Markus Diefenthaler (Jefferson Lab)

Compute Resource Estimates

ePIC Streaming Computing Model

The ePIC Streaming Computing Model

Marco Battaglieri[1], Wouter Deconinck[2], Markus Diefenthaler[3], Jin Huang[4], Sylvester Joosten[5], Jeferey 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.

Abstract

This document provides a current view of the ePIC Streaming Computing Model. With datataking a decade in the future, the majority of the content should be seen largely as a proposed plan. The primary drivers for the document at this time are to establish a common understanding within the ePIC Collaboration on the streaming computing model, to provide input to the October 2023 ePIC Software & Computing review, and to the December 2023 EIC Resource Review Board meeting. The material should be regarded as a snapshot of an evolving document.

Report: Initial version of a plan set to develop over the next decade.

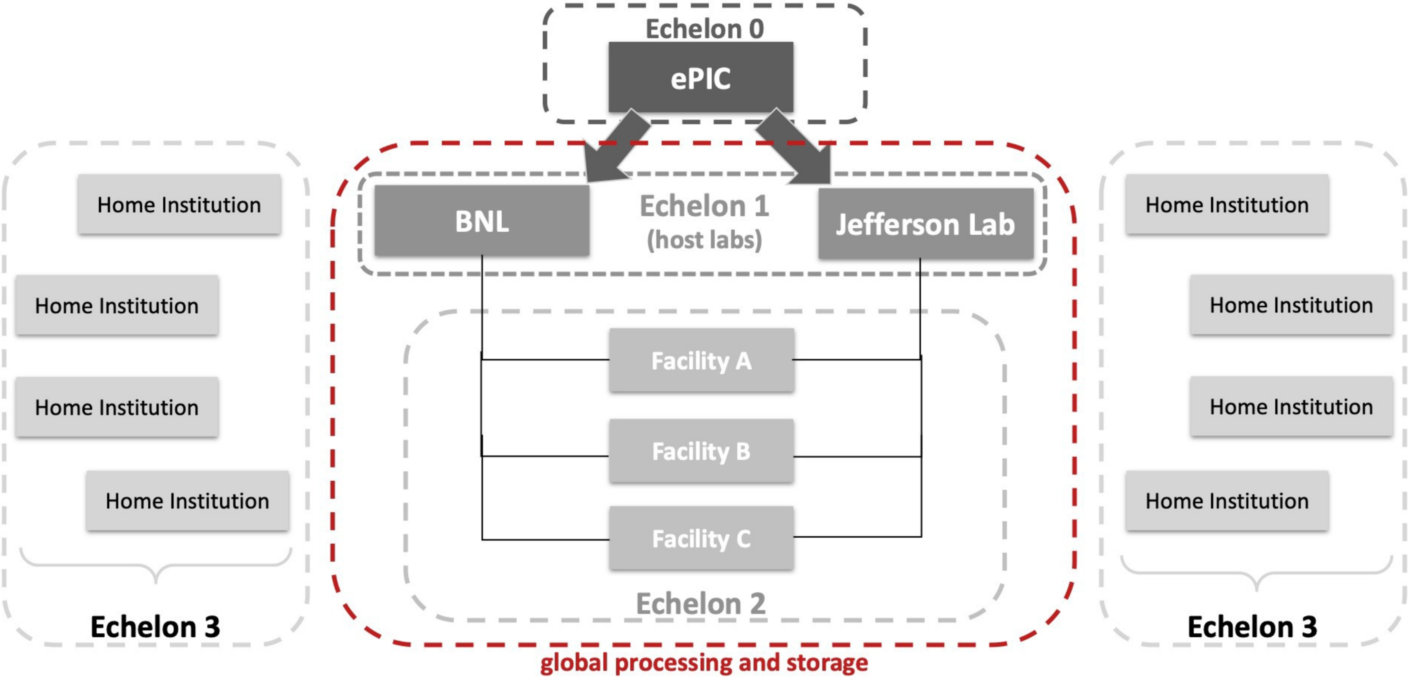

Echelon 0: ePIC experiment.

Echelon 1: Crucial and innovative partnership between host labs.

Echelon 2: Global contributions.

Echelon 3: Full support of the analysis community.

Towards a Quantitative Computing Model

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		✓	✓	

ToDo: Estimate compute resources for each use case.

Computing Scale

Check numbers.

Based on the number of electronic channels in the detector, and the occupancy you expect, what is the expected frame size?

- The streaming data is transferred in frames, collecting all data of **0.6ms**.
- The frame size based on our current detector readout design is **10MB** when running at peak luminosity and in standard operating conditions.

How many events do we expect to record and simulate per year, respectively?

- Assuming a **50% up-time for ½ year**, we will record **15.5 billion frames in a year**.
- The event rate at peak luminosity is **500kHz**, which gives roughly **4 x 10^{12} events** (**60% background**, **40% bunch crossing related**):
  - This will of course be much lower at start of operations, where the luminosity will be lower (but relatively speaking the background rate is expected to be higher).
  - The actual physics events related to key EIC observables is only a very small fraction of the total physics bunch crossings. The **expected number of DIS events / physics event of interest for one year of running at peak luminosity is ~ 10^{10}**. This is the number that drives our **simulation needs**, and we expect to simulate **10x events for each event of interest, yielding O(10^{11})** simulated events. While considerable (~ 60k core years on today's hardware), this should be a realistic target in a decade.

How many core-seconds on a typical modern machine does our reconstruction and simulation take today, respectively?

- Our current simulations of background embedded events take **~17s for simulation** and **~ 2s for reconstruction**, per event.

Use Case: First Full Reconstruction

	Low	High	
Luminosity	1.00E+33	1.00E+34	cm^-2s^-1
Weeks of running	26	26	weeks
Operation Efficiency	50%	50%	
Data Rate to Storage	10.0	100	Gbps
Raw Data Storage (no duplicates)	16	157	PB
Data Productions	1.6	15.7	PB
Total Storage (no duplicates)	17.3	173.0	PB
Reconstruction time / core	1	1	s/event
Reconstruction iterations	3	3	
Total reconstruction time / core	3	3	s/event
Event size	33	33	kB
Number of events produced	477	4765	Billion Events
CPU-core hours	397	3971	Mcore-hrs
Cores needed to process in 26 weeks	91	909	k-cores