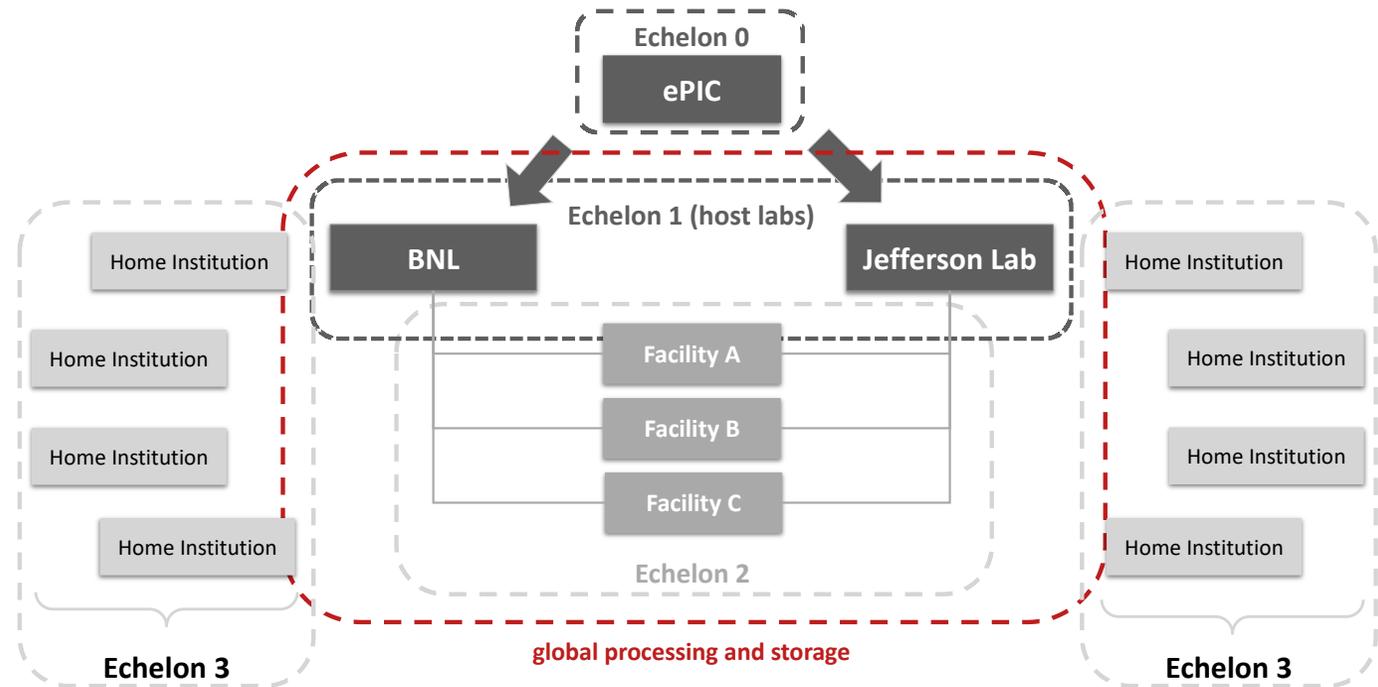
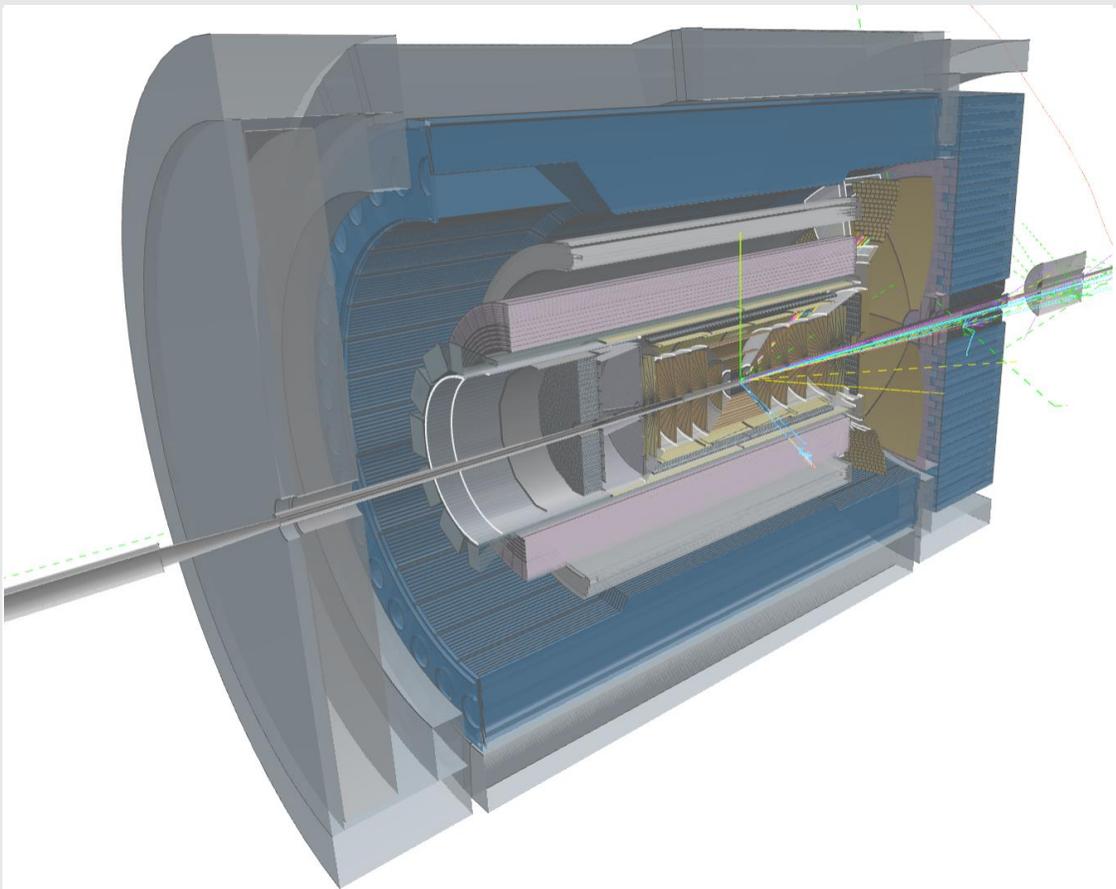


Computing Model Overview



Markus Diefenthaler (Jefferson Lab)

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 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.

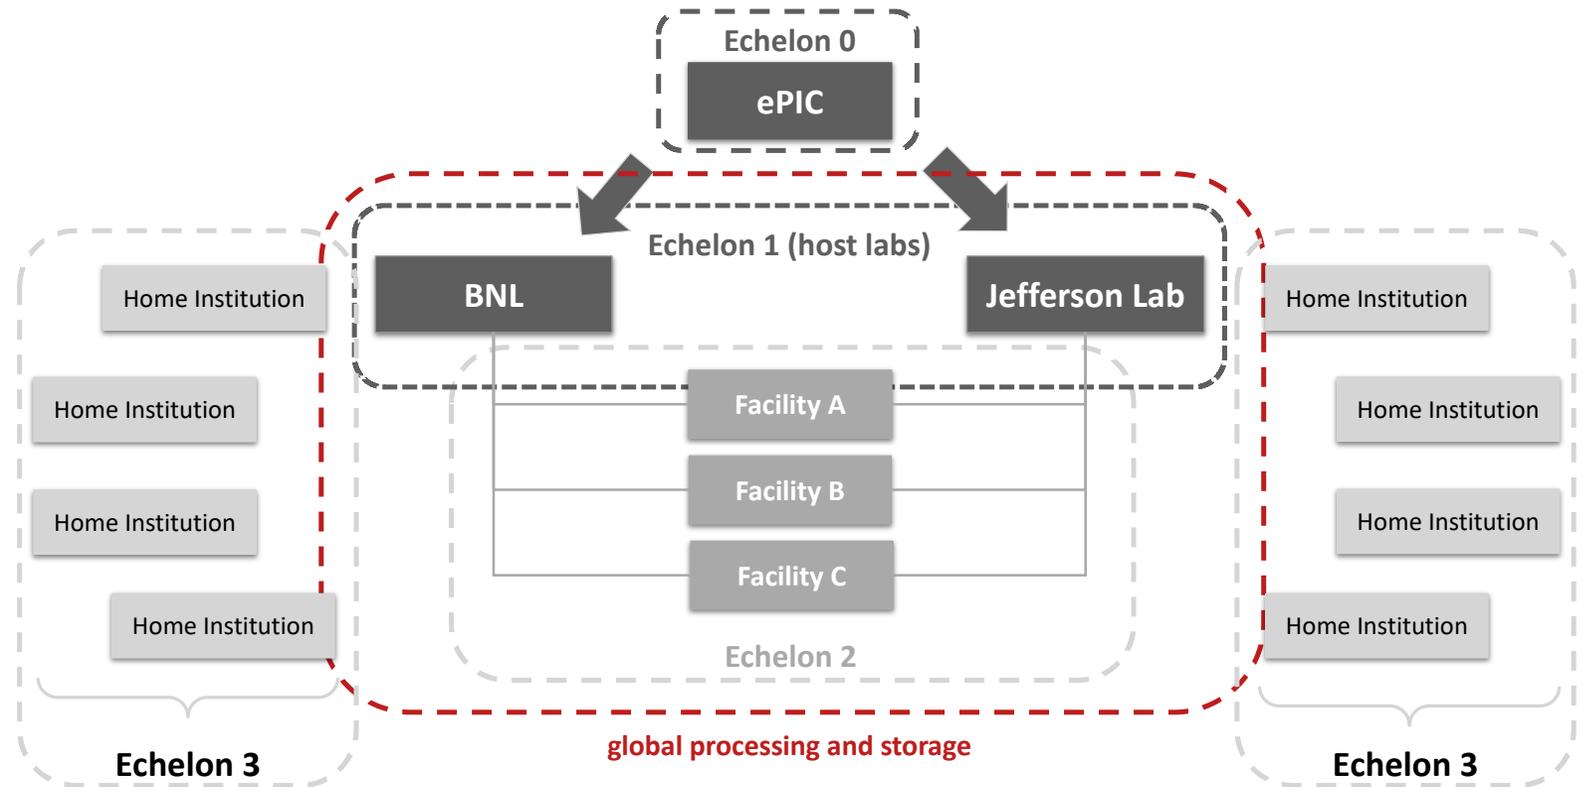
⁵Argonne National Laboratory, Lemont, IL, USA.

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.

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.

Enabling Next-Generation Compute-Detector Integration

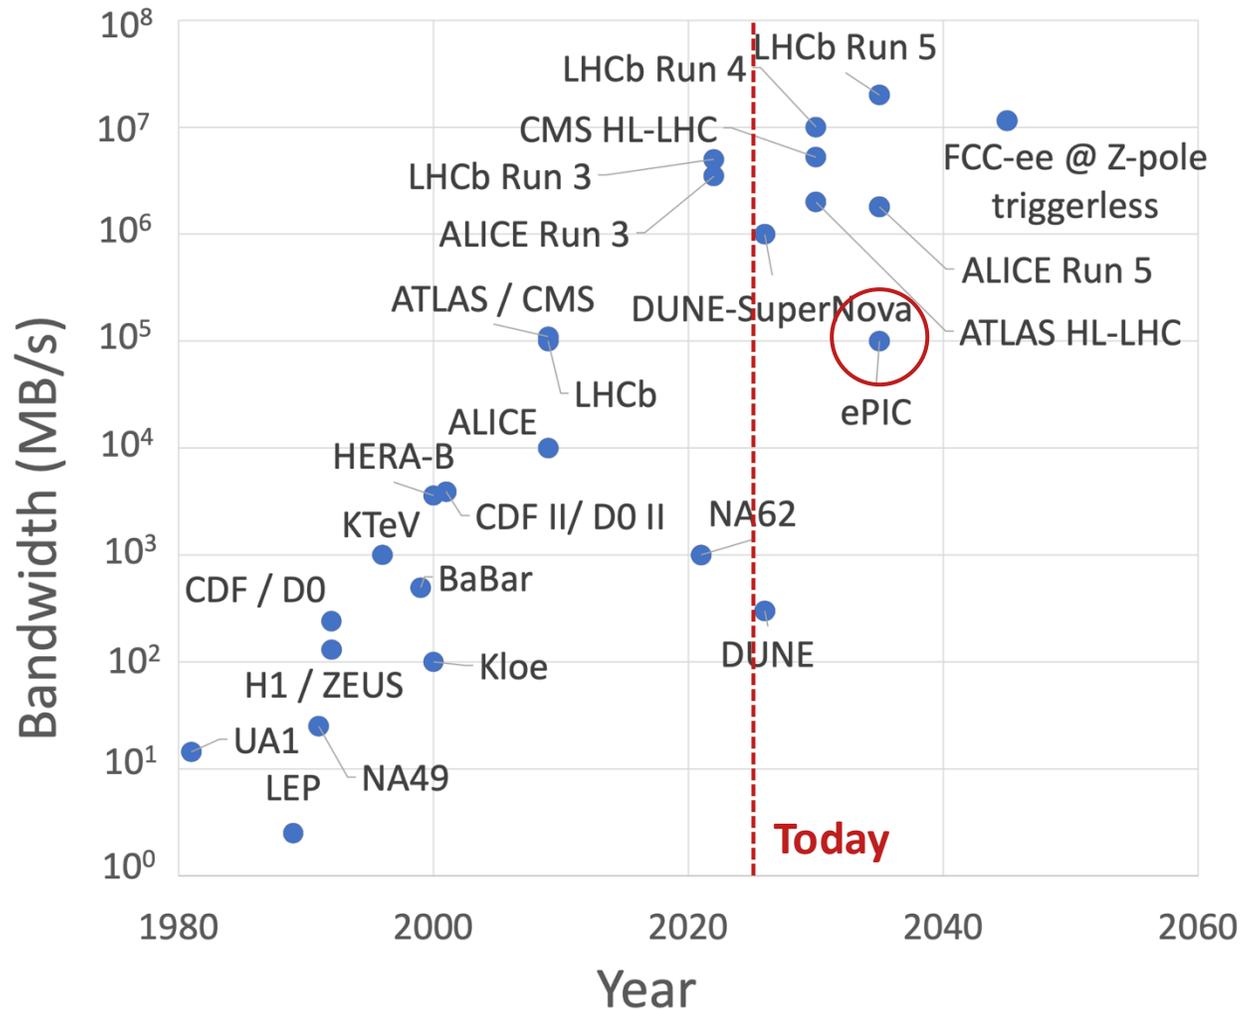
- **Maximize Science:** Capture every collision signal, including background.
 - **High-precision measurements:** Control of systematic uncertainties is critical.
 - Event selection using all available detector data for **holistic reconstruction**:
 - **Eliminate trigger bias** and provide accurate estimation 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).

ePIC Within the Global Particle Physics Experiments Landscape



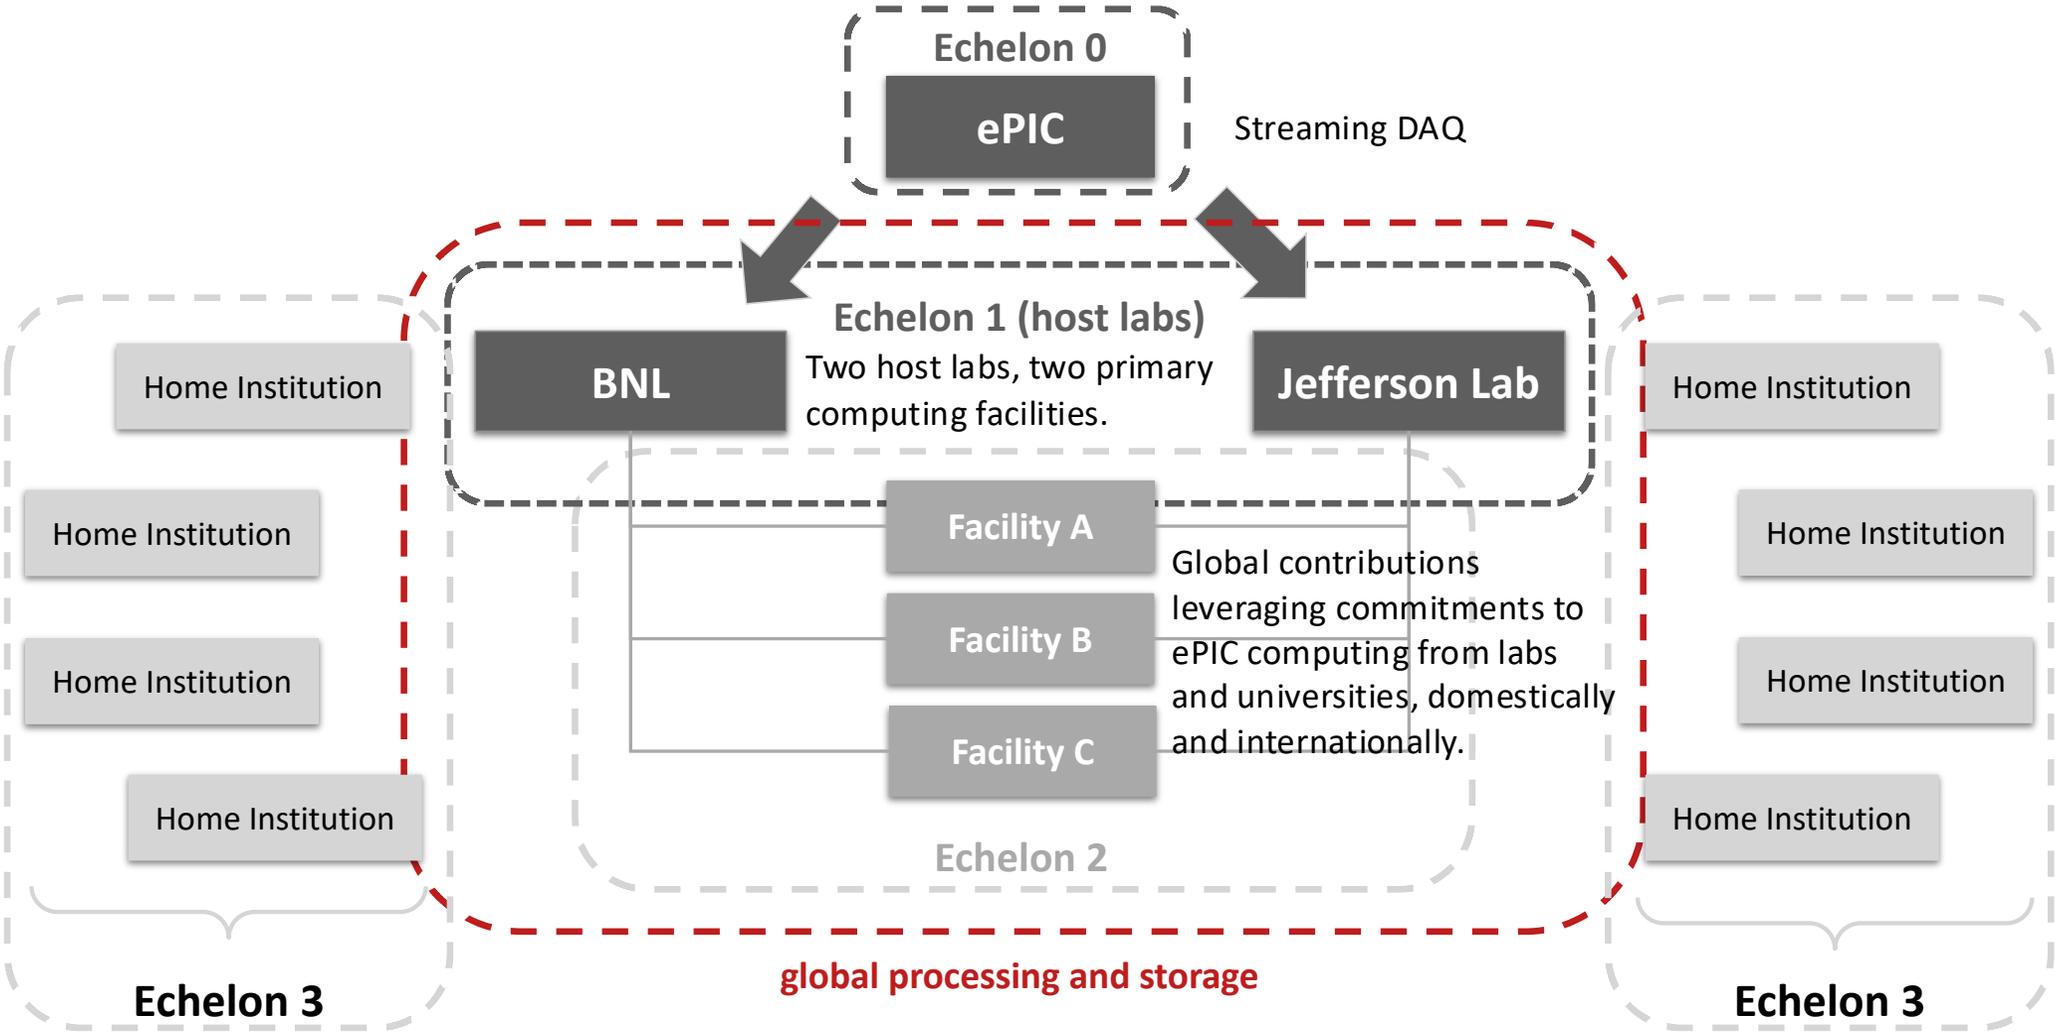
Streaming Readout

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/>

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 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.

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.

Topics for Discussion

- **E0 to E1 Latency**
 - Monitoring results available within 10s.
 - Additional requirements under discussion.
- **Network Resilience and Redundancy**
 - Reliability of E0-E1 data transfer.
 - Streaming continuity and failover strategy:
 - E1 sites are symmetric peers.

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.

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.
- **Testbeds** 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.

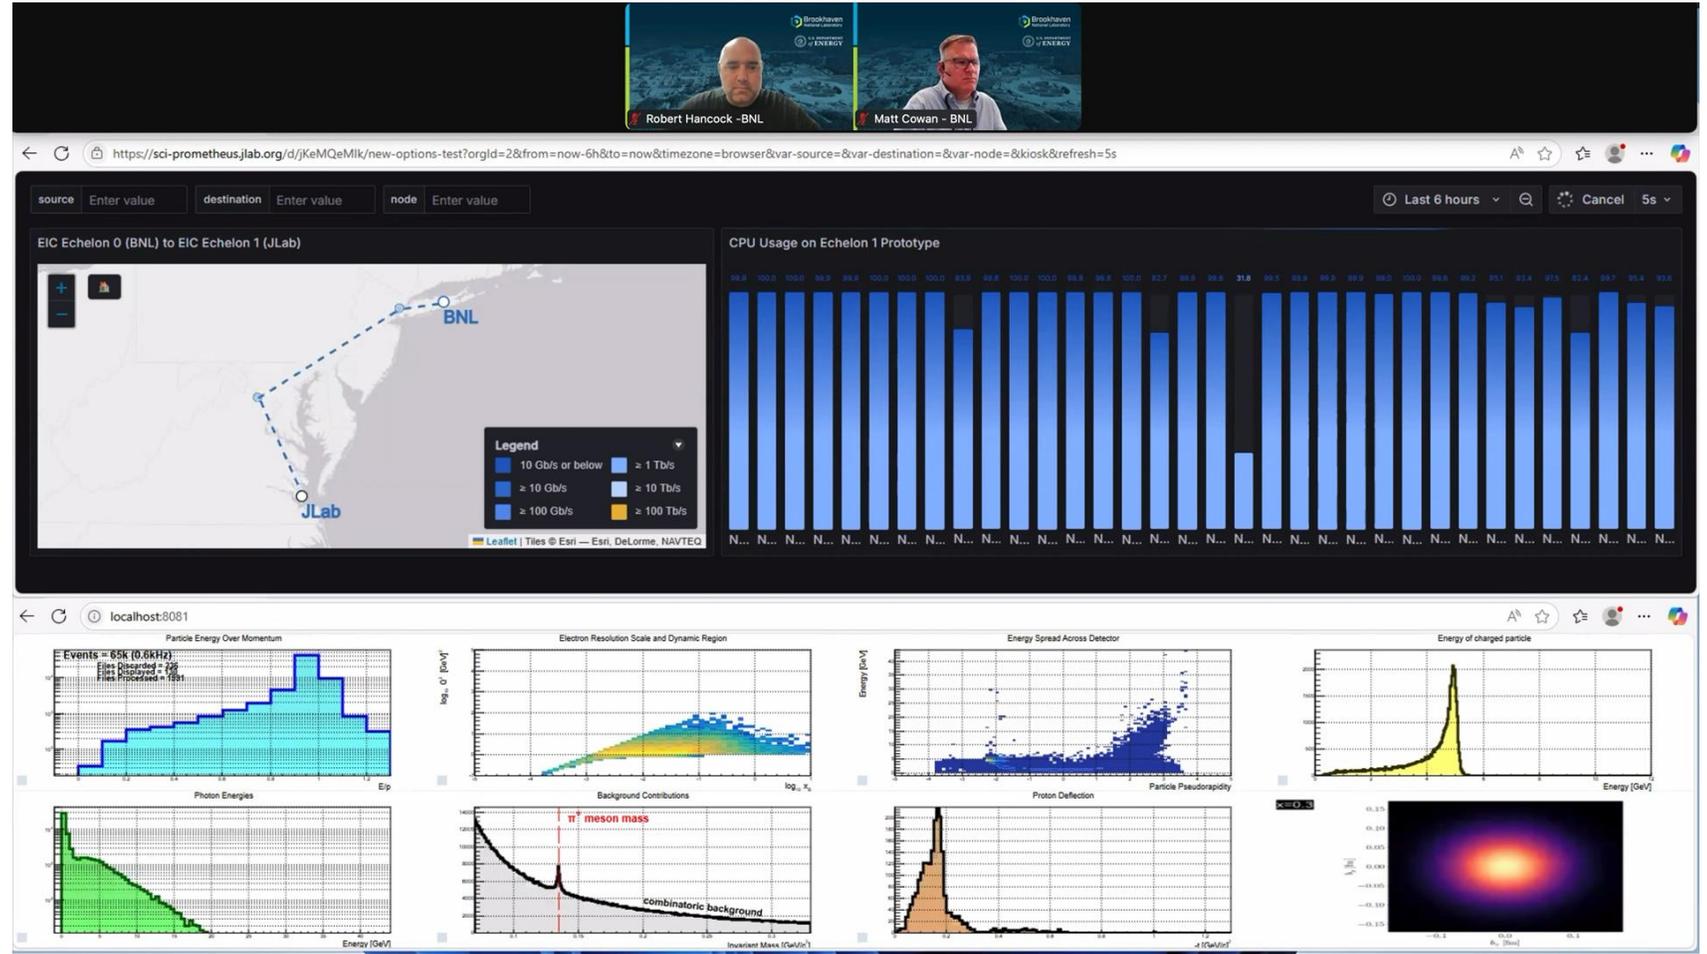
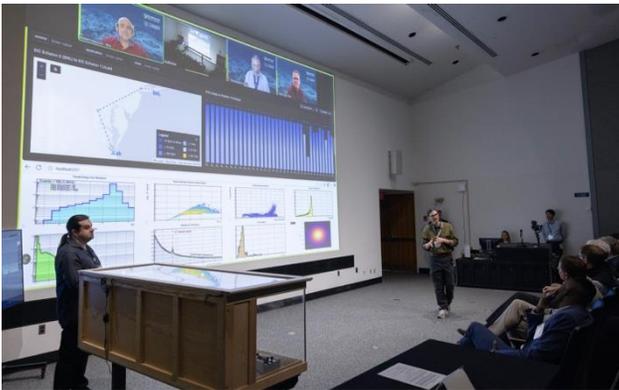
Streaming analysis: Demonstrate simulation data production streaming to E2 site.

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

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.