ePIC Streaming Computing Model

Markus Diefenthaler (Jefferson Lab) for the ePIC Collaboration

Review Charge

This charge covers an assessment of the ePIC computing model in preparation for the December 2023 RRB meeting. The scope of this review also includes the organization of the newly formed ECSJI:

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?**
2. **Are the plans for integrating international partners' contributions adequate at this stage of the project?**
3. Are the plans for software and computing integrated with the HEP/NP community developments, especially given data taking in ten years?
4. Are the resources for software and computing sufficient to deliver the detector conceptional and technical design reports?
5. Are the ECSJI plans to integrate into the software and computing plans of the experiment sufficient?

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

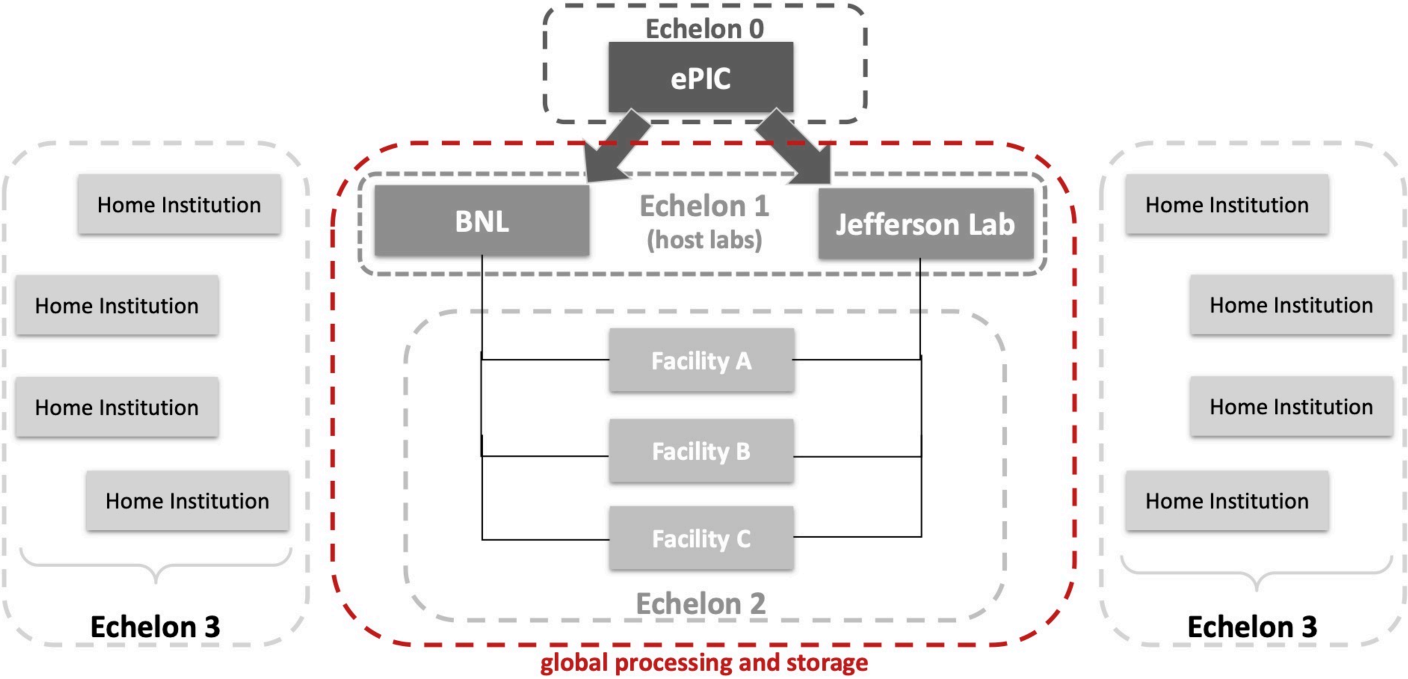

Current View of the Computing Model

In a decade: Start of the ePIC operations

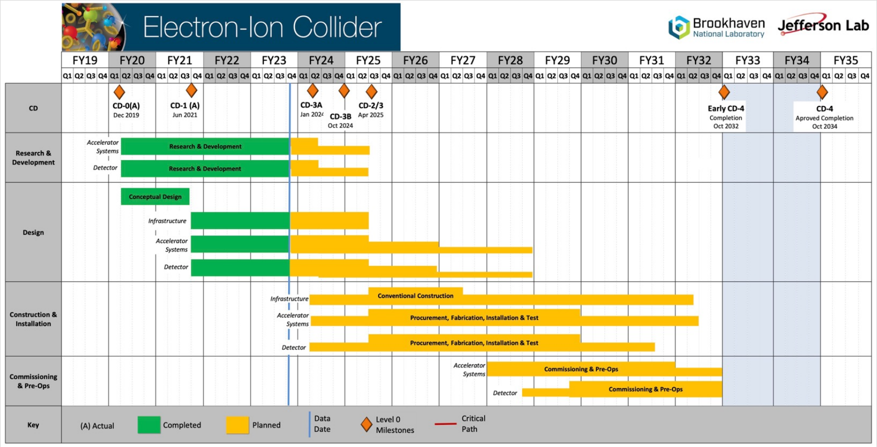

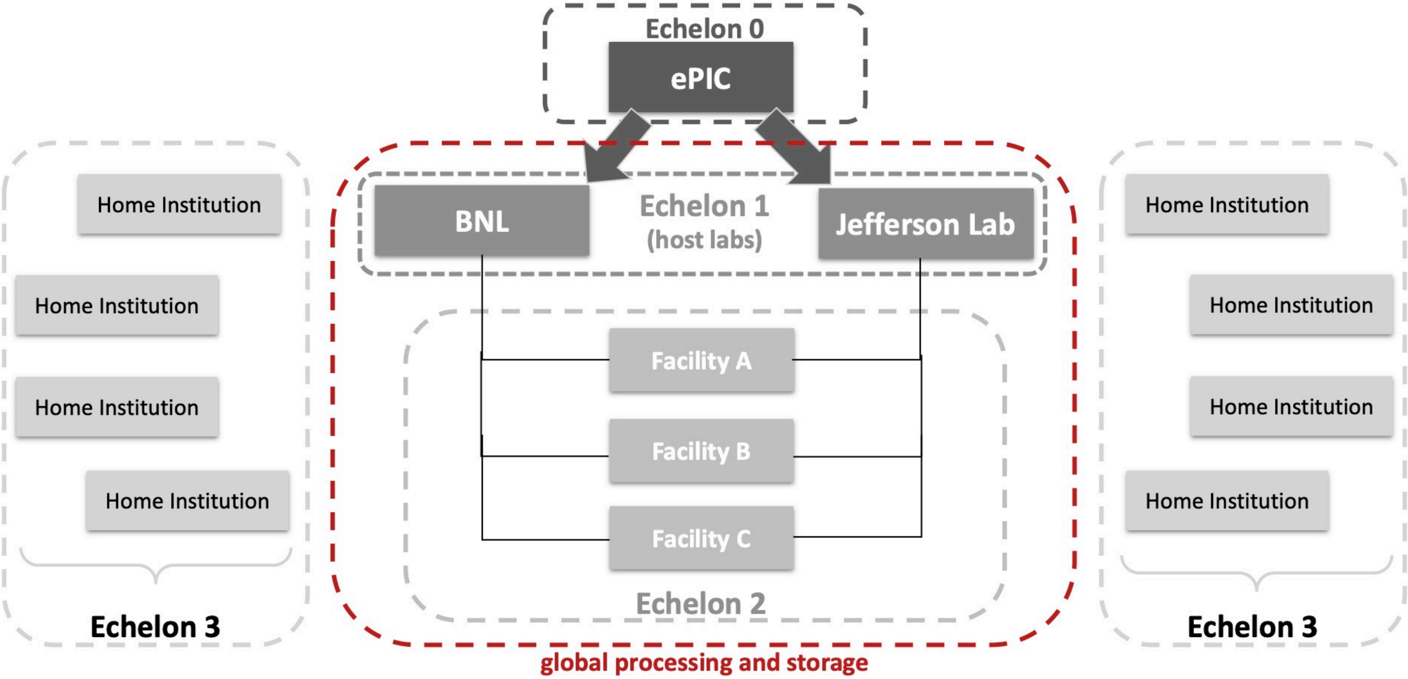

Echelon 0

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

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

Streaming Readout: Trigger-Less Data Acquisition

Definition of Streaming Readout

- Data is digitized at a fixed rate with thresholds and zero suppression applied locally.
- Data is read out in continuous parallel streams that are encoded with information about when and where the data was taken.
- Event building, filtering, monitoring, and other processing is deferred until the data is at rest in tiered storage.

Advantages of Streaming Readout

- Simplification of readout and increased flexibility (no custom trigger hardware and firmware).
- Continuous data flow provides detailed knowledge of backgrounds.
- Streamline workflows and take advantage of other emerging technologies:
    - AI for autonomous experimentation and control.
    - Reconstruction of physics events from holistic detector information.

Optimize Physics Reach

Integrated interaction and detector region (90 m)

Get ~100% acceptance for all final state particles, and measure them with good resolution. All particles count!

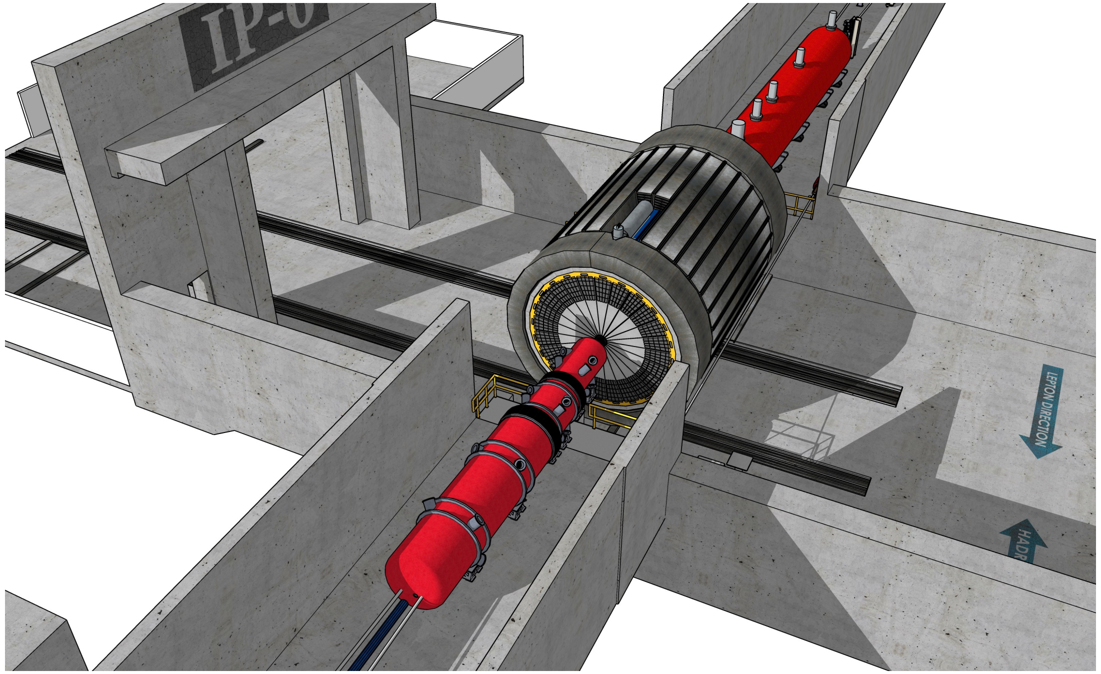

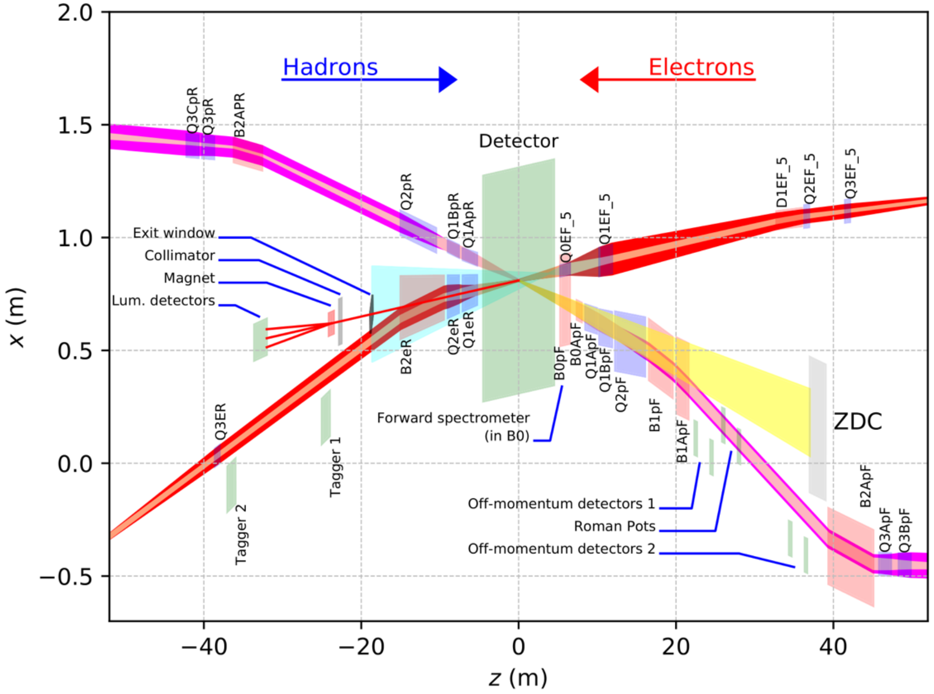

Compute-Detector Integration

Extend integrated interaction and detector region into detector readout (electronics), data acquisition, data processing and reconstruction, and physics analysis.

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 Due to Moderate Signal Rate:

- **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.

	EIC	RHIC	LHC → HL-LHC
Collision species	$\vec{e}+\vec{p}, \vec{e}+A$	$\vec{p}+\vec{p}/A, A+A$	$p+p/A, A+A$
Top 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**
Top collision rate	500 kHz	10 MHz	1-6 GHz
$dN_{ch}/d\eta$	0.1-Few	~3	~6
Charged particle rate	**4M N_{ch}/s**	**60M N_{ch}/s**	**30G+ N_{ch}/s**

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**.
- **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.

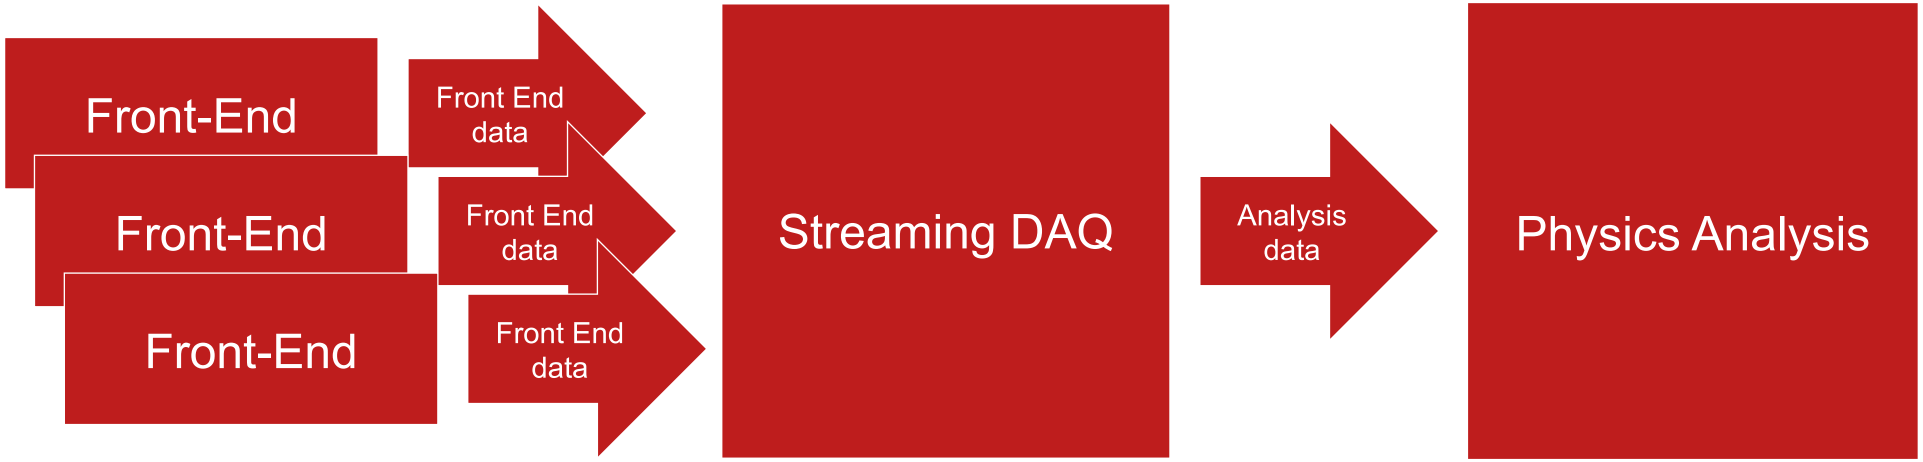

Front-End

Front-End

Front-End

Front End data

Front End data

Front End data

Streaming DAQ

Analysis data

Physics Analysis

Data Flow: **100 Tbps** → **10 Tbps** → **0.1 Tbps**

Report on “ePIC Streaming Computing Model”

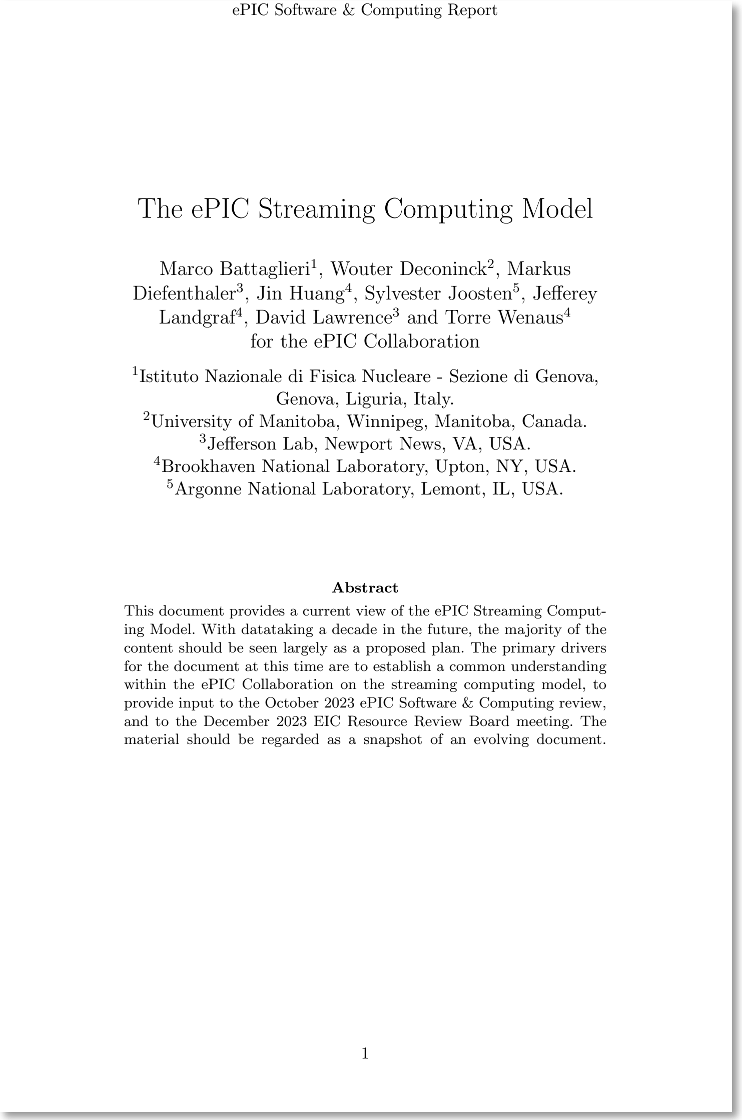

ePIC Software & Computing Report

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

For the review material, we have submitted a report on the "ePIC Streaming Computing Model”:

https://indico.bnl.gov/event/20481/page/566-report-by-epic

The structure of this presentation aligns with that of the report.

We provide references to various sections in the report for more detailed information on the topics.

Sec. 4	Reference to Sec. 4 in the report.

Computing Use Cases

Interface between Streaming DAQ and Streaming Computing

- **Ongoing Discussion**: Where does the interface lie between “online” and “offline”?
- **Working Definition**: Point where data moves to archival storage.
- **Challenges**: Technical and sociological differences at this point:
  - **Pre archival storage**: Risk of permanently losing data in case of error or reduced live time.
  - **Post archival storage**: Less stringent requirements and latencies. More open environment.

Scope of Use Case Discussion:

- Discussion of use cases centers on “offline” computing.

Use Cases:

- Stored Data Streaming and Monitoring
- Alignment and Calibration
- Prompt Reconstruction
- First Full Reconstruction
- Reprocessing
- Simulation
- Analysis
- Modeling and Digital Twin

Computing Use Cases: Stored Data Stream, Monitoring, and Calibrations

Use Case: Stored Data Streaming and Monitoring

Sec. 4.2

- **Primary Duty**: Archive incoming raw data from DAQ.
- **Data Replication**: Geographically separate replicas at BNL and JLab.
- **Real-Time Workflows**: Calibration and prompt processing.
- **Monitoring**: Automated and user-interface based validation, alarming.

Use Case: Alignment and Calibration

Sec. 4.3

- **Requirement**: Prompt alignment and calibration for rapid turnaround of reconstruction.
- **Workflows**: As automated and autonomous as possible during operation.
- **Data Products**: Sent to a globally accessible conditions database.
- **Initial Alignment and Calibration**: Only Echelon 1.
- **Refinements**: Can proceed in multiple Echelons.

Computing Use Cases: Reconstruction

Unique Requirement of Streaming Based Processing:
- Reconstruction of physics events from time frames as produced from the DAQ.

→ **Details on slide 28.**

Use Case: Prompt Reconstruction

Sec. 4.4

- **Real-Time**: Events reconstructed in near real-time.
- **Availability**: Crucial for quick raw data to production turnaround:
    - Required for monitoring, diagnostics, and quick-turnaround calibrations.
- **Limitations**: Confined to Echelon 1 due to strict latency requirements.

Use Case: First Full Reconstruction

Sec. 4.5

- **Resource Constraints**: Echelon 1 may not handle complete processing.
- **Shared Work**: Echelon 2 involved for remaining tasks.
- **Time-Scale**: Driven by calibrations. Maximum acceptable completion is 2-3 weeks.

Use Case: Reprocessing

Sec. 4.6

- **Forms**: Full reprocessing, re-reconstruction, regeneration of analysis object data.
- **Batch-Style**: Suited for batch processing.
- **Resources**: Utilizes Echelons 1-2 and other opportunistic resources.

Computing Use Cases: Simulation, Analysis, and Modeling

Use Case: Simulation

Sec. 4.7

- **Types**: Monte Carlo physics and detector simulation.
- **Volume**: At least one order of magnitude more simulated than real events.
- **Data Structure**: Similar to that of real data.
- **Workflow Strategy**: Mimic streaming data workflows to gain experience.

Use Case: Analysis

Sec. 4.8

- **Analysis Use Cases**: Broad science program with numerous observables.
- **Techniques**: Unfolding or joint theoretical-experimental analyses on the event level.
- **Resource Needs**: Varying Echelon requirements based on study complexity.

Use Case: Modeling and Digital Twin

Sec. 4.9

- **Modeling**: E.g., streaming data used for AI background modeling.
- **Digital Twin**: Complementary to detector simulations for real-time experimental control or optimization of experimental conditions.
- **Resources**: Echelon 1 and 2 primarily used for modeling workflows.

Use Case to Echelon Mapping

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
Stored Data Streaming and Monitoring	✓	✓		
Alignment and Calibration		✓	✓	
Prompt Reconstruction		✓		
First Full Reconstruction		✓	✓	
Reprocessing		✓	✓	
Simulation		✓ *	✓	
Analysis		✓ *	✓	✓
Modeling and Digital Twin		✓	✓	

*** Opportunistically**

Echelon 0

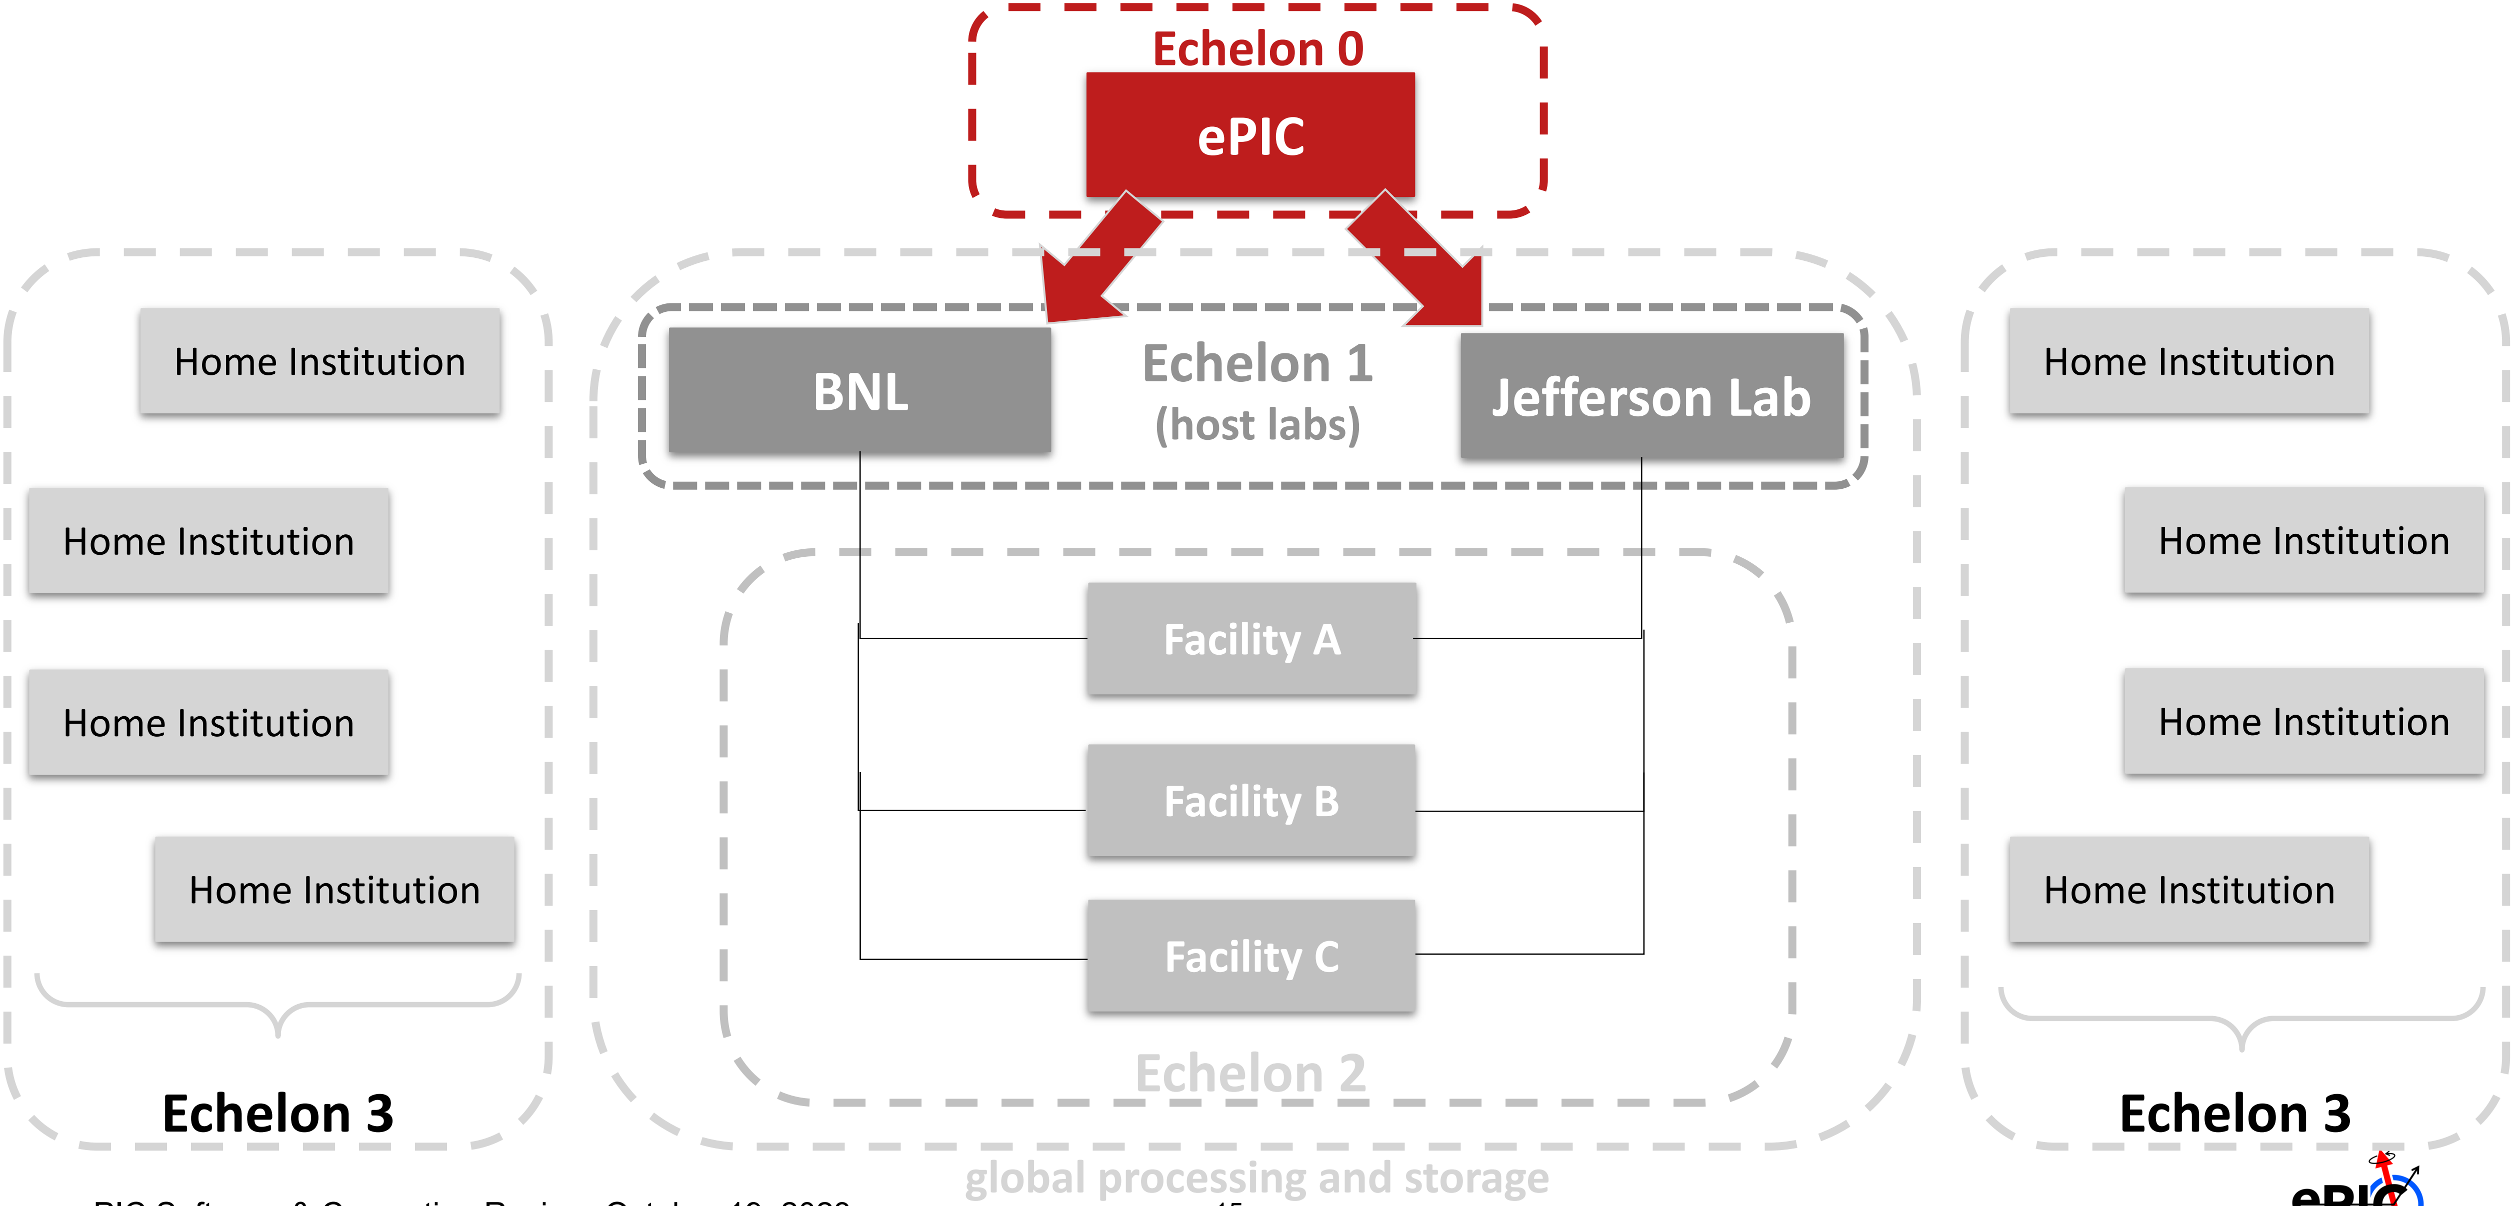
Echelon 0
ePIC
BNL
Echelon 1
(host labs)
Jefferson Lab
Home Institution
Home Institution
Home Institution
Home Institution
Home Institution
Home Institution
Home Institution
Home Institution
Facility A
Facility B
Facility C
Echelon 2
Echelon 3
Echelon 3
global processing and storage

Echelon 0: Expected Worst-Case Data Rate Contributions

Sec. 3.10

Fig. 8

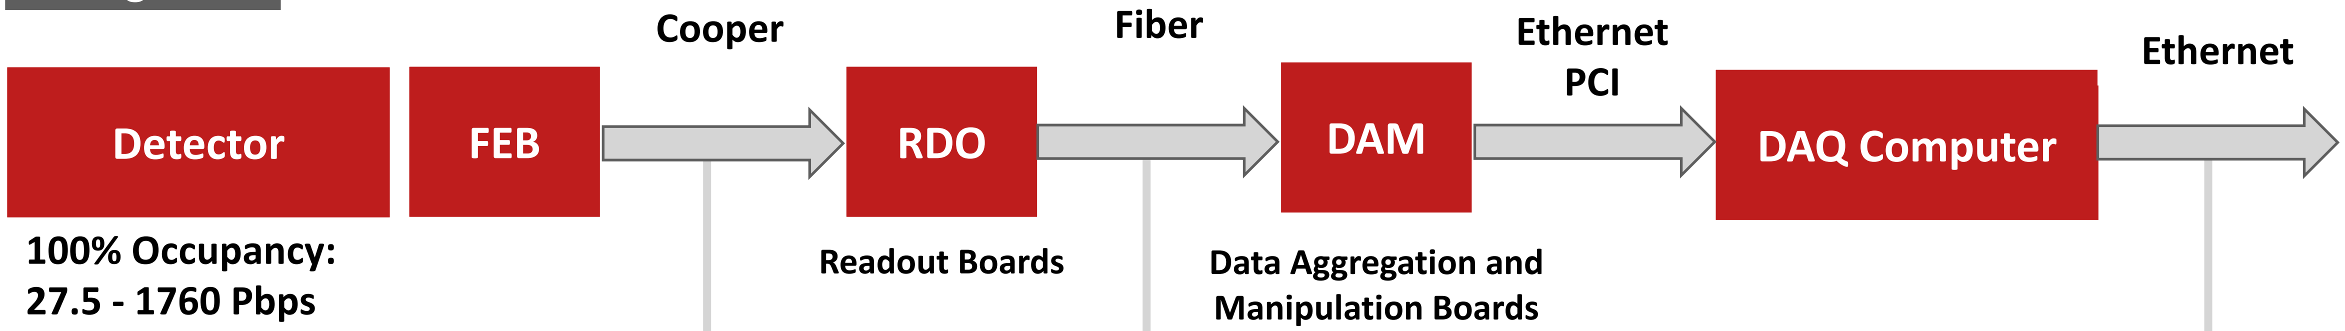

100% Occupancy:
27.5 - 1760 Pbps

Readout Boards

Data Aggregation and
Manipulation Boards

Aggregate	2.0 Tbps
Signal + Background	400 Gbps
Noise	1.6 Tbps

Aggregate	2.0 Tbps
Per RDO (Avg.)	0.7 Gbps

Aggregate	96 Gbps
Signal	38 Gbps
Synchrotron Radiation	0.1 Gbps
Electron Beam	22 Gbps
Hadron Beam	4 Gbps
Noise	32 Gbps

Echelon 0: Networking and Storage Requirements

- **Max Luminosity and Bandwidth**:
  - Expected maximum luminosity ≈ 10^{34} $cm^{-2}s^{-1}$.
  - Corresponding bandwidth ≈ 100 Gbps.
- **Design Considerations**: 100 Gbps between Echelon 0 and Echelon 1 at the host labs.
- **Outgoing Bandwidth**: Required: 200 Gbps

Table 2

Resource	Type	Amount
Outgoing bandwidth	Raw data	200 Gbps
	Monitoring, slow controls, misc. meta data	≤ 1 Gbps
	TOTAL	**≤ 201 Gbps**
Incoming bandwidth	Monitoring, calibrations	≤ 1 Gbps
Storage	Disk (outgoing data buffer w/ 24hr)	1PB

Echelon 1

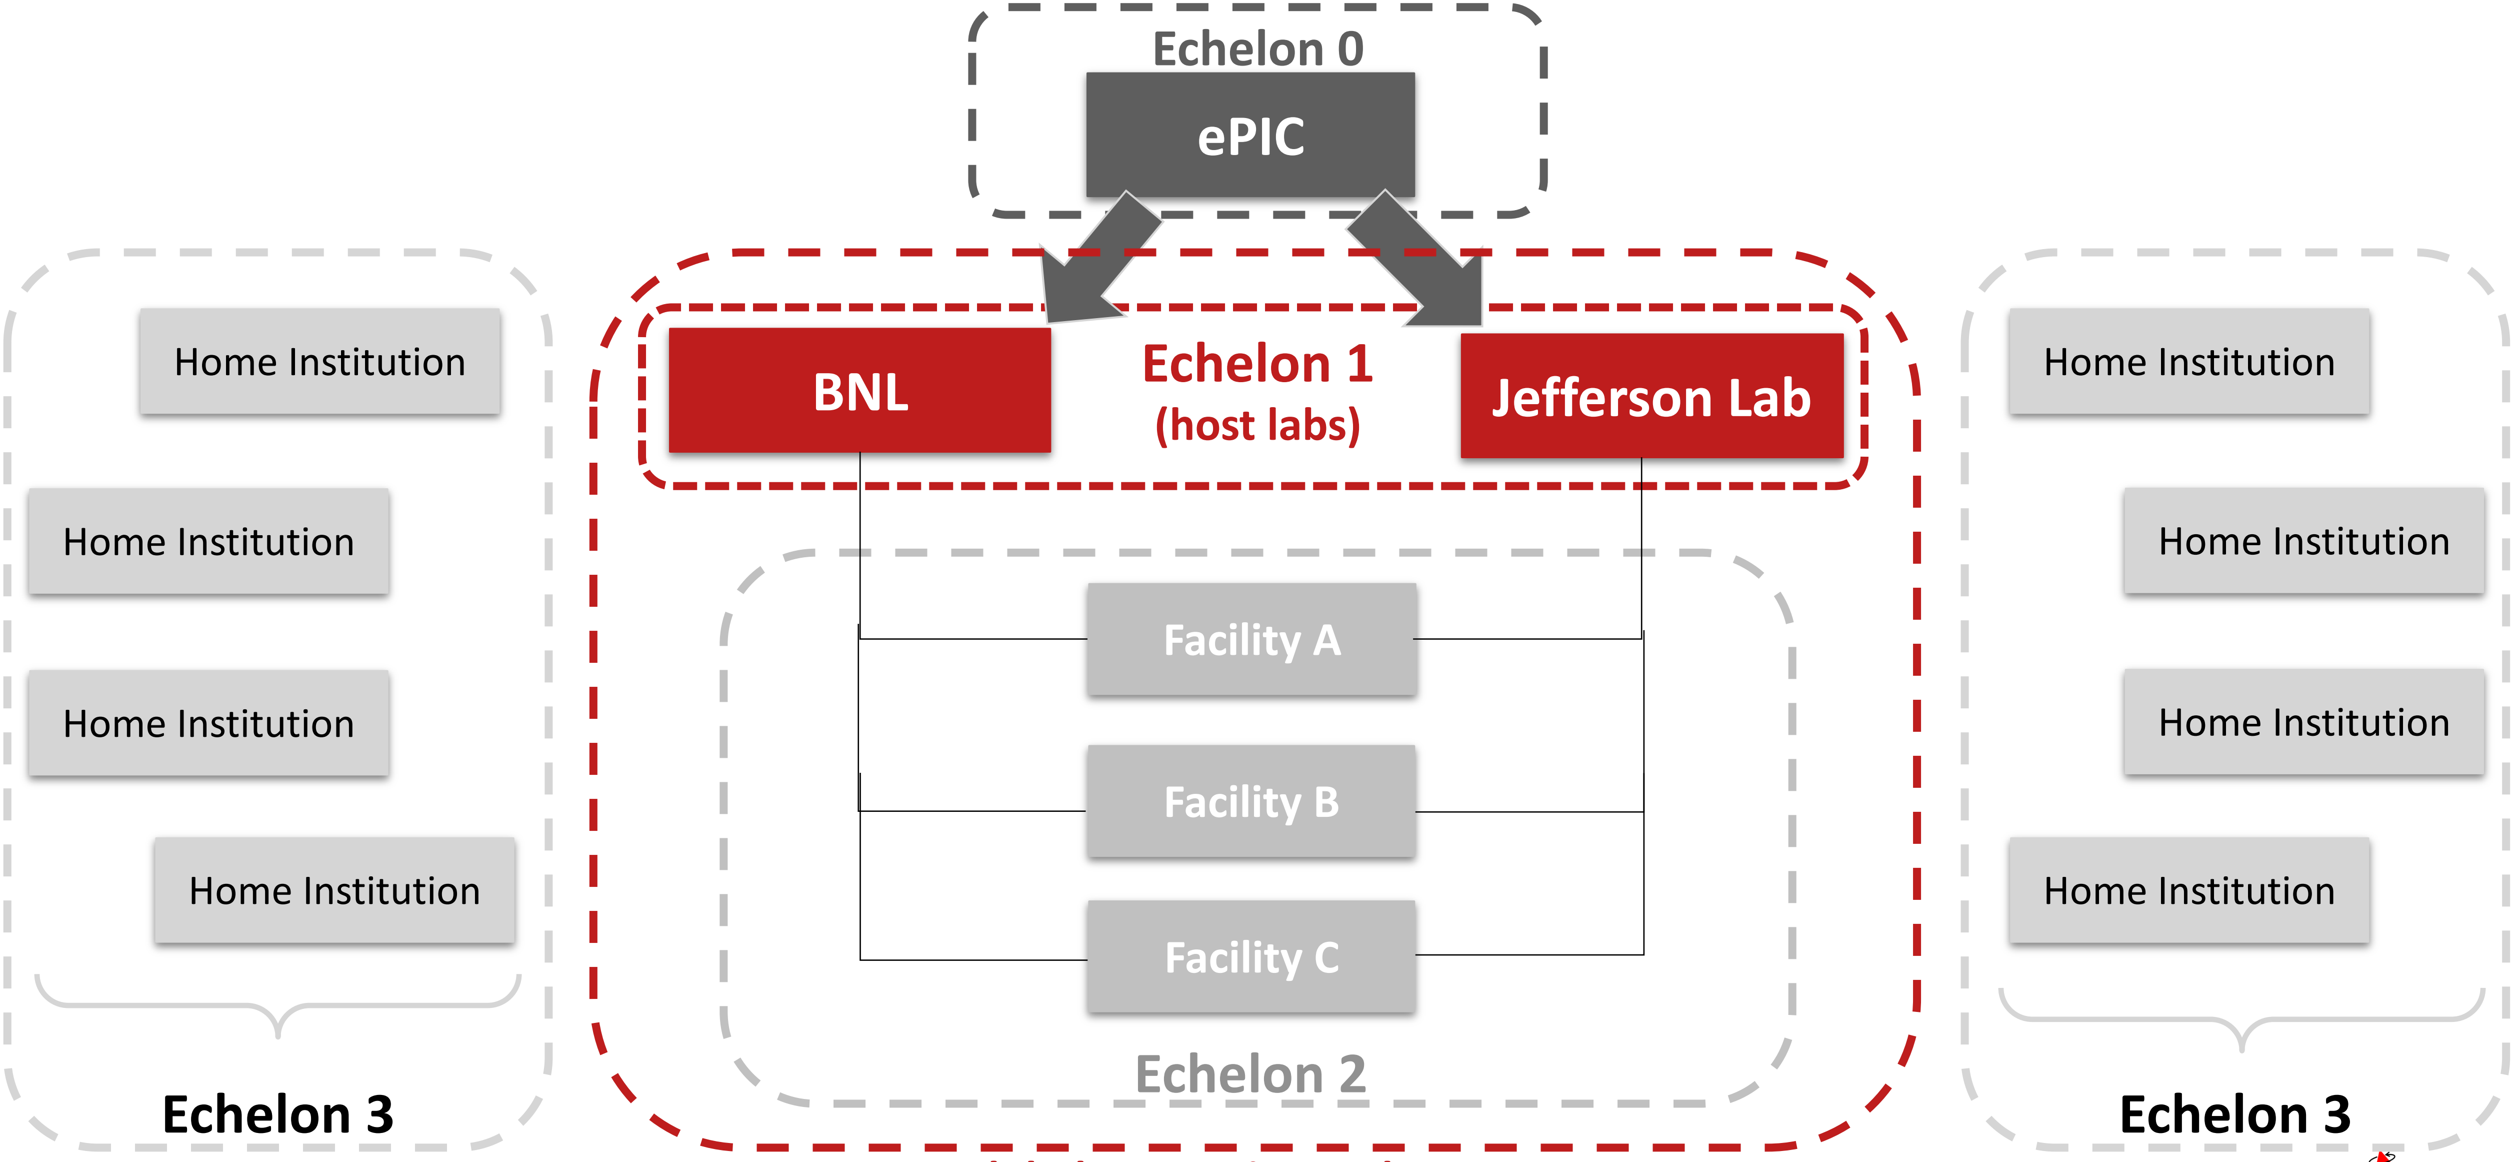

Echelon 0

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 1: Networking and Storage Requirements

Sec. 5.3.1

Sec. 5.3.2

Table 3

Resource	Type	Amount
Outgoing bandwidth	Raw data – *immediate* (1/6 of total)	17 Gbps
	Raw data – *replay* (contingency)	50 Gbps
	Monitoring, slow controls, misc. meta data	1 Gbps
	TOTAL	**≤ 68 Gbps**
Incoming bandwidth	Monitoring, calibrations, slow controls (from Echelons 1–2)	1 Gbps
Storage	Disk (temporary)	1PB
Raw+recon. only, no sim.	Disk (permanent)	20PB / yr
	Tape	220PB/yr

Values shown are for a single Echelon 1. There will be two.

Echelon 1: Comment on Compute Resources

Sec. 5.3.4

- **Status**: Current estimates on computing resources not reliable:
  - For accurate estimates, a prototype for holistic reconstruction of physics events from 1ms time slices is needed.
  - Need for reliable estimates on the fraction of background events and their impact on reconstruction.
  - Understanding speed at which these events can be discarded is crucial.

  - Importance of defining alignment and calibration methods for each subsystem.
  - Detailed discussions on fast alignment and calibration techniques are necessary.

- **Goal**: Reliable compute resource estimates before TDR.

The same comments pertain to the computing resources for Echelon 2.

Compute Resource Estimates

Sec. 5.3.4

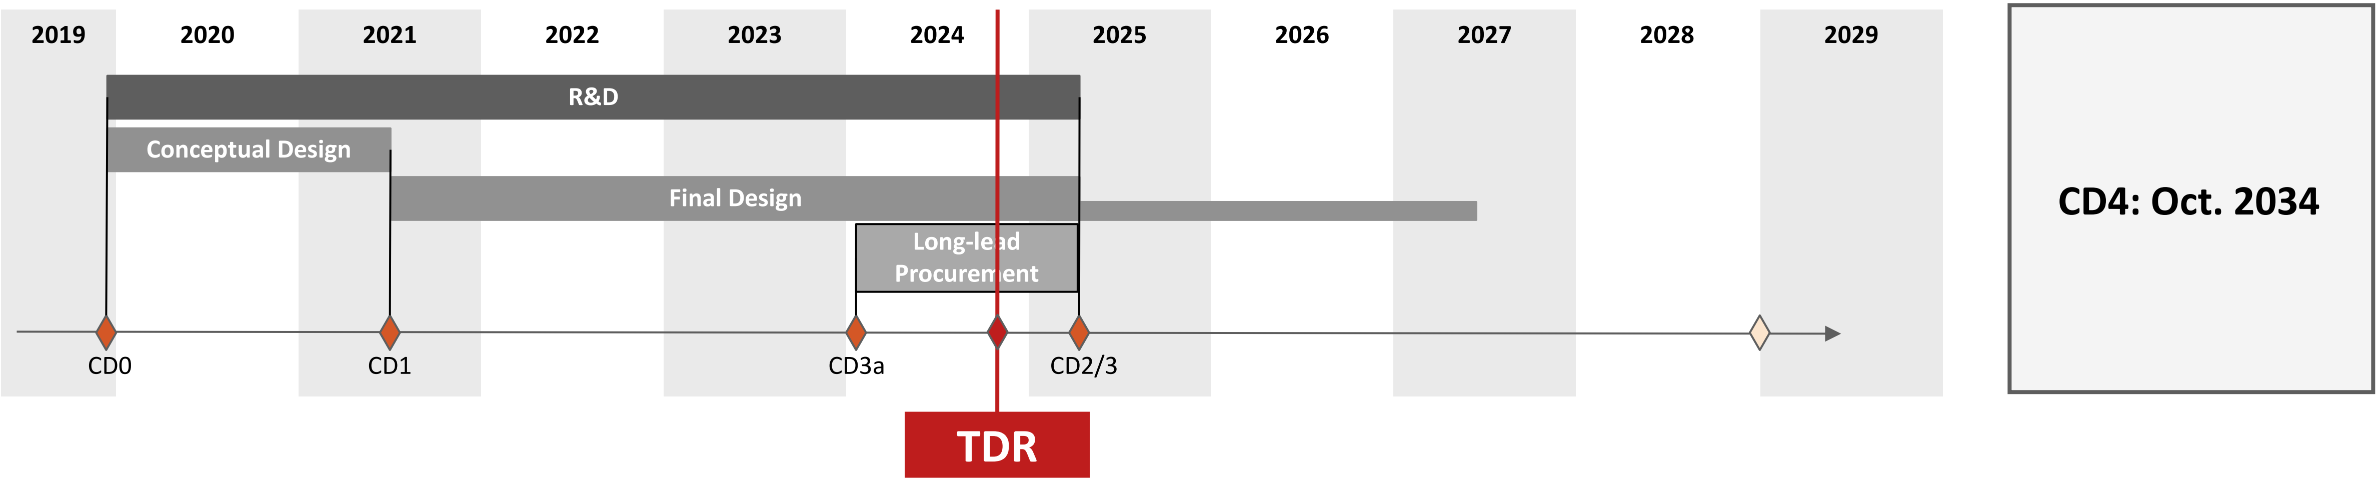

- **Goal**: Reliable compute resource estimates before TDR.
- **Planning and milestones**, as outlined in Sec. 9.2 , cover various aspects:
  - **By 12/23**: Define data model and format of time slices.
  - **By 03/24**: Detailed simulation of Streaming DAQ.
  - **By 06/24**: Prototype holistic event reconstruction from time slices.
  - **By 09/24**: Detailed compute resource estimates for all use cases from Slides 11–13.

Currently, we rely mainly on opportunistic compute resources for our simulation campaigns with an concurrent core count of 5-10k. Expect to grow substantially once we fold in Streaming DAQ simulations.

Echelon 2

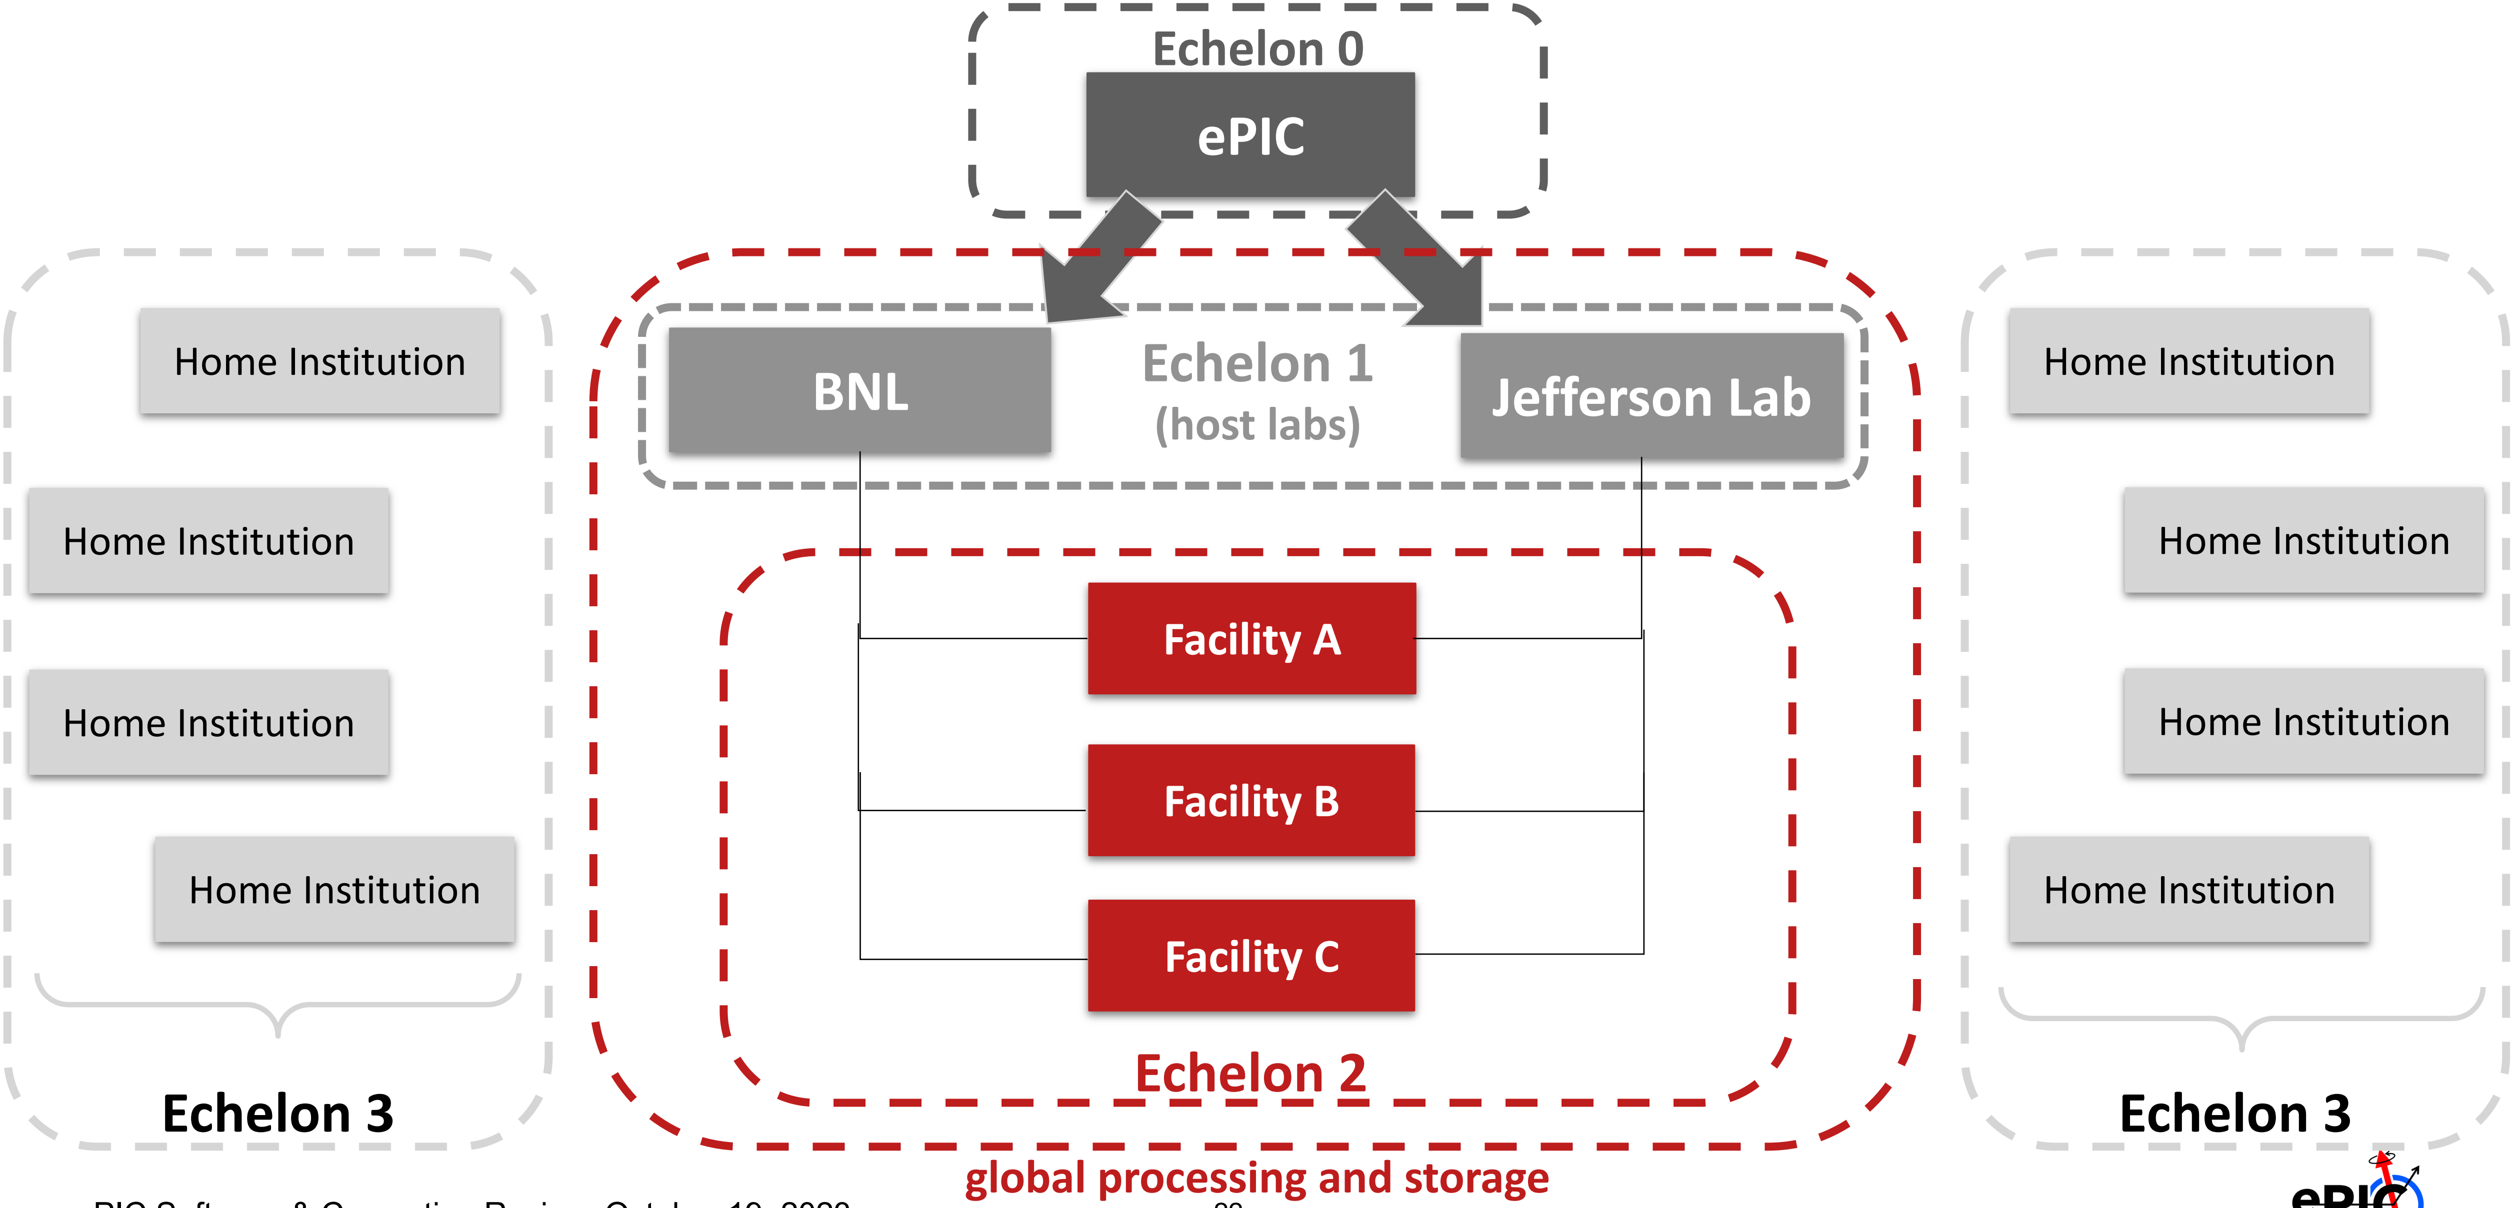
Echelon 0
ePIC
Echelon 1
(host labs)
BNL
Jefferson Lab
Facility A
Facility B
Facility C
Echelon 2
global processing and storage
Home Institution
Home Institution
Home Institution
Home Institution
Echelon 3
Home Institution
Home Institution
Home Institution
Home Institution
Echelon 3

Echelon 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**.

International computing contributions are essential.

→ **See also Andrea's talk.**

Lessons Learned from LHC

- Echelon 2 resources must have **MOUs specifying service requirements**.
- These MOUs **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.

Echelon 3

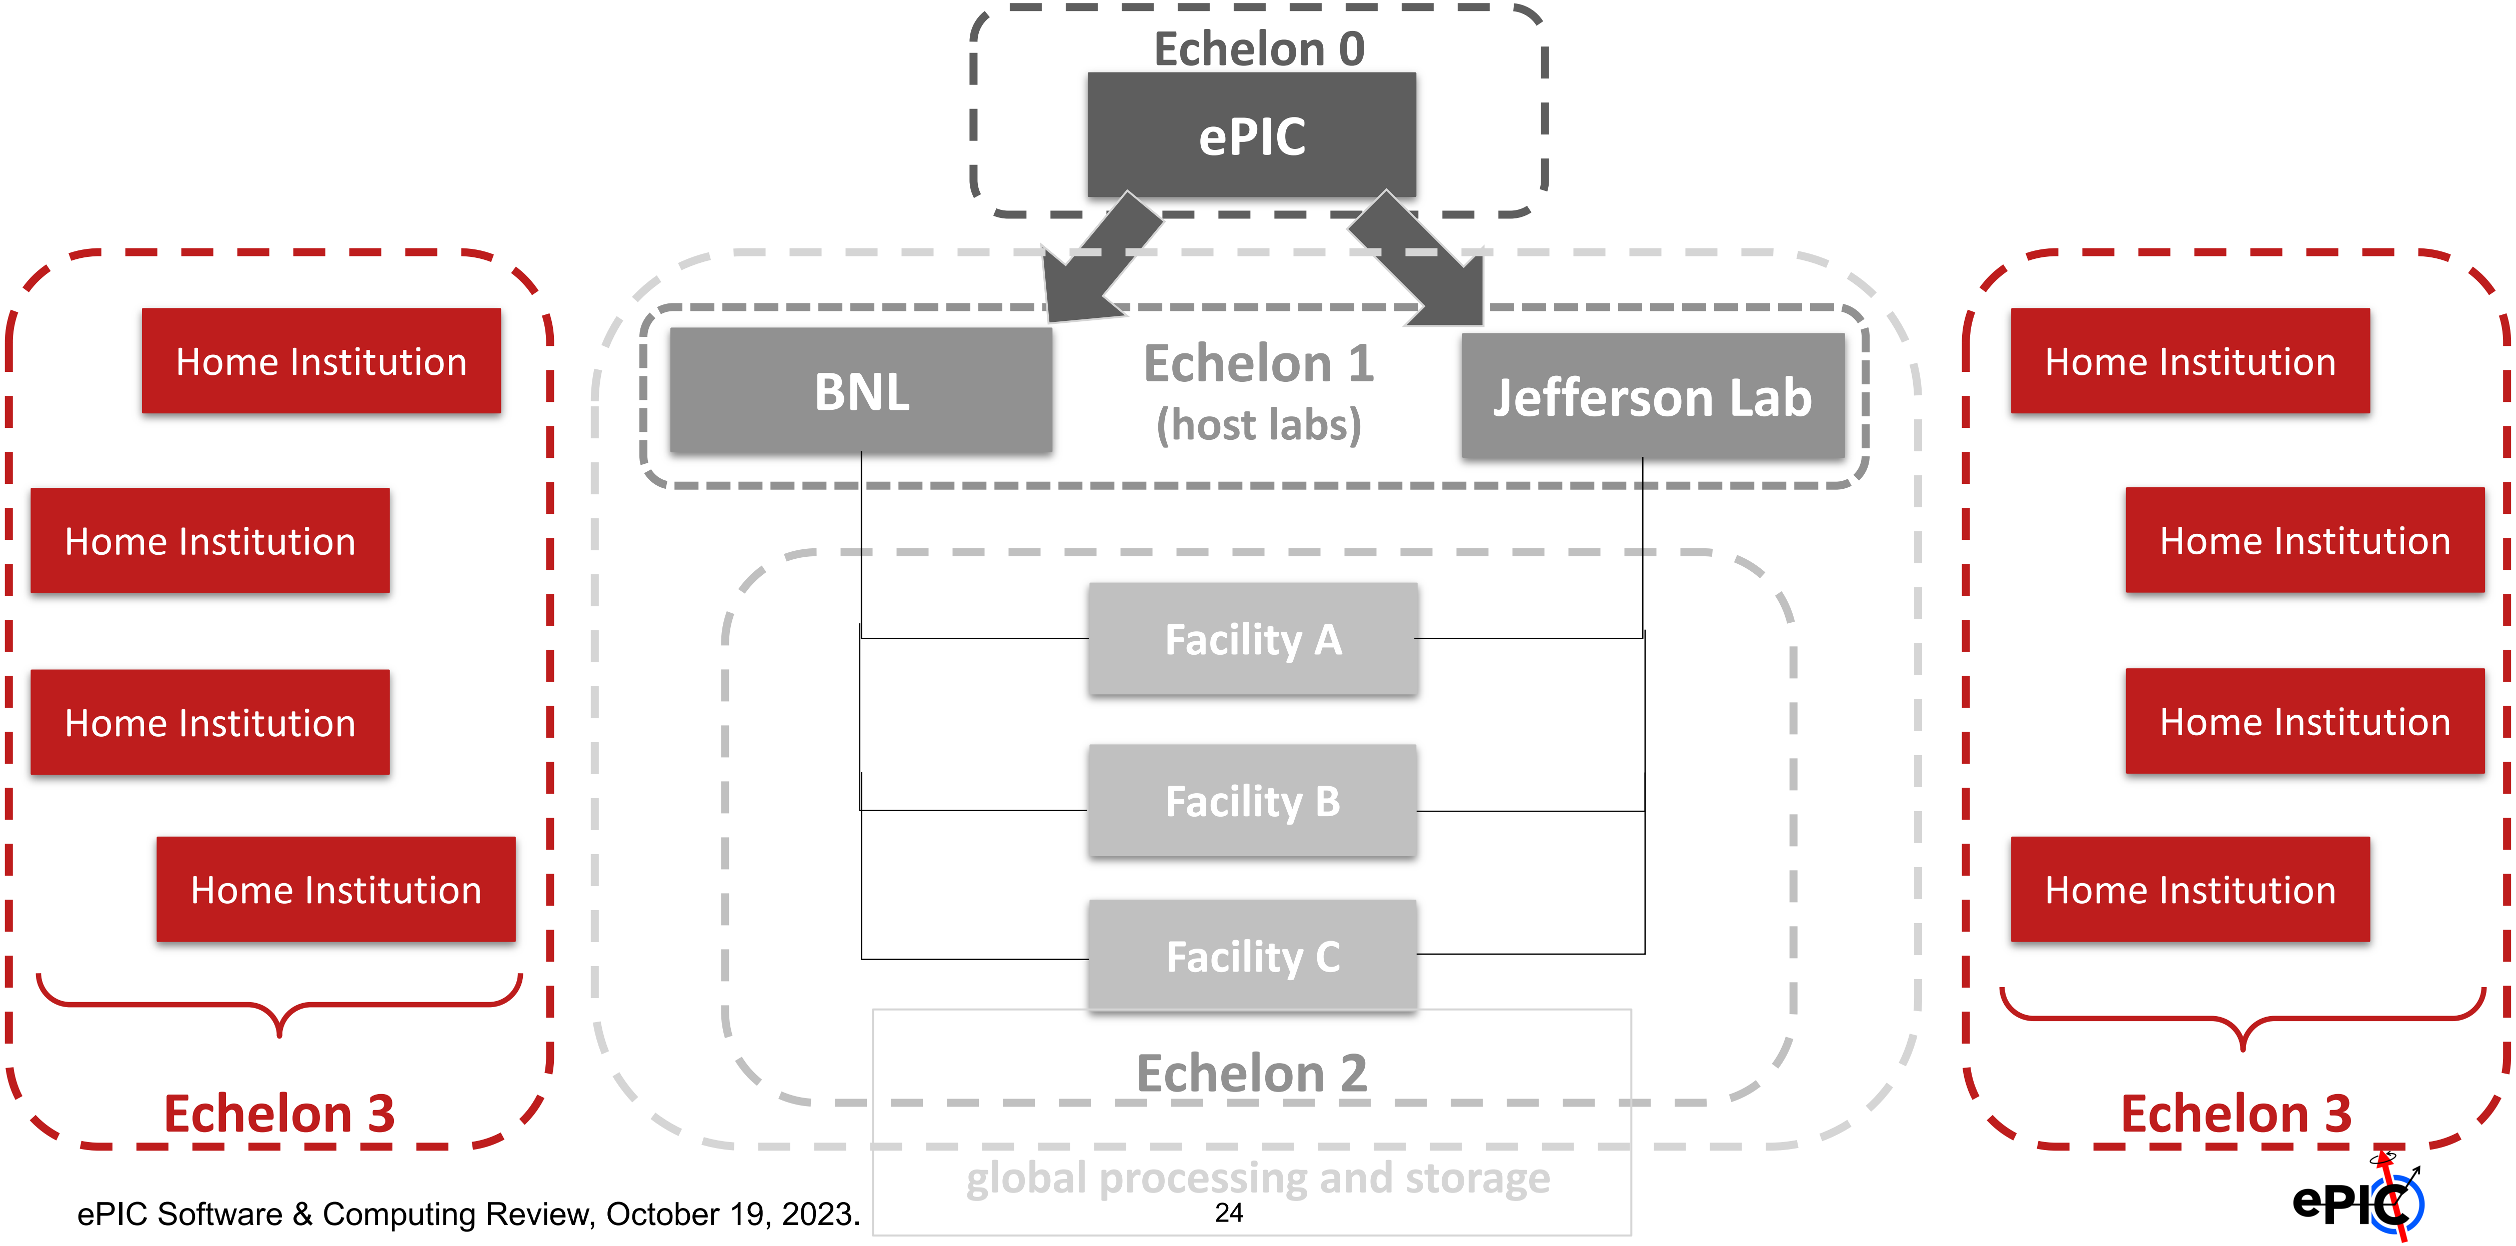

Echelon 0
ePIC
Echelon 1
(host labs)
BNL
Jefferson Lab
Facility A
Facility B
Facility C
Echelon 2
global processing and storage
Home Institution
Home Institution
Home Institution
Home Institution
Echelon 3
Home Institution
Home Institution
Home Institution
Home Institution
Echelon 3

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.

Opportunistic Resources

- **Use of Open Science Grid (OSG):**
  - Currently the most productive resource for ePIC. Expected to continue playing a role (simulation productions).
  - Stable concurrent core count of 5-10k.
    - Expect to grow substantially once we fold in Streaming DAQ simulations.

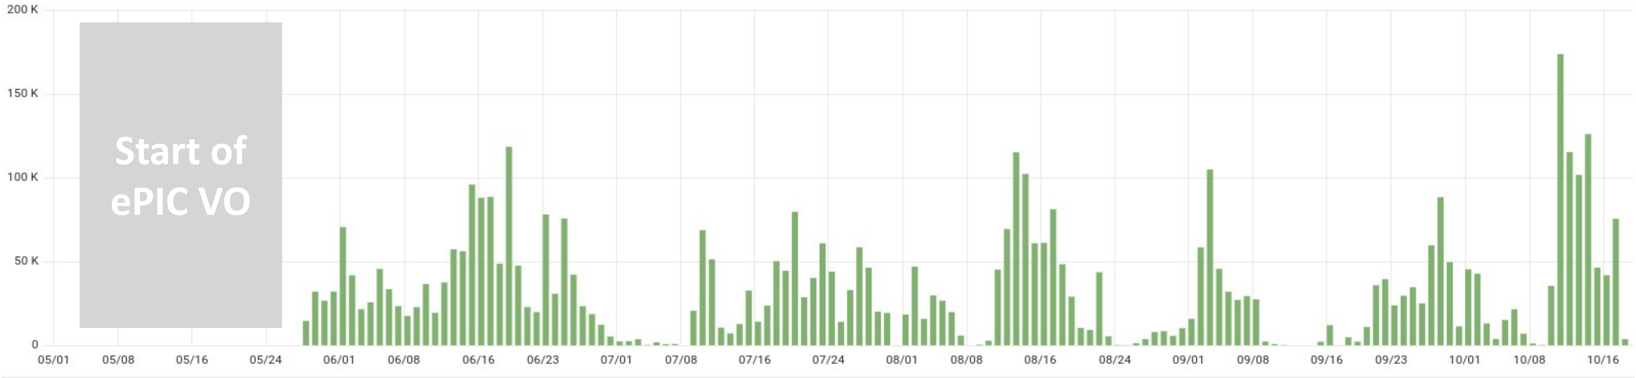

- **Opportunistic Resources:**
  - Utilization of opportunistic resources like OSG is foreseen as ePIC science will likely be compute-limited.
  - WLCG evolving in also supporting non-LHC experiments.
- **Commercial Clouds:**
  - Actively used by other science communities.
  - Cost-effectiveness is promising, especially for fast-turnaround use cases.
  - ePIC is monitoring these developments for future consideration.

Special Resources

Sec. 5.6

EIC SOFTWARE:
Statement of Principles

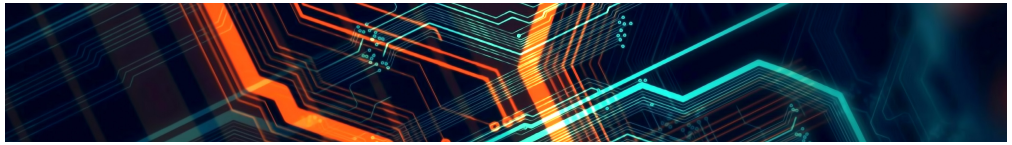

③ **We will leverage heterogeneous computing:**

- 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.

Special Resources:

- Includes non-x86 architectures like ARM, and accelerators like GPUs and TPUs.
- ePIC infrastructure aims to be flexible and extensible to support emerging architectures.
- ARM already supported. Will play important role due to its cost-effectiveness and ease of porting.
- FPGAs will be used for low-level data processing in Echelon 0.

Artificial Intelligence:

- Tracking the technologies for AI R&D and applications.
- Likely to exploit specialized accelerators like TPUs.

Large Supercomputers (LCFs): Developed by the DOE and NSF, often include special resources.

- Effectiveness for ePIC use will be evaluated on a case-by-case basis.

Authorization and Access

- **ePIC Collaboration Priorities**: Ensuring access for every collaborator to collaboration resources:
    - Data,
    - Websites,
    - Collaborative tools,
    - Information systems,
    - Document repositories,
    - And so on.
- Accessibility as a **leading or determining factor**:
    - Influences choice of tools and services.
    - Influences hosting decisions.
- **Examples:**
    - GitHub chosen for code repository: https://github.com/eic
    - Cloud-based Mattermost instance: https://chat.epic-eic.org

ePIC, the ECSJI, and the EIC RRB

The ePIC collaboration supports the formation of **EIC Computing and Software Joint Institute (ECSJI)** and its associated bodies, including the **EIC International Computing Organization (EICO)**, to oversee and coordinate the **complex computing fabric of ePIC and the EIC**:

- Crucial and innovative Echelon 1 partnership between the host labs,
- Global contributions represented at Echelon 2,
- Full support of the analysis community at Echelon 3.

- ECSJI plays a crucial role in the computing aspects of the EIC.
- The experiments design and develop their computing models and software, consistent with the computing fabric developed under the oversight of the ECSJI.
- ePIC computing operations are developed within the ePIC Collaboration, in close consultation and collaboration with ECSJI and other computing resource providers.
- Oversight and review mechanisms are in place for both ECSJI and ePIC.
- The ePIC Software & Computing Coordinator serves as the primary liaison between ePIC and ECSJI.
- The EIC Resource Review Board (RRB) is responsible for overseeing resources, including software and computing.
  - ePIC reports its computing and software status, as well as multi-year resource requirements, to the RRB for review and decision-making.

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.

Workflow Management and Data Management

Sec. 6.2 Sec. 6.3

- **Technology Adoption**: Take advantage of rapidly evolving streaming frameworks and tools while avoiding lock-in.
  - **Common features**, e.g., distributed parallel model or workflow descriptions, minimize risk of lock-in.
  - **Resource Management**: Systems like Apache Storm and Spark manage resources directly, while others like PanDA overlay on batch or dynamic processing.
  - **Real-time data flow optimization** must supported by both the facility and the streaming WMS.
  - **Resource flexibility** across various use cases and workflows, readily re-purposable.

- **Batch Processing Support**: Workflows like simulation and reprocessing benefit from conventional batch processing.
- **International Collaboration**: ePIC's global nature necessitates workflow tools to support global computing resource use, excluding prompt reconstruction.

Rucio for Data Management

- Broad acceptance in HEP and NP.
- Currently being integrated and tested in ePIC.
- **Authentication in Rucio**
  - Migration to SciToken-based authentication and authorization (AA) mechanisms.
  - Uniform authorization across scientific computing infrastructures.
- **Rucio and Data Movement Tool**
  - Diverse data movement tools: XRootD, S3 and other object storage, FTS, Globus, httpd.
  - Rucio encapsulates fragmentation and avoids system lock-in.

Milestones

Sec. 9.2.1

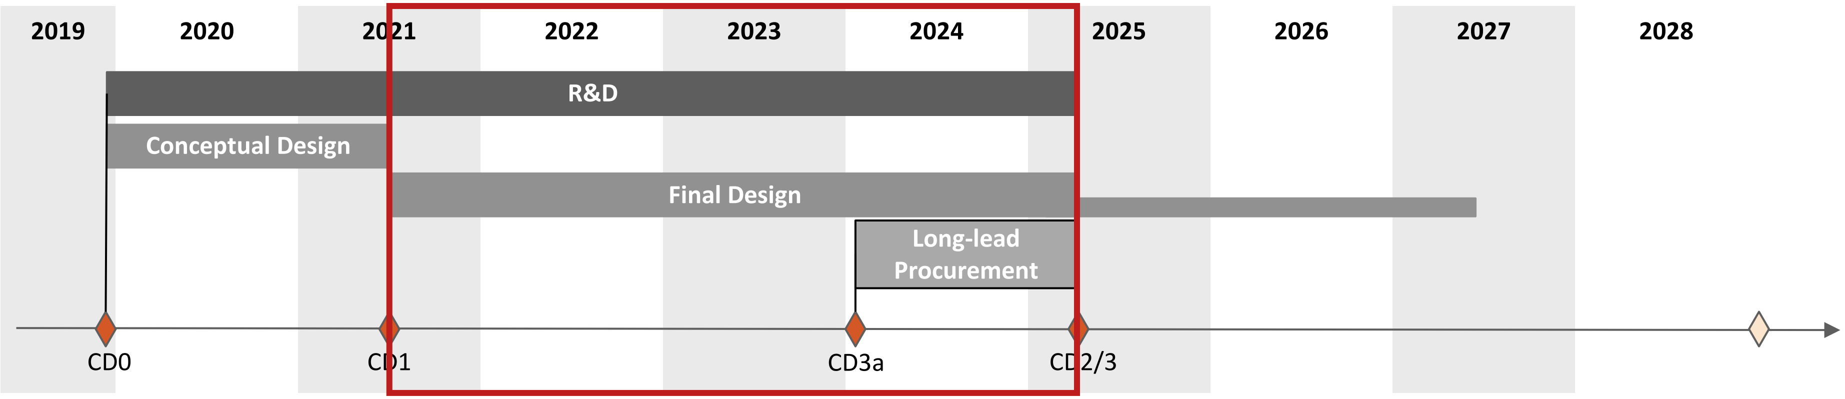

CD2/3 and TDR

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

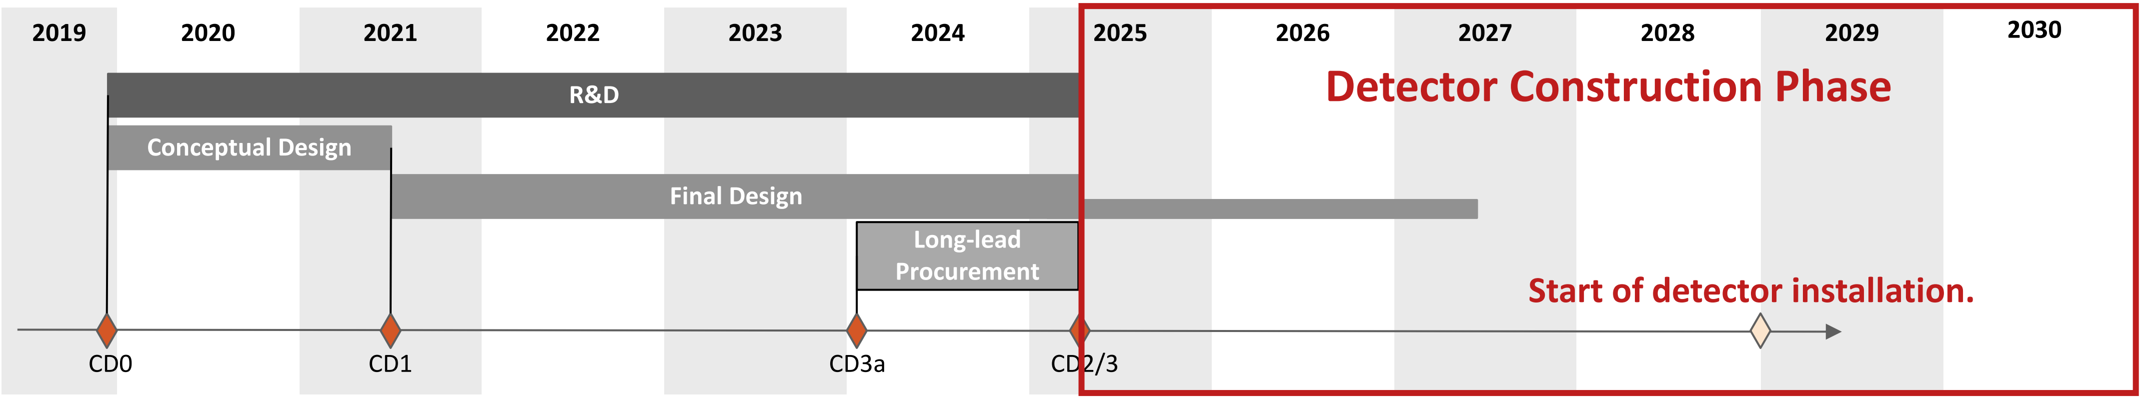

Detector Construction Phase

Start of detector installation.

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

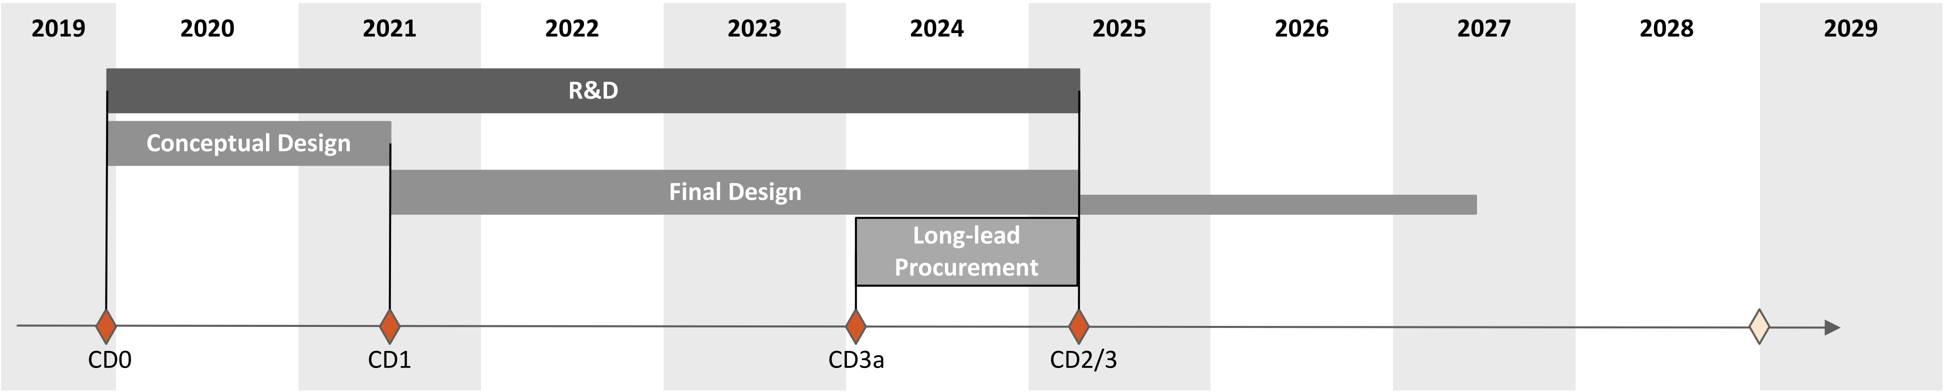

CD4: Oct. 2034

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.

Summary

- **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.

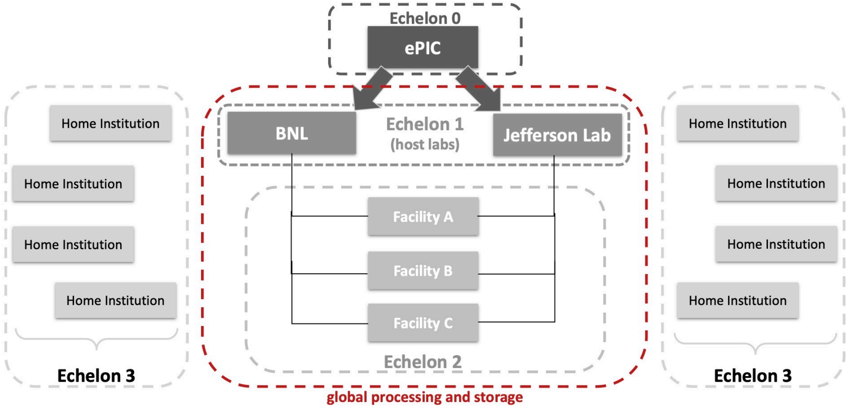

- ePIC supports the ECSJI formation to oversee and coordinate the complex computing fabric of the EIC.
- **High level milestones** ensures that the agile development process is continuously confronted with real world exercising of the software and the developing realization of the computing model:
  - Priority always given to meeting near-term needs.
  - 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 0: Scale of the Streaming DAQ

Sec. 3.7

Fig. 7

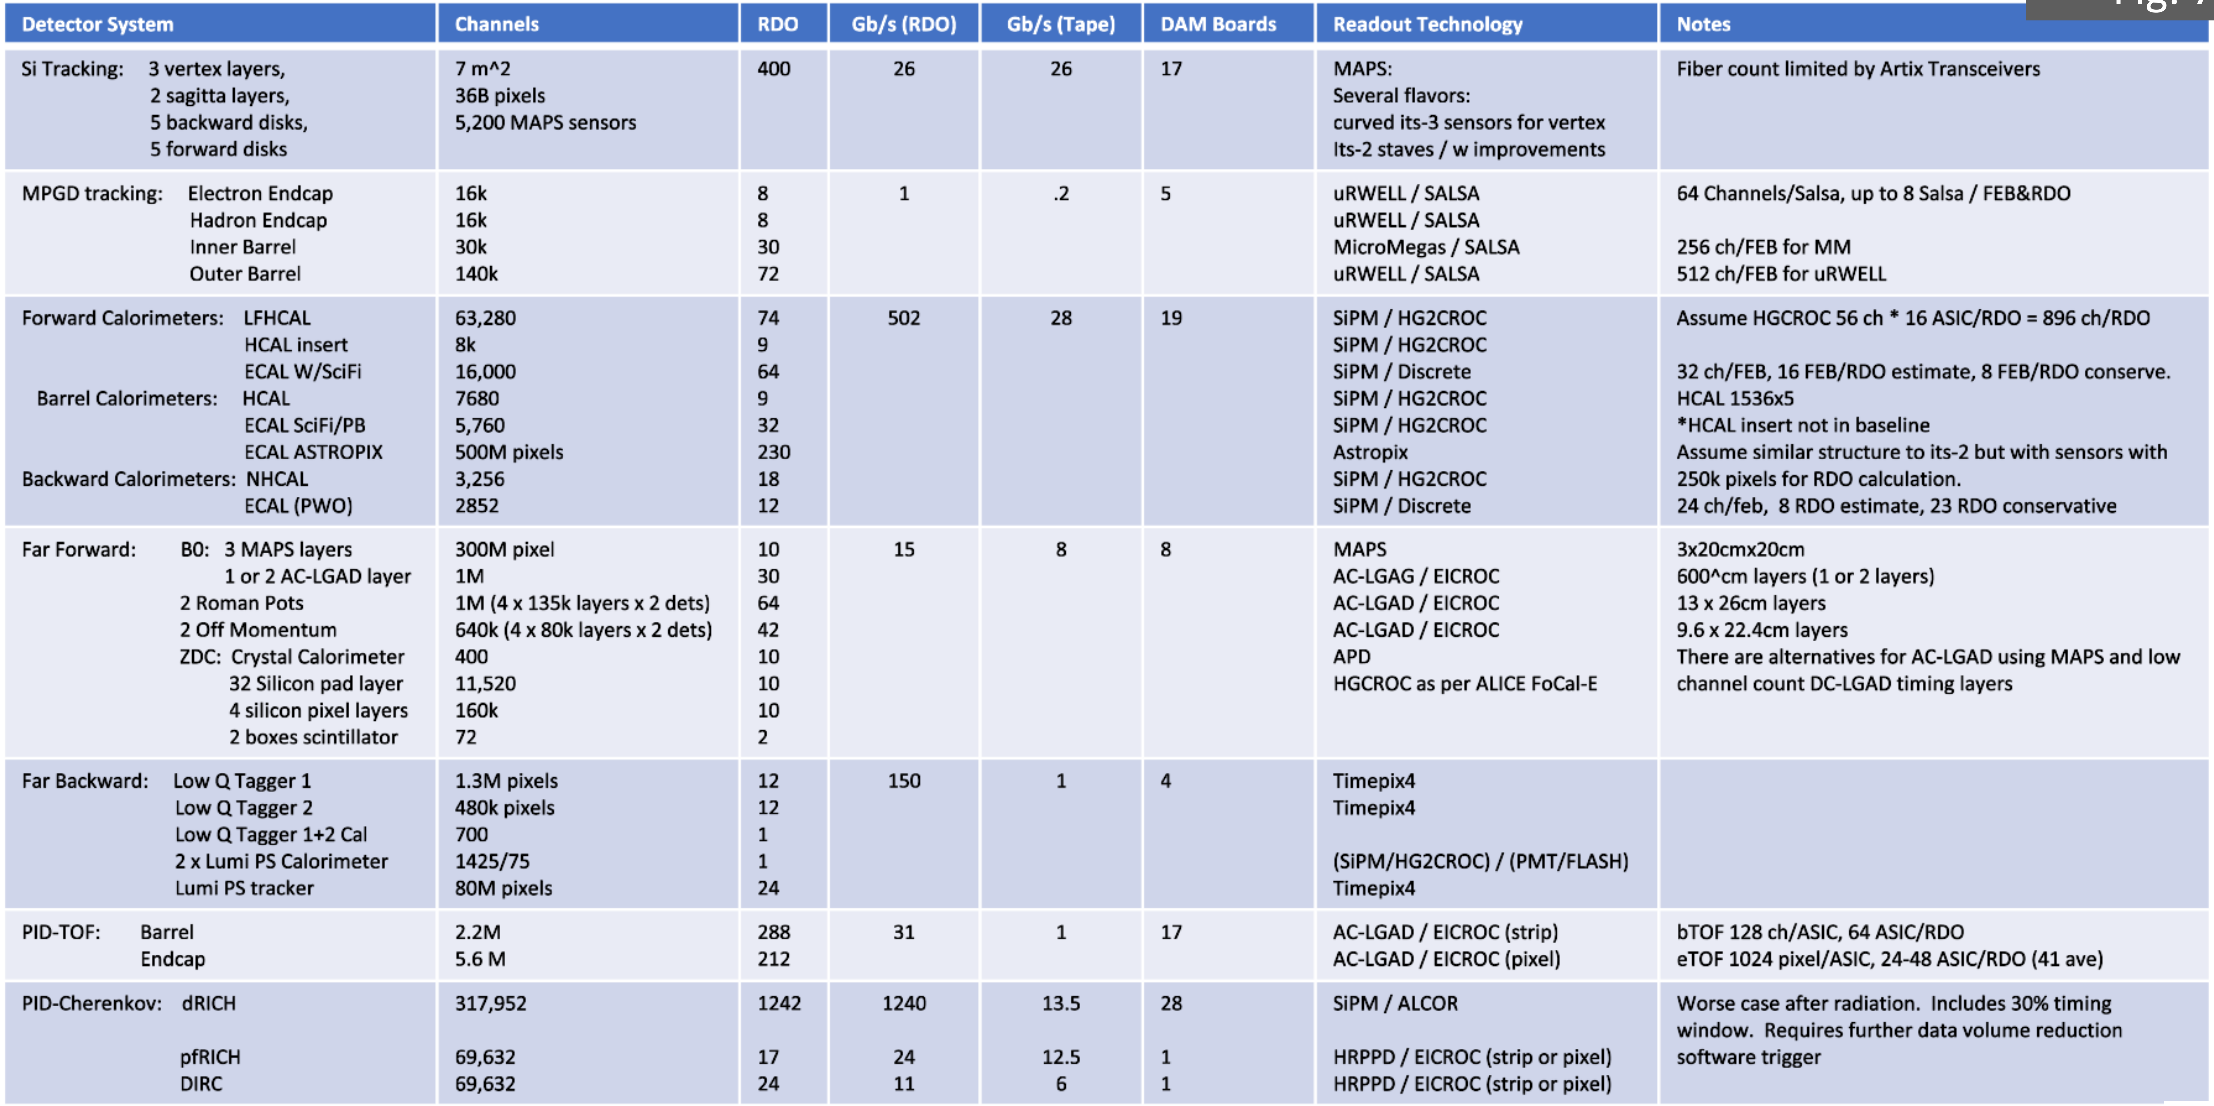

Detector System	Channels	RDO	Gb/s (RDO)	Gb/s (Tape)	DAM Boards	Readout Technology	Notes
Si Tracking: 3 vertex layers, 2 sagitta layers, 5 backward disks, 5 forward disks	7 m^2 36B pixels 5,200 MAPS sensors	400	26	26	17	MAPS: Several flavors: curved its-3 sensors for vertex Its-2 staves / w improvements	Fiber count limited by Artix Transceivers
MPGD tracking: Electron Endcap Hadron Endcap Inner Barrel Outer Barrel	16k 16k 30k 140k	8 8 30 72	1	.2	5	uRWELL / SALSA uRWELL / SALSA MicroMegas / SALSA uRWELL / SALSA	64 Channels/Salsa, up to 8 Salsa / FEB&RDO 256 ch/FEB for MM 512 ch/FEB for uRWELL
Forward Calorimeters: LFHCAL HCAL insert ECAL W/SciFi Barrel Calorimeters: HCAL ECAL SciFi/PB ECAL ASTROPIX Backward Calorimeters: NHCAL ECAL (PWO)	63,280 8k 16,000 7680 5,760 500M pixels 3,256 2852	74 9 64 9 32 230 18 12	502	28	19	SiPM / HG2CROC SiPM / HG2CROC SiPM / Discrete SiPM / HG2CROC SiPM / HG2CROC Astropix SiPM / HG2CROC SiPM / Discrete	Assume HGCROC 56 ch * 16 ASIC/RDO = 896 ch/RDO 32 ch/FEB, 16 FEB/RDO estimate, 8 FEB/RDO conserve. HCAL 1536x5 *HCAL insert not in baseline Assume similar structure to its-2 but with sensors with 250k pixels for RDO calculation. 24 ch/feb, 8 RDO estimate, 23 RDO conservative
Far Forward: B0: 3 MAPS layers 1 or 2 AC-LGAD layer 2 Roman Pots 2 Off Momentum ZDC: Crystal Calorimeter 32 Silicon pad layer 4 silicon pixel layers 2 boxes scintillator	300M pixel 1M 1M (4 x 135k layers x 2 dets) 640k (4 x 80k layers x 2 dets) 400 11,520 160k 72	10 30 64 42 10 10 10 2	15	8	8	MAPS AC-LGAG / EICROC AC-LGAD / EICROC AC-LGAD / EICROC APD HGCROC as per ALICE FoCal-E	3x20cmx20cm 600^cm layers (1 or 2 layers) 13 x 26cm layers 9.6 x 22.4cm layers There are alternatives for AC-LGAD using MAPS and low channel count DC-LGAD timing layers
Far Backward: Low Q Tagger 1 Low Q Tagger 2 Low Q Tagger 1+2 Cal 2 x Lumi PS Calorimeter Lumi PS tracker	1.3M pixels 480k pixels 700 1425/75 80M pixels	12 12 1 1 24	150	1	4	Timepix4 Timepix4 (SiPM/HG2CROC) / (PMT/FLASH) Timepix4	
PID-TOF: Barrel Endcap	2.2M 5.6 M	288 212	31	1	17	AC-LGAD / EICROC (strip) AC-LGAD / EICROC (pixel)	bTOF 128 ch/ASIC, 64 ASIC/RDO eTOF 1024 pixel/ASIC, 24-48 ASIC/RDO (41 ave)
PID-Cherenkov: dRICH pfRICH DIRC	317,952 69,632 69,632	1242 17 24	1240 24 11	13.5 12.5 6	28 1 1	SiPM / ALCOR HRPPD / EICROC (strip or pixel) HRPPD / EICROC (strip or pixel)	Worse case after radiation. Includes 30% timing window. Requires further data volume reduction software trigger