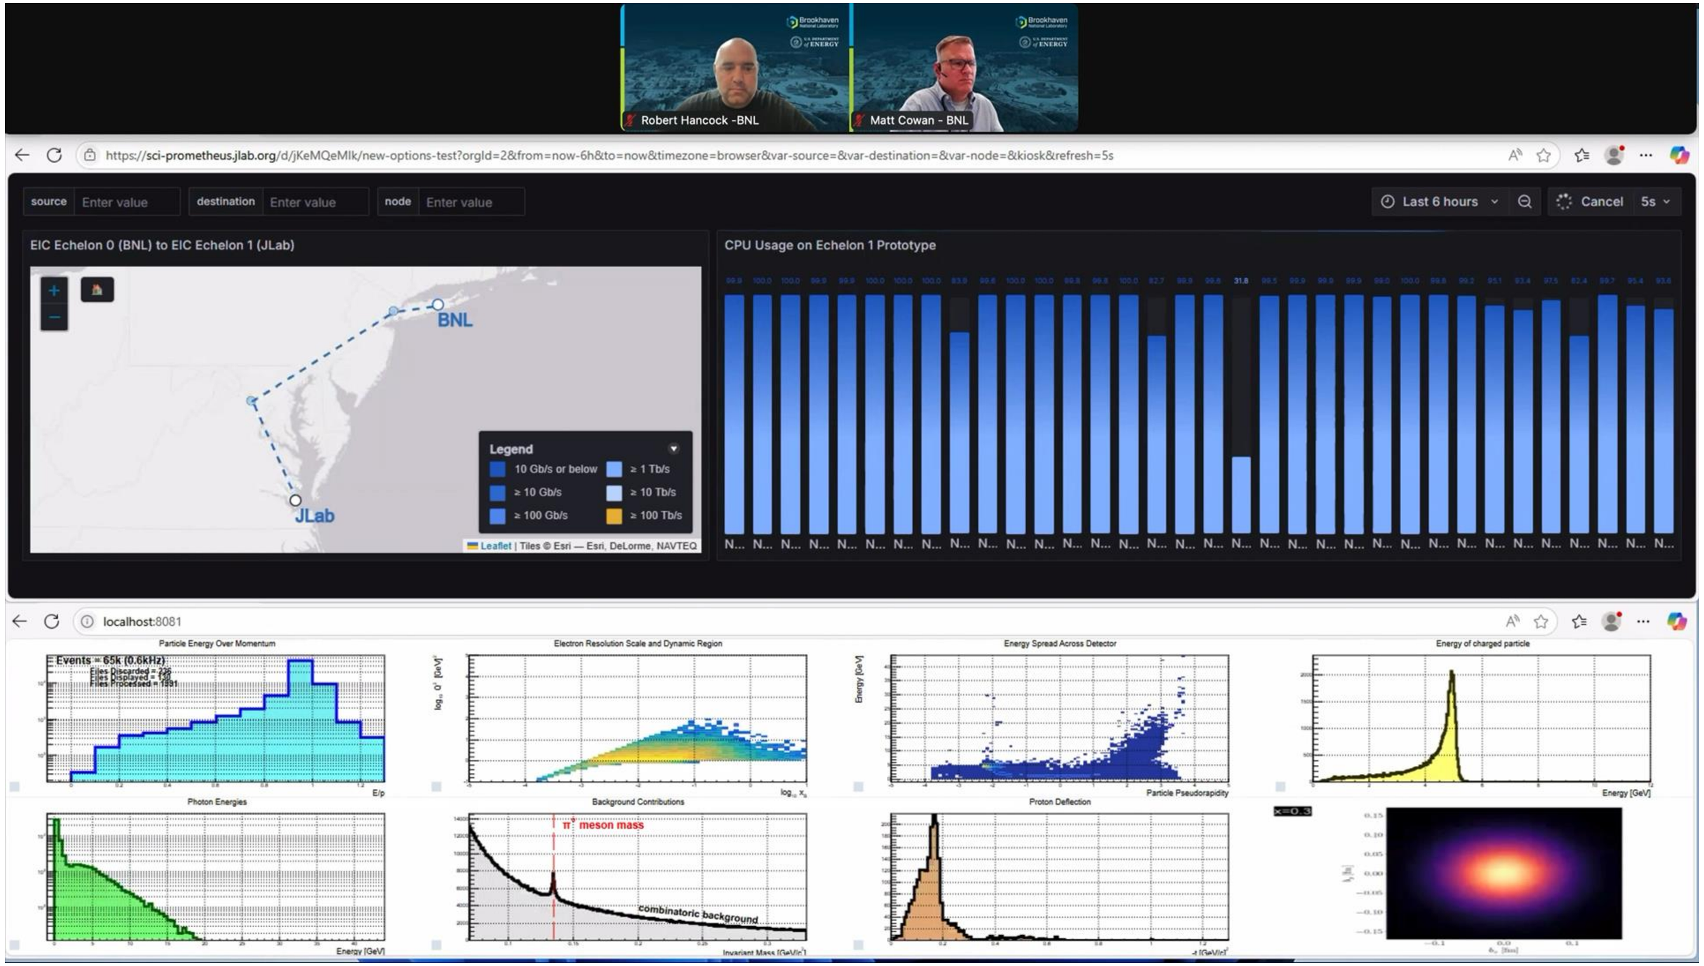

We successfully showcased **real-time data transfer from BNL to JLab** using the ESnet-JLab FPGA Accelerated Transport (EJFAT) Load Balancer and its **processing at JLab**. It also prompted collaboration with BNL to test current tools and establish a clear network path.

Streaming DAQ and Computing Milestones

FY25	FY26	FY27	FY28	FY29	FY30	FY31	
PicoDAQ	MicroDAQ	MiniDAQ	Full DAQ-v-1	Production DAQ			DAQ
Streaming Orchestration			Streaming Challenges				
AI-Empowered Streaming Data Processing			Analysis Challenges				Computing
			Distributed Data Challenges				
AI-Driven Autonomous Calibration			AI-Driven Autonomous Alignment, Calibration, and Control				AI

- **Compute-Detector Integration**:
  - Joint deliverables between **DAQ** and **computing** to develop integrated systems for detector readout, data processing, and ultimately physics analysis.
  - **Key role of AI(/ML)**: Empowering data processing and enabling autonomous experimentation and control.

- **FY28Q1 deliverables**:
  - fully functional testbed for streaming orchestration,
  - autonomous calibration workflow for one detector system,
  - AI/ML-empowered streaming reconstruction.

Distributed Computing for ePIC

- **Echelon 1** sites uniquely perform the **low-latency streaming workflows**:
  - 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.
- **Priority**: Establishing EIC International Computing Organization (EICO):

ECSAC recommended a staged implementation starting with the Collaboration and Overview Boards. See ECSJI presentation for details on EICO.

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 Model

Building Echelon 2 capacity through early integration in simulation campaigns.

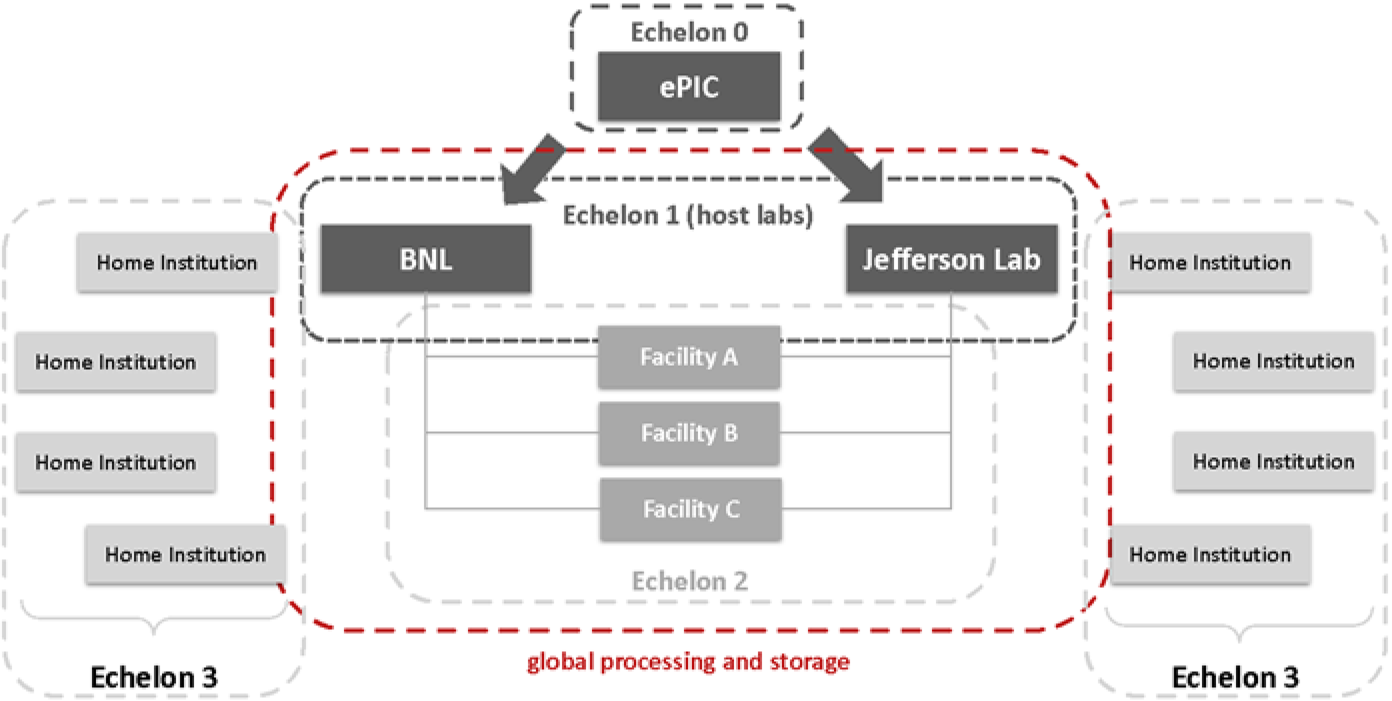

- Central infrastructure at **JLab**:
  - HTCondor submit node
  - Rucio main server
  - Rucio storage element
- Central infrastructure at **BNL**:
  - HTCondor submit node
  - Rucio storage element

- Active integration of **Canada** and **Italy** as compute providers integrated through OSG.

- Commissioning of **Canada** as Rucio storage element provider.
- Planning in progress for **Japan** and **Taiwan** as storage providers.

- Discussions on potential compute and storage contributions are underway with **France** and the **United Kingdom**.

Summary

Organization

- **Deputy Coordinator and convener roles were updated**; will support the next phase of the Software & Computing activities.
- **Committee recommendations** on **code reviews** and **Continuous Integration** are being implemented.
- **Priorities continue to be defined and refined** within collaboration meetings.

ECSAC Review

- **ECSAC Review, commended progress** notably in streaming data processing, simulation campaigns, and software readiness for the pre-TDR.
- ECSAC recommended **multi-year planning**, a **data lifecycle strategy**, **continued testbed development**, **institute engagement to address workforce gaps**, and a **staged implementation of EICO**.

Computing Model and Testbeds

- **Streaming Computing Model** is defined, reviewed, and actively guiding ongoing developments. Its integration of computing and detector systems of ePIC experiment **maximizes scientific output** and **accelerates scientific discovery**.
- Effort **transitioning from design to implementation**, with active and functional prototypes.
- **Deliverables are aligned with EIC Project / Streaming DAQ milestones**.
- ePIC will be **compute-intensive experiment**; **substantial role for Echelon 2** foreseen.

Backup

Science Drivers for Streaming Readout at ePIC

Broad ePIC Science Program:

- Plethora of observables, with less distinct topologies where every event is significant.

Moderate Signal Rate:

	EIC	RHIC	LHC → HL-LHC
Collision species	$\vec{e}+\vec{p}, \vec{e}+A$	$\vec{p}+\vec{p}/A, A+A$	$p+p/A, A+A$
Maximum x-N C.M. energy	140 GeV	510 GeV	13 TeV
Peak x-N luminosity	10^{34} cm^{-2} s^{-1}	10^{32} cm^{-2} s^{-1}	$10^{34} \rightarrow 10^{35}$ cm^{-2} s^{-1}
x-N cross section	**50 µb**	**40 mb**	**80 mb**
Maximum collision rate	500 kHz	10 MHz	1-6 GHz
$dN_{ch}/d\eta$	0.1-Few	~3	~6
Charged particle rate	**4×10^6 N_{ch}/s**	**6×10^7 N_{ch}/s**	**3×10^{10}+ N_{ch}/s**

Enabling Next-Generation Compute-Detector Integration

- **Maximize Science**: **Capture every collision signal, including background.**
    - **High-precision measurements**: Control of systematic uncertainties is critical.
    - 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 two weeks for data for physics analyses**.
    - Timeline driven by alignment and calibration.
    - Subsystem experts indicate a two-week turnaround is feasible.
- **Technologies**: Compute-detector integration using:

Streaming Readout for continuous data flow of the full detector information.	**Artificial Intelligence** for rapid processing (autonomous alignment, calibration, and validation).	**Heterogeneous Computing** for acceleration (CPU, GPU).

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**

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.

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**

~350 PB to store data of one year.

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

Requirements for Streaming Orchestration

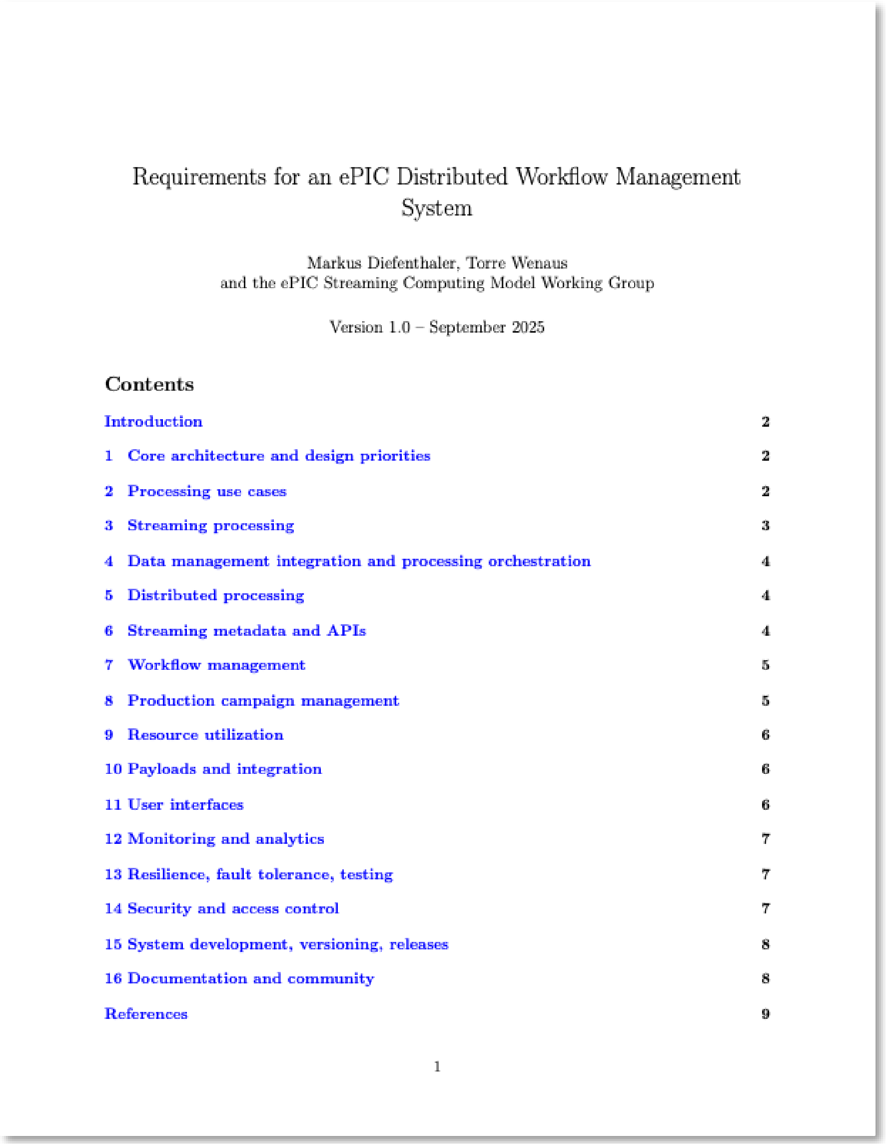

Requirements for an ePIC Distributed Workflow Management System

Markus Diefenthaler, Torre Wenaus
and the ePIC Streaming Computing Model Working Group

Version 1.0 – September 2025

Contents

1

Requirements Document

- Builds upon and guides further development of the **ePIC Streaming Computing Model.**
- Developed collaboratively by the **ePIC Streaming Computing Model WG**.
- Informed by **lessons learned from other experiments** and streaming systems.

Key Themes

- **Scalable and Automated Workflows**: Low overhead, automated orchestration, and real-time processing across E1–2.
- **Streaming-First Design**: Native support for near real-time processing.
- **Integrated Data Management**: Tight coupling with Rucio-based DDM for data-driven workflows and provenance tracking.
- **Flexible & User-Centric Interfaces**: CLI, REST, and web interfaces with support for custom dashboards and diagnostics.
- **Robust Monitoring & Resilience**: Real-time analytics, fault tolerance, and automated recovery mechanisms.
- **Community and Documentation Focus**: Open development, transparent processes, and collaborative design.