ePIC Software & Computing Report

Markus Diefenthaler (Jefferson Lab) for the ePIC Collaboration

Our Philosophy

- We focus on **modern scientific software & computing practices** to ensure the **long-term success of the EIC scientific program** throughout all CD milestones.
  - Strong emphasis on modular, orthogonal tools.
  - Integration with HTC/HPC, CI workflows, and enable use of standard data science toolkits.
- We **leverage cutting edge sustainable community software** where appropriate, **avoiding the "not invented here" syndrome**.
  - Can build our software on top of a mature, well-supported, and actively developed software stack by using modern community tools, e.g. from CERN, the HPC community, and the data science community.
  - Actively collaborate with external software projects, while externalizing some support burden to external projects.
- We embrace these practices today to avoid starting our journey to EIC with technical debt.
- **We are writing software for the future, not the lowest common denominator of the past!**

ePIC

ePIC Software & Computing Organization

Spokesperson's Office

Detector | **Software & Computing** | Physics

Guiding Principles:

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

Software & 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)
Sylvester Joosten (ANL)

Development Working Groups:

- Physics and Detector Simulation
- Reconstruction
- *Analysis Tools*

Deputy Coordinator (Infrastructure)
Torre Wenous (BNL)

Infrastructure Working Groups:

- Streaming Computing Model
- *Multi-Architecture Computing*
- *Distributed Computing*

Since May: Monthly Simulation Productions

Operations | Development | Infrastructure

Objectives

1. Achieve **continuous deployment** of the software used for detector and physics simulations.
2. Ensure **regular updates** of simulation productions for detector and physics studies in preparation for the TDR (and all other CD milestones).
3. Implement **timely validation and quality control** for simulation productions on datasets that require substantial time and resources.

Train: Major central campaign at a fixed (monthly) schedule.

Charter: Special interest runs for working groups.

Taxi: Bespoke runs for individual users.

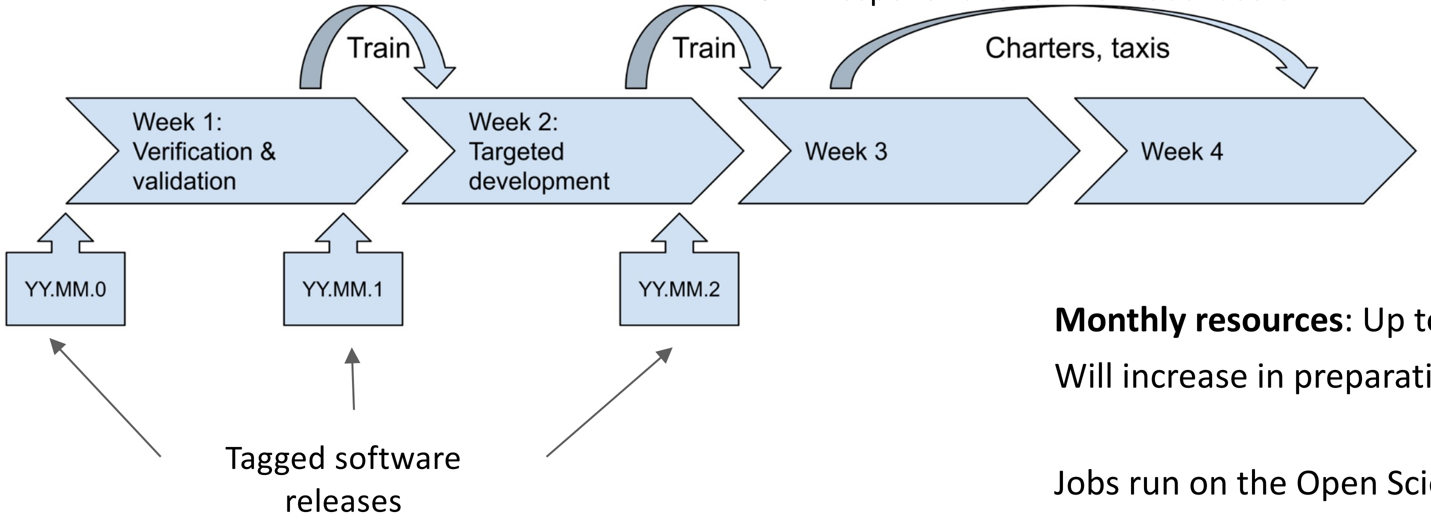

Monthly resources: Up to 10k continuously used cores (about 500 core-years).

Will increase in preparation of TDR.

Jobs run on the Open Science Grid (OSG), with output stored at host labs.

Production WG Conveners

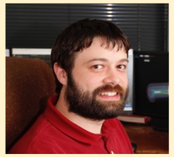

Thomas Britton (Jefferson Lab)

Sakib Rahman (U. Manitoba)

User Learning

Onboarding and training:

- Delivered **online tutorials** (with up to 800 views!).
- Established very active **Helpdesk** on Mattermost.

In-person tutorials at the January collaboration meeting, covering five key topics:

1. **Collaborative Development**

   Eic-shell Easy to get started locally… in only 1 line!

   ```
   curl -L get.epic-eic.org | bash
   ```

   Based on container images, the same images are used for simulation campaigns.

2. **Working with Simulations in Python or ROOT**
3. **Developing Detector Simulations**
4. **Developing Reconstruction Algorithms**
5. **Developing Benchmarks by Validation WG**

User Learning WG Conveners

Kolja Kauder (BNL)

Holly Szumila-Vance (Jefferson Lab)

Validation WG Conveners

Torri Jeske (Jefferson Lab)

Dmitry Kalinkin (U. Kentucky)

Open, Collaborative Development

Enable Access Without Restrictions

- ePIC collaboration members include over 170 institutions worldwide.
- **Data publicly available** at host labs.
- **Simple, flat data structures** (i.e. could be a csv), stored as ROOT files:
    - Straightforward analysis in ROOT or Python without the need for data structure libraries.
    - Encourage collaboration with computer, data, and other scientists outside of NP and HEP.
- **Public software repositories** on GitHub.

Encourage Upstream Contributions

- Requirements of well-formed HepMC as input has resulted in real improvements to multiple MCEGs used by EIC community.
- Various upstream contributions to software packages, e.g., ACTS, DD4hep, Spack, uproot.

Encourage Social Coding

- CI platform provides the incentive for developers to commit code frequently: achieving data management and analysis preservation goals.
- Pull request reviews to ensure higher quality code and build developer skills.

Community Building

Regular workshops to drive forward priority targets and provide an avenue for new collaboration members to actively engage.

ePIC AI Strategy

AI is key part of all Software & Computing WGs in ePIC:

Focus on AI Development as Part of Simulation Campaigns:
- Integration of AI methods in monthly simulation campaigns as measure of AI progress.
- Ongoing work on centralization of training, management of model parameters, and workflow integration.
- **Strategic Development**: Emphasis on algorithms for fast calibrations and PID.

Collaborative Efforts:
- Knowledge transfer with NHEP experiments on AI integration in production workflows.
- Introduction of an AI challenge at the forthcoming collaboration meeting
  - Results will be showcased at CERN meeting in April 2024.

Future Directions:
- Data and analysis preservation for AI approaches.
- Distributed learning for the distributed streaming computing model of ePIC.
- Work with the theory community in NHEP on ways to advance data analysis and interpretation using AI.

Physics and Detector Simulation WG Conveners

Kolja Kauder (BNL)

Chao Peng (ANL)

Reconstruction Framework & Algorithms WG Conveners

Derek Anderson (Iowa State)

Shujie Li (LBL)

Streaming Computing Model WG

Operations | Development | **Infrastructure**

- Works in synergy with Electronics and DAQ WG:
  - And Streaming Readout workshop series: Streaming Readout XII or XIII will be in Japan.
- Defined requirements and high-level design for a **computing model** that enables **rapid data processing for physics analyses** in 11 WG meetings since July.
  - Synergies with SPADI-Alliance. 
- Started documenting a streaming computing model that can be redefined further with international partners.
- Initial version of the ePIC Streaming Computing Model has been presented in recent ePIC Software & Computing Review:

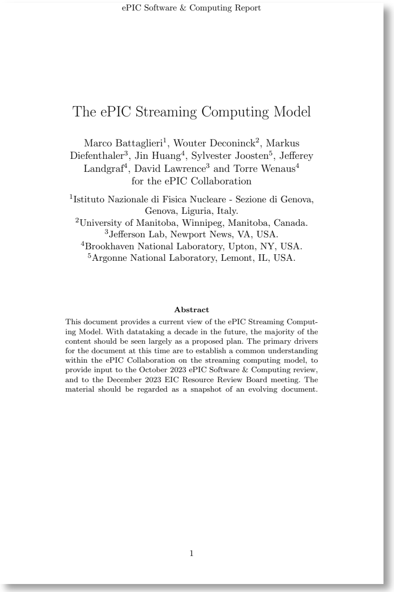

ePIC Software & Computing Report

The ePIC Streaming Computing Model

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

1

Direct Link

Streaming Computing Model WG Conveners


Marco Battaglieri (INFN Genova)

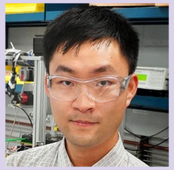

Jin Huang (BNL)

Jeff Landgraf (BNL)

Compute-Detector Integration to Maximize Science

Broad ePIC Science Program:
- Plethora of observables, with less distinct topologies where every event is significant.
- High-precision measurements: Reducing systematic uncertainties of paramount importance.

Streaming Readout Capability:
- Event selection using all available detector data for **holistic reconstruction**:
  - **Eliminate trigger bias** and provide accurate estimation of uncertainties during event selection.
- **Capture every collision signal**, including background:
  - Ideal for broad ePIC Science Program and multi-purpose detector.
  - Ideal to **reduce background** and related systematic uncertainties in measurement.

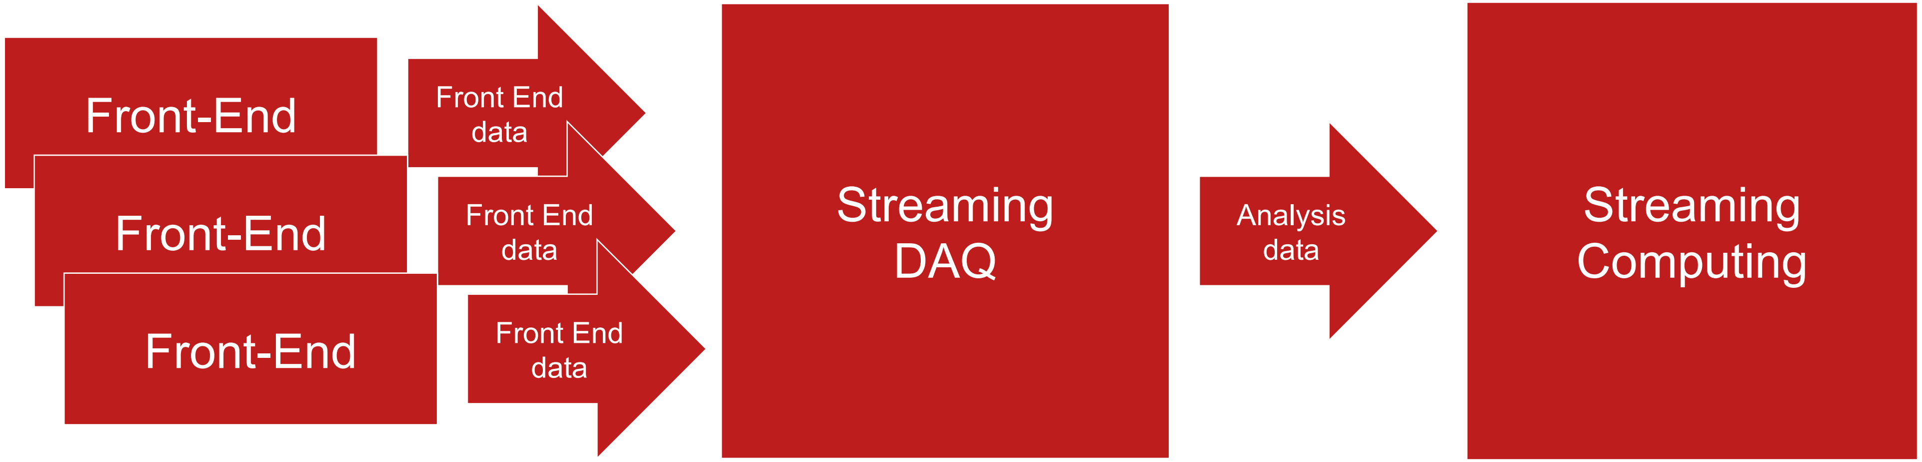

Data Flow: **400 Tbps** → **2 Tbps** → **0.1 Tbps**

Compute-Detector Integration to Accelerate Science

- **Problem** Data for physics analyses and the resulting publications available after O(1year) due to complexity of NP experiments (and their organization).
  - Alignment and calibration of detector as well as reconstruction and validation of events time-consuming.
- **Goal** **Rapid turnaround of 2-3 weeks for data for physics analyses**.
  - Timeline driven by calibrations.
- **Solution** Compute-detector integration using:

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

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

Heterogeneous computing for acceleration.

ePIC Streaming Computing Model

Four Tiers:

Echelon 0: ePIC Experiment

Echelon 1: Host Labs

Echelon 2: Global processing and data facilities, includes HPC and HTC resources.

Echelon 3: Home institute computing

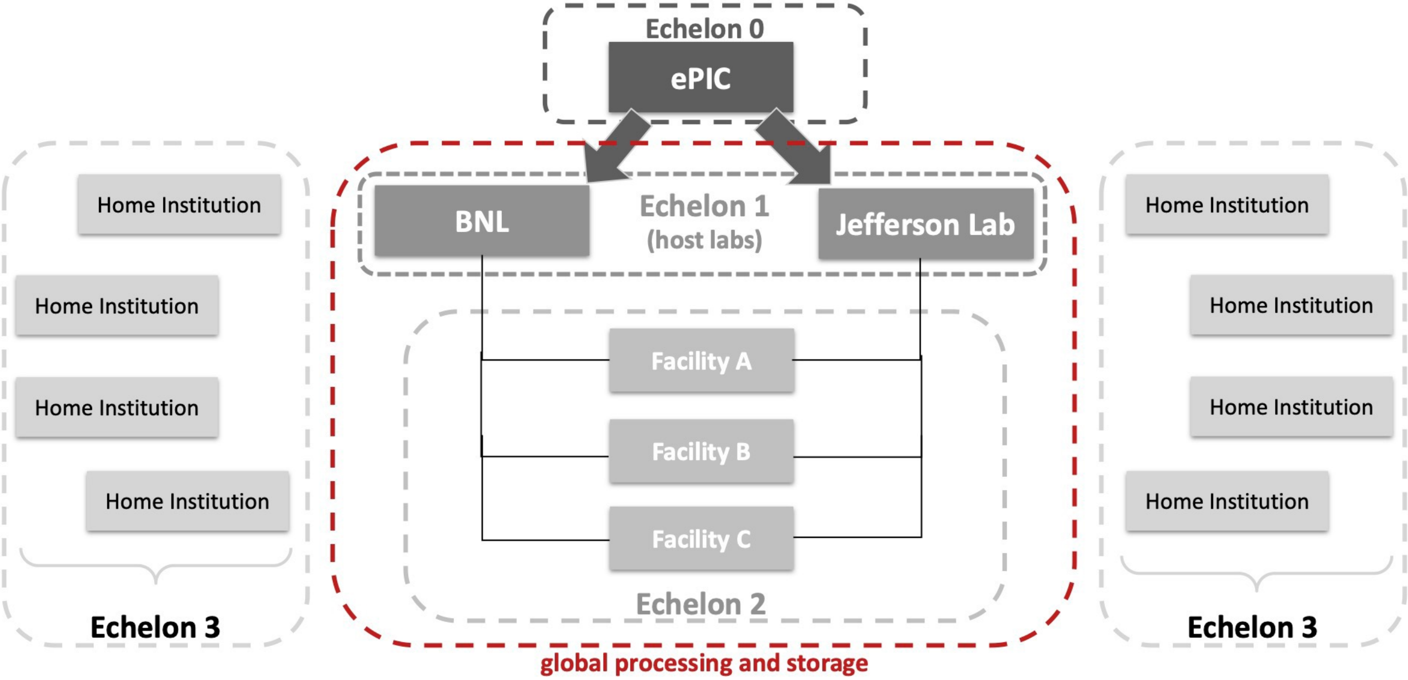

Successful review of ePIC Software & Computing by EIC Computing and Software Advisory Committee (ECSAC).

From review close out: *"The ePIC collaboration is doing all the right things to make progress towards exploiting in-kind contributions to Computing & Software for ePIC."*

Use Cases

Echelon 0: ePIC Experiment
Echelon 1: Host Labs
Echelon 2: Global Processing and Data Facilities
Echelon 3: Home Institute Computing

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		✓	✓	

Echelon 2

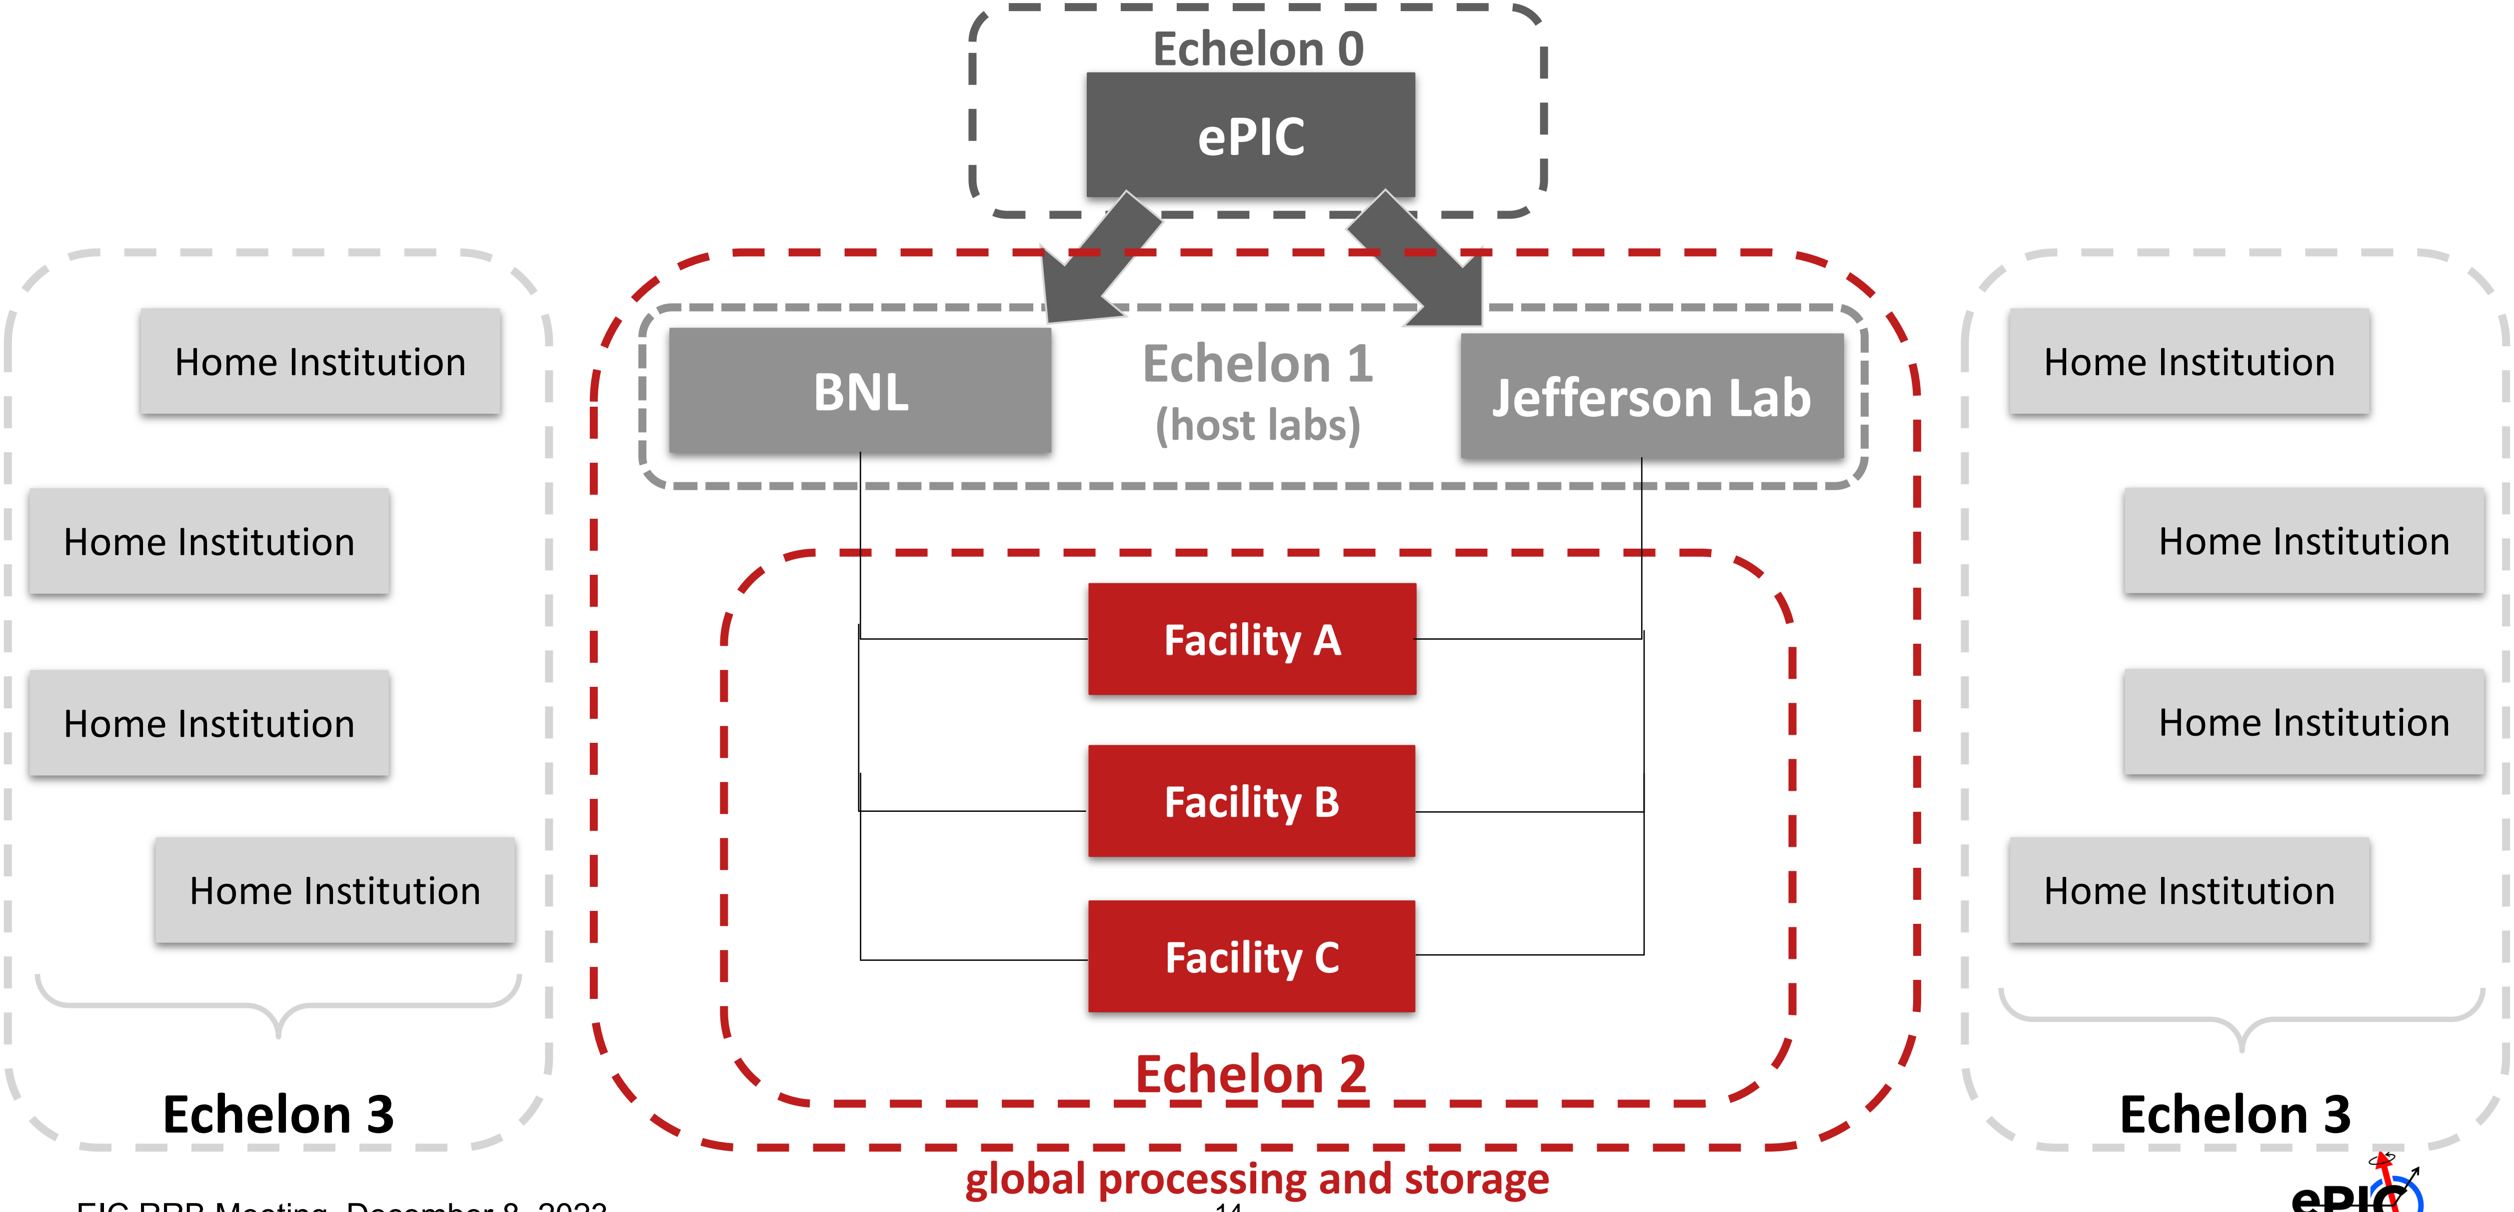
Echelon 0
ePIC
Echelon 1
(host labs)
BNL
Jefferson Lab
Facility A
Facility B
Facility C
Echelon 2
Home Institution
Home Institution
Home Institution
Home Institution
Echelon 3
Home Institution
Home Institution
Home Institution
Home Institution
Echelon 3
global processing and storage

Echelon 2: Global ePIC Computing

ePIC is an international collaboration and so is its computing:

- **Echelon 2** includes **global resources contributed by collaborating institutions**.
- Achieving scientific goals relies on effectively using Echelon 2's resources.
- Design of computing model aims for **effective integration and management**.
- ePIC supports the **ECSJI** formation to oversee and coordinate the complex computing fabric of the EIC.

International computing contributions are essential.

Lessons Learned from LHC

- Echelon 2 resources need to **assure technical compatibility with the ePIC computing model**.
- **ePIC commits to facility integration** and robust testing/validation, includes monitoring and diagnostics.

Connectivity

- Echelon 2 sites connect equally to both Echelon 1 Host Labs.
- Connection is ultimately through ESnet network backbone.
- **Interconnected Mesh**:
  - Echelon 2 sites also interconnect based on their network environment.
  - Lessons from LHC show this interconnected mesh model is more effective than a hierarchical model.

Current Estimate for Compute Resources

Streaming DAQ sends data in **1ms time frames**.
Each time frame corresponds to 10MB of data.

Based on our current detector readout design and when running at peak luminosity and in standard operating conditions.
40% of data bunch crossing related, 60% background.

In a year, we will record 15.5 billion frames.

Assuming a 50% up-time for 6 months.

Number of expected events (assuming a 50% up-time for 6 months):

- The event rate at peak luminosity is 500kHz, which gives roughly 4 x 10^{12} events.
  - Lower at start of operations, where the luminosity will be lower (but relatively speaking background rate is expected to be higher).
- The expected number of physics events of interest for one year of running at peak luminosity is ~ 10^{10}.
  - The actual physics events is only a very small fraction of the total physics bunch crossings.

Number of simulation events:

- 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.

Core-seconds for simulation and reconstruction (on a typical modern machine):

- Our current simulations including background take ~17s for simulation and ~ 2s for reconstruction, per event.
  - Simulation and reconstruction on event level only.
- **Unknown**: How much this will change once changing to streaming data processing?
- **Priority target for TDR**: Prototype of event reconstruction from realistic frames.

International Contributions

- **From the review close out**:
  "There are clearly very significant opportunities in in-kind computing infrastructure contributions."
- Canada, Italy, and the United Kingdom are engaged as a proof of concept in this context.
- Computing centers of these countries were already included in large-scale simulation efforts for the EIC.

Digital Research Alliance of Canada | Alliance de recherche numérique du Canada

INFN

UKRI Science and Technology Facilities Council

Collaboration Interest in Software & Computing

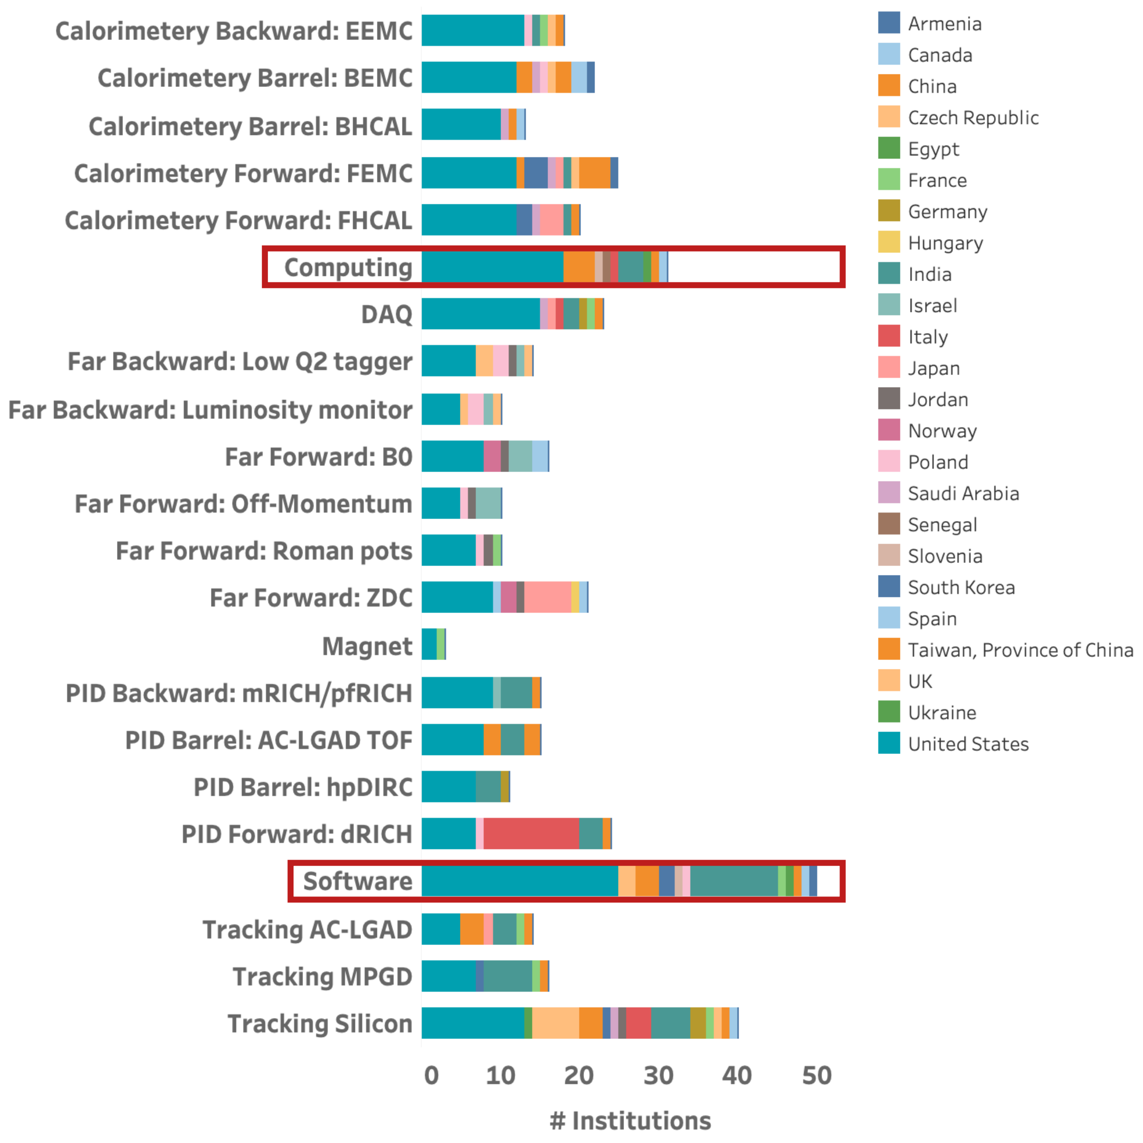

The interest does not yet reflect actual, active participation in Software & Computing.

Key priorities for the TDR include conducting **simulation campaigns**, **validating our software and simulation outputs**, and developing **reconstruction algorithms**.

Additionally, we will focus on **advancing distributed computing capabilities** and will start **integrating resources from international partners** in our resources on the Open Science Grid.

Efforts on these key priorities will benefit substantially from an increase of the workforce.

Further Reading

EIC SOFTWARE: Statement of Principles

1. **We aim to develop a diverse workforce, while also cultivating an environment of equity and inclusivity as well as a culture of belonging.**
2. **We will have an unprecedented compute-detector integration:**
   - We will have a common software stack for online and offline software, including the processing of streamed data and its time-ordered structure.
   - We aim for autonomous alignment and calibration.
   - We aim for a rapid, near-real-time turnaround of the raw data to online and offline productions.
3. **We will leverage heterogeneous computing:**
   - We will enable distributed workflows on the computing resources of the worldwide EIC community, leveraging not only HTC but also HPC systems.
   - EIC software should be able to run on as many systems as possible, while supporting specific system characteristics, e.g., accelerators such as GPUs, where beneficial.
   - 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.
4. **We will aim for user-centered design:**
   - We will enable scientists of all levels worldwide to actively participate in the science program of the EIC, keeping the barriers low for smaller teams.
   - EIC software will run on the systems used by the community, easily.
   - We aim for a modular development paradigm for algorithms and tools without the need for users to interface with the entire software environment.

The ePIC Streaming Computing Model

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

1

Closeout Report of EIC Computing & Software Review

EIC Computing & Software Advisory Panel

ECSAC Members:
Frank Wuerthwein (chair), Mohammad Al-Turany, Pere Mato, Heidi Schellmann, David Brown, Simone Campana, Christoph Paus

ePIC Representatives:
Rosi Reed, Sylvester Joosten, Markus Diefenthaler, Torre Wenaus, Andrea Bressan

ECSJI Representatives:
Amber Boehnlein, Eric Lancon

Host Lab Representatives:
Amber Boehnlein, Eric Lancon,

The EIC Computing & Software Advisory Committee (ECSAC) met with ePIC and ECSJI representatives for 1.5 days at the SURA facilities in Washington D.C. Lab Directors joined us for a session via video. ECSAC received a number of excellent presentations, with plenty of time for detailed Q&A. ECSAC provided overnight questions for further clarifications, that the ePIC team answered.

We thank the organizers of this review, Amber Boehnlein and Eric Lancon for inviting us to hear about this exciting new physics program, for their strong organization of the review, and for their clear instructions and charge. We thank the ePIC representatives and presenters for their time in preparing their material, and their clear presentations. And we thank the **ePIC collaboration** for supporting the presenters, and their responses to the questions from the Committee.

Overall, we think the ePIC computing & software preparations are in excellent shape for this early in the process. In the following we list the charge questions, and provide for each of them findings, comments, and recommendations.

1. **At this stage, approximately ten years prior to data collection, is there a comprehensive and cost-effective long-term plan for the software and computing of the experiment?**

Yes, to the extent that this can be determined at this point.

Findings:

Summary

- The ePIC software stack is a **modern** and **modular toolkit** built from NP/HEP community tools and components from HPC and Data Science; ePIC is an **active** member of the **NP/HEP software and computing community**.

- **Streaming Readout of the ePIC Detector to maximize and accelerate science:**
  - ePIC aims for **rapid turnaround of 2-3 weeks for data for physics analyses**.
  - Timeline driven by calibrations.
- **Four tiers of the ePIC Streaming Computing Model computing fabric:**
  - **Echelon 0**: ePIC experiment and its streaming readout.
  - **Echelon 1**: Crucial and innovative partnership between host labs.
  - **Echelon 2**: Essential global contributions.
  - **Echelon 3**: Full support of the analysis community.

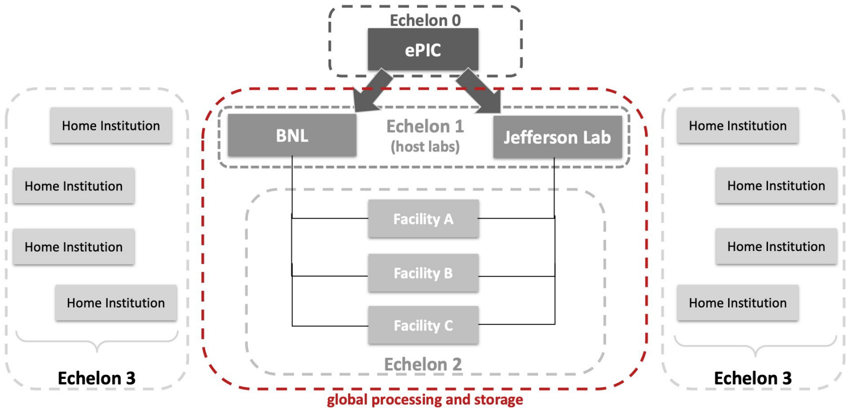

- **High level milestones** ensures that the agile development process is continuously confronted with real world exercising of the software and the developing realization of the computing model:
  - Priority always given to meeting near-term needs. ePIC leverages monthly production campaigns, CI-driven benchmarks, and timeline-based prioritization to ensure timely completion of the simulation studies for the Technical Design Report.
  - Longer range timeline progressively exercising the streaming computing model to deliver for the needs of the CD process, for specific applications, e.g. test beams, for scaling and capability challenges, and ultimately for the phases of datataking.

Backup

Echelon 3

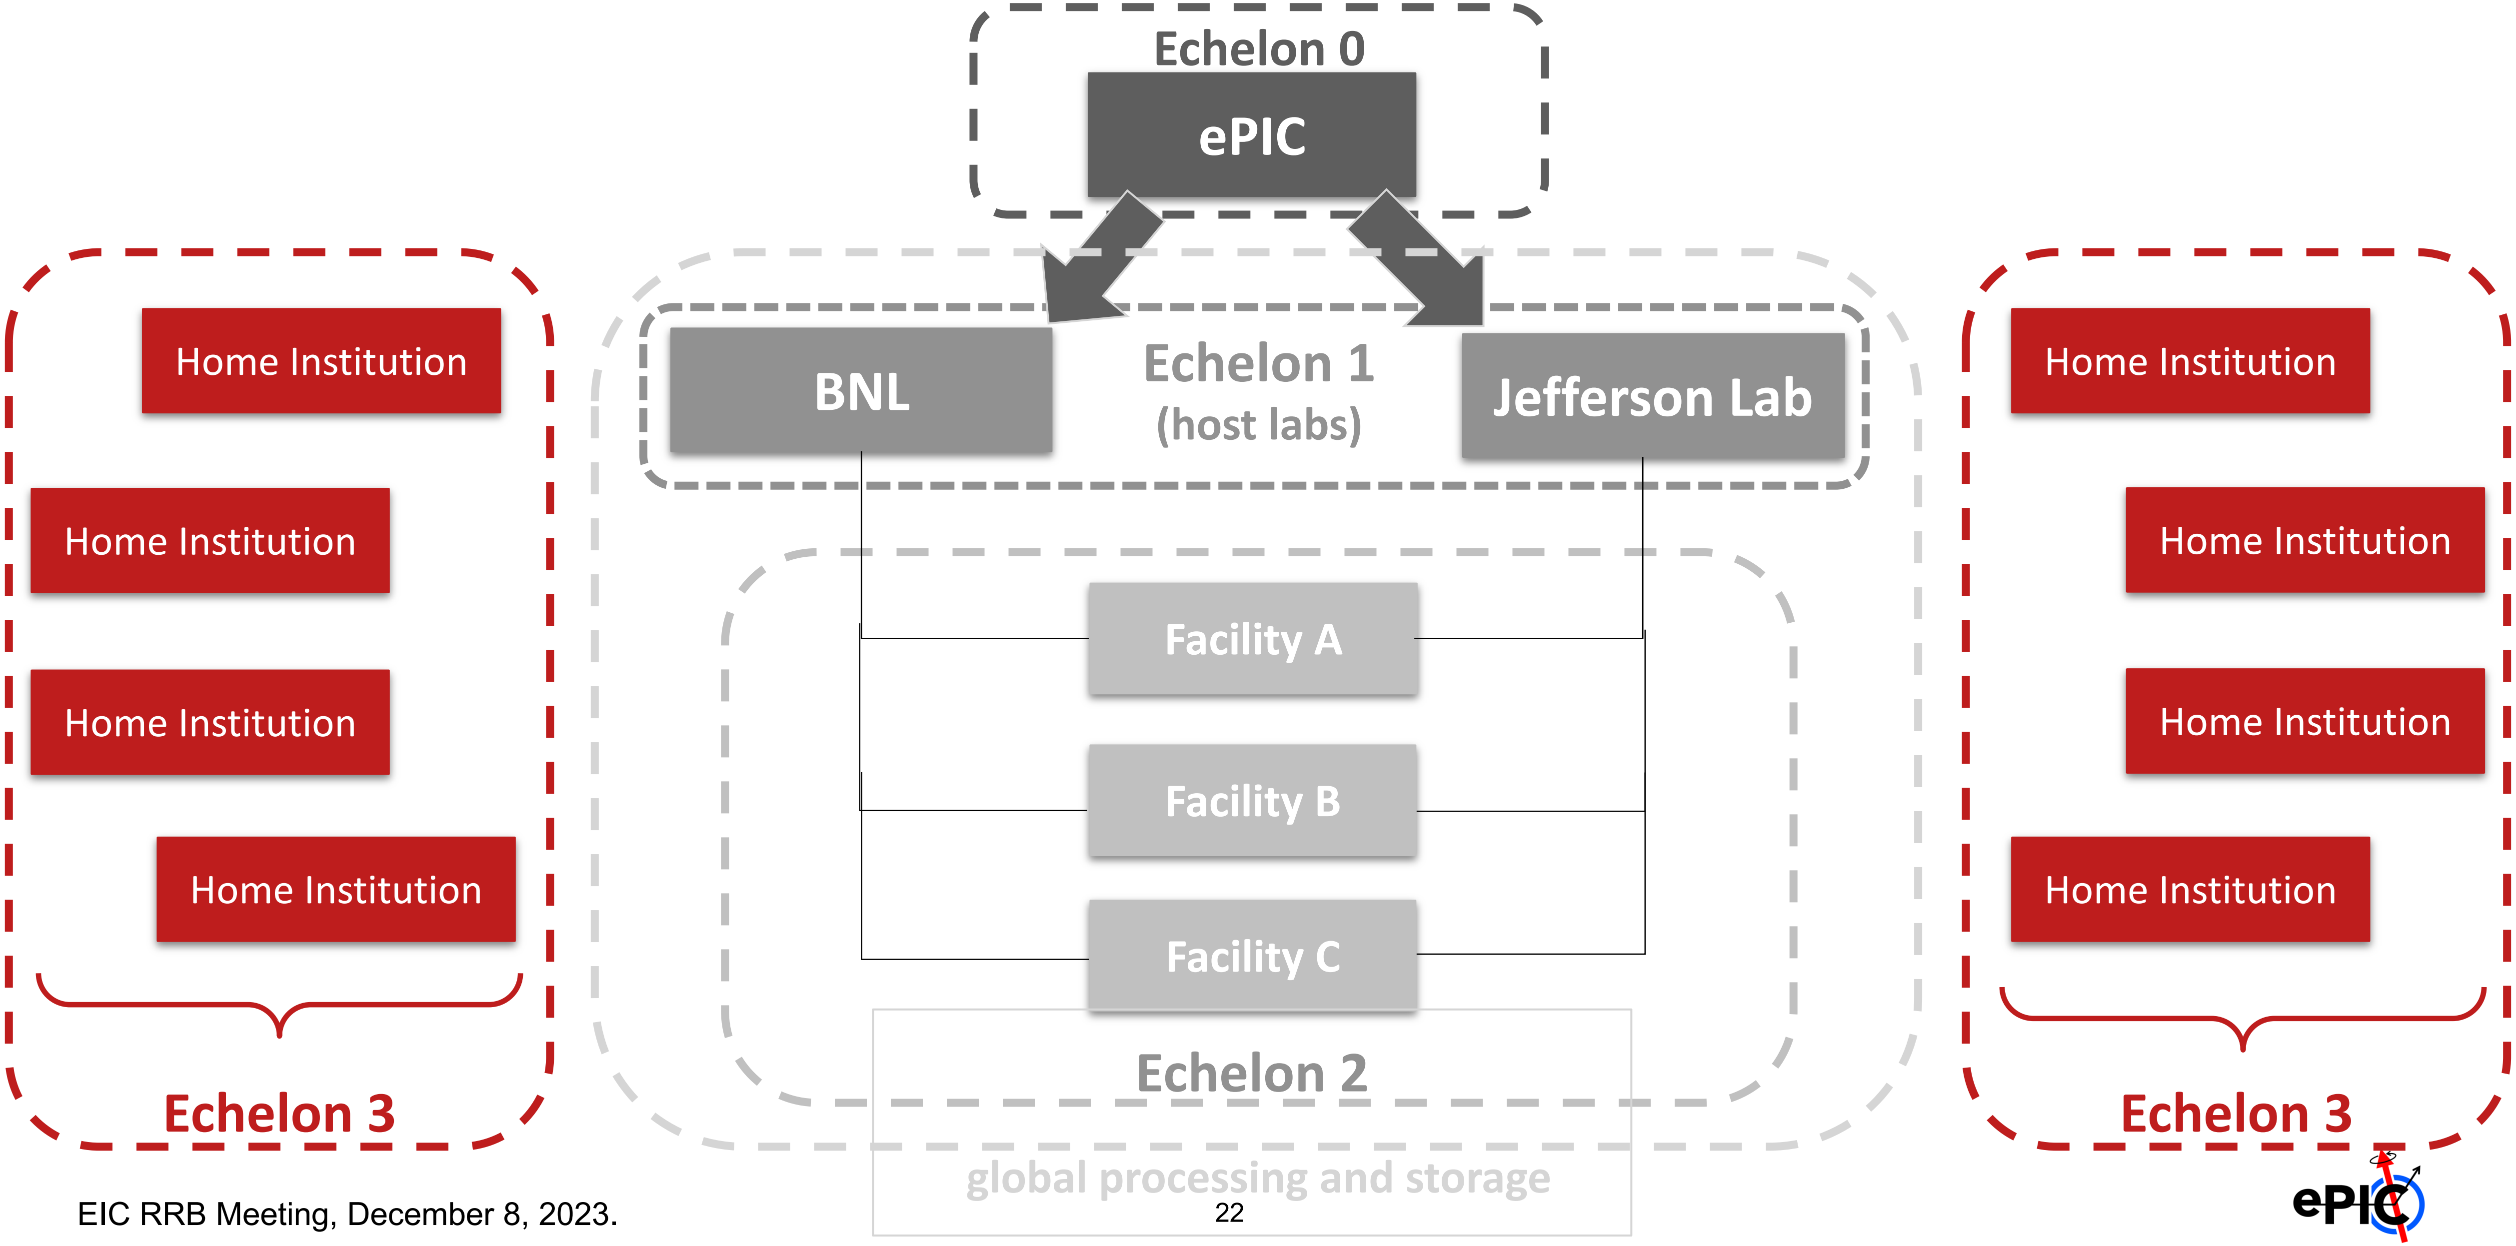
Echelon 0
ePIC
BNL
Echelon 1
(host labs)
Jefferson Lab
Facility A
Facility B
Facility C
Echelon 2
global processing and storage
Home Institution
Home Institution
Home Institution
Home Institution
Echelon 3
Home Institution
Home Institution
Home Institution
Home Institution
Echelon 3

Echelon 3: Home Institute Computing

- **Echelon 3**: Component in the computing model where collaborators directly interact with the computing system:
  - Users can access ePIC computing through various platforms like institutional clusters, work desktops, and personal laptops.
  - The role of Echelon 3 is to serve these diverse use cases.
- **Echelon 3 Resources**:
  - Echelon 3 resources are both global and local to the user, similar to Echelon 2.
  - The resources available in E3 are numerous, diverse, volatile, and often have restrictions on their use.
  - These resources are not intended to be managed as Collaboration resources.
- The collaboration will provide tools, interfaces, connection points, data access mechanisms, and support to make Echelon 3 resources effective for ePIC analysis.

Milestones

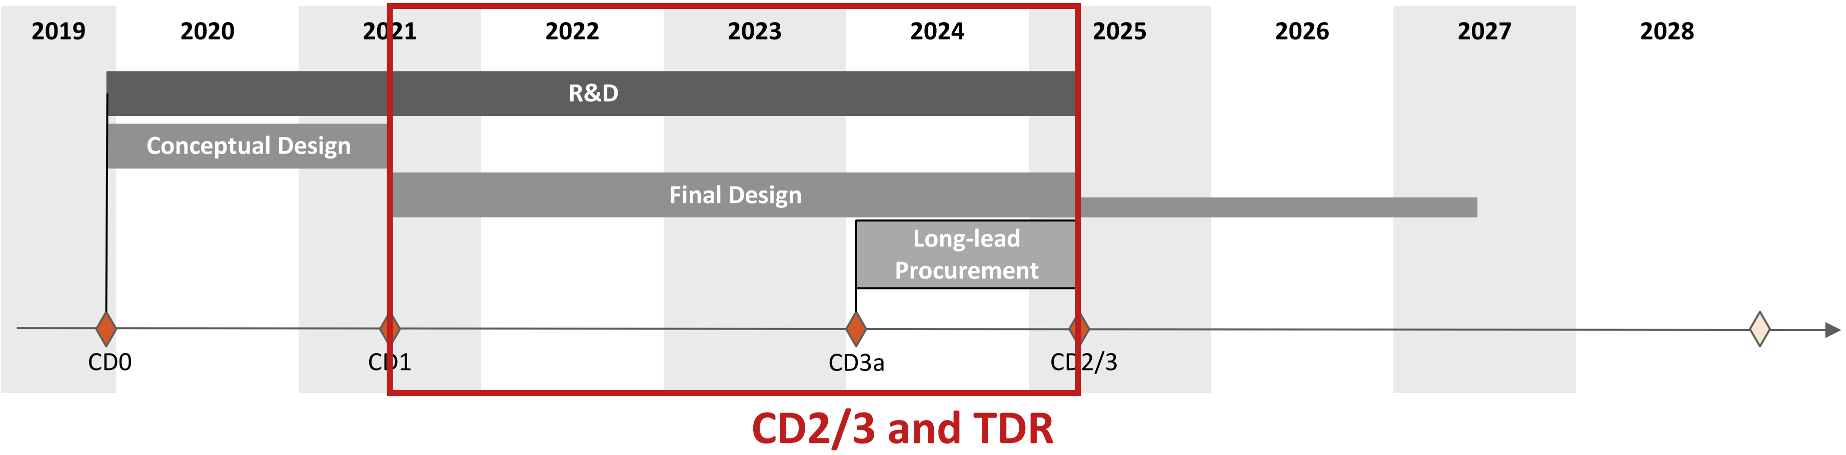

Milestones Prior to CD2/3 and TDR

- Software and simulation readiness for TDR preparation (and subsequent phases of the CD process).
- Provide for each use case detailed estimates on the compute resources; update the networking and storage estimates according to format of streaming data format that is currently being defined.

Milestones

Sec. 9.2.2

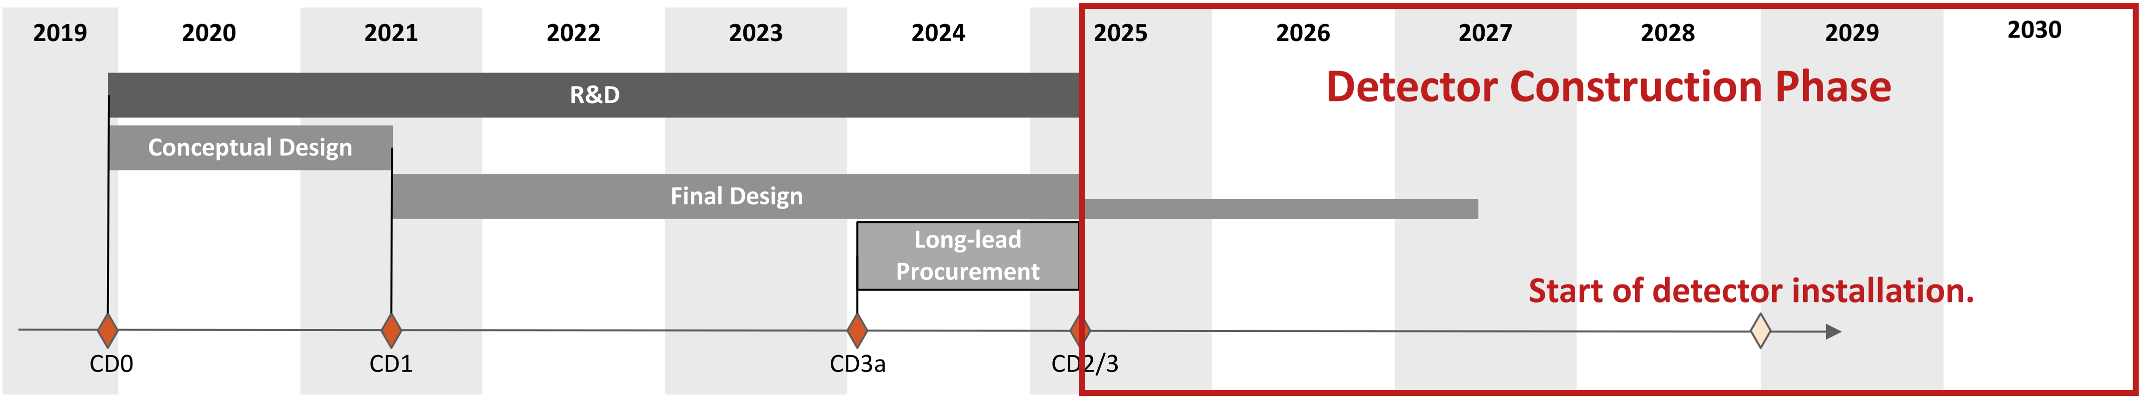

Milestones During Detector Construction Phase

- Provisioning DAQ and software sufficient for test beams, which can serve as small scale real-world testbeds for the developing DAQ and software.
- Streaming challenges exercising the streaming workflows from DAQ through offline reconstruction, and the Echelon 0 and Echelon 1 computing and connectivity.
- 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 ePIC's computing requirements.
- Analysis challenges exercising autonomous alignment and calibrations.
- Analysis challenges exercising end-to-end workflows from (simulated) raw data to exercising the analysis model.

Milestones

Sec. 9.2.3

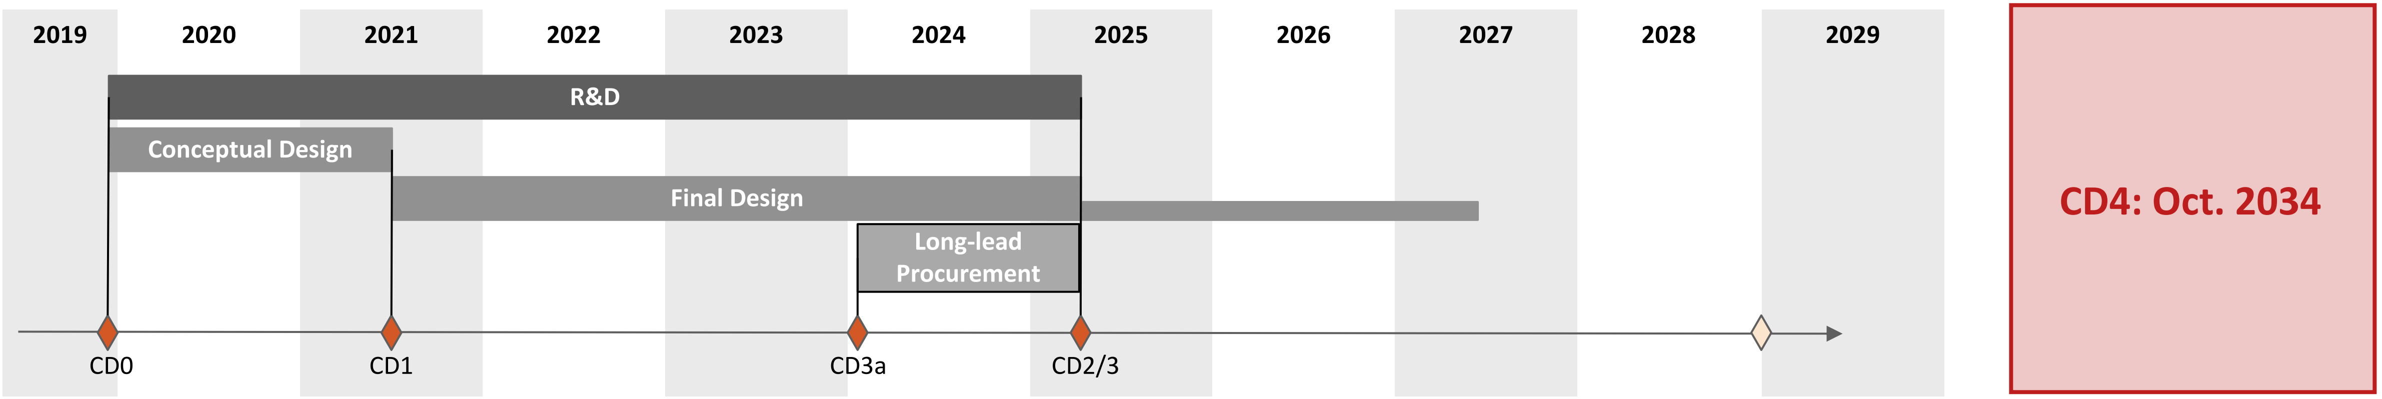

Milestones During Detector Commissioning

- This phase has unique expectations and requirements compared to steady-state operation:
  - Utilization of semi-triggered data-taking modes.
  - Initial calibrations.
  - Gradual extension of first pass processing from Echelon 1 to Echelon 2.
- Careful planning of software & computing efforts and leveraging experience from data and analysis challenges during the detector construction phase essential.

Milestones during Early Datataking Phase

- Simpler and more conservative approaches will be used during initial data-taking phase.
- The ePIC Streaming Computing Model will be gradually deployed and validated.

ePIC AI Strategy

AI is key part of all Software & Computing WGs in ePIC:

Focus on AI Development as Part of Simulation Campaigns:
- Software and computing plan and integration with NHEP community developments successfully reviewed by EIC Computing and Software Advisory Committee.
- Integration of AI methods in monthly simulation campaigns as measure of AI progress.
- Ongoing work on centralization of training, management of model parameters, and workflow integration.
- **Strategic Development**: Emphasis on algorithms for fast calibrations for streaming computing workflows and PID.

Collaborative Efforts:
- Knowledge transfer with NHEP experiments on AI integration in production workflows.
- Introduction of an AI challenge at the forthcoming collaboration meeting
    - Results will be showcased at CERN meeting in April 2024.

Future Directions:
- Data and analysis preservation for AI approaches.
- Distributed learning for the distributed streaming computing model of ePIC.
- Work with the theory community in NHEP on ways to advance data analysis and interpretation using AI.

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 datataking 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.