ePIC Software & Computing Report

D. Anderson, M. Battaglieri, T. Britton, C. Chatterjee, M. Diefenthaler, S. Gardner, T. Gunji, T. Jeske, D. Kalinkin, S. Kay, J. Landgraf, S. Li, A. Prozorov, S. Rahman, H. Szumila, and T. Wenaus on behalf of ePIC Software & Computing.

ePIC Software & Computing Organization

Coordinators and WG Conveners

Development

Dmitry Kalinkin





Simulation WG: Simon Gardner, Sakib Rahman. **Reconstruction WG**: Derek Anderson, Chandra Chatterjee, Shujie Li.

Infrastructure

Torre Wenaus



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

Operations

Holly Szumila-Vance





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

Roles

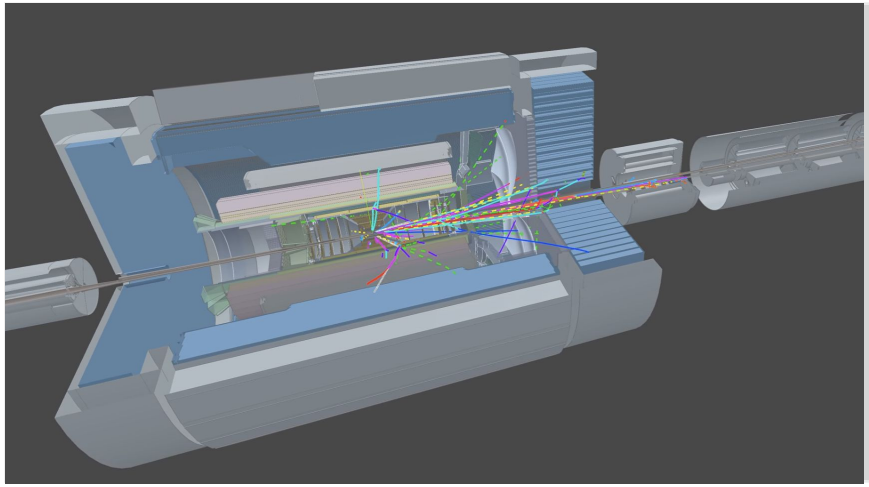

Enable Science

- Development, operations, and infrastructure for preTDR and future milestones.
- **Activity:** Monthly software releases and simulation campaigns tailored to collaboration needs.

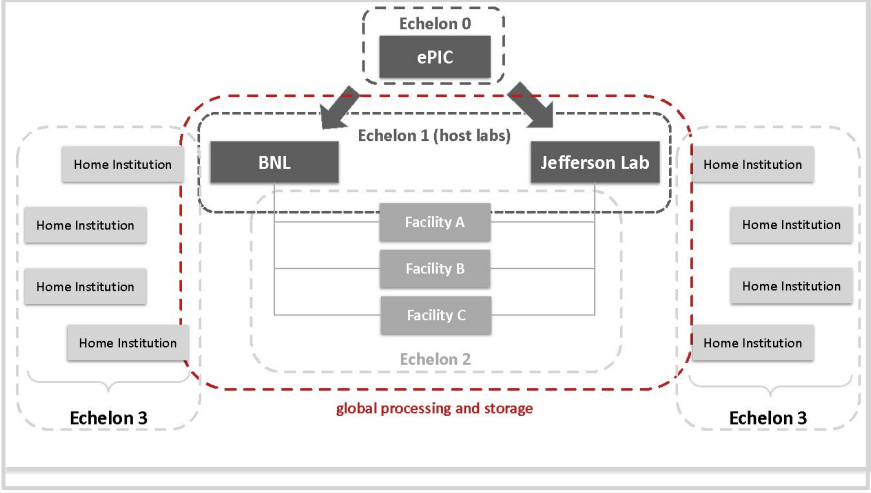

Compute–Detector Integration

- Maximize and accelerate science with the ePIC computing model
- **Activity:** Define and test DAQ–computing interfaces to mitigate integration risks. Develop and deliver a functional testbed to validate workflow management for autonomous detector calibration

ePIC Software Releases

What are the Software Releases?

- New tags for our simulation geometry and reconstruction softwares incorporating latest developments
- Progressive software stack updates
- Container pushed to DockerHub, ghcr.io, CVMFS
- Releases presented at the Wednesday meetings (26.01, 26.02, 26.03/26.04), and on ~compsw

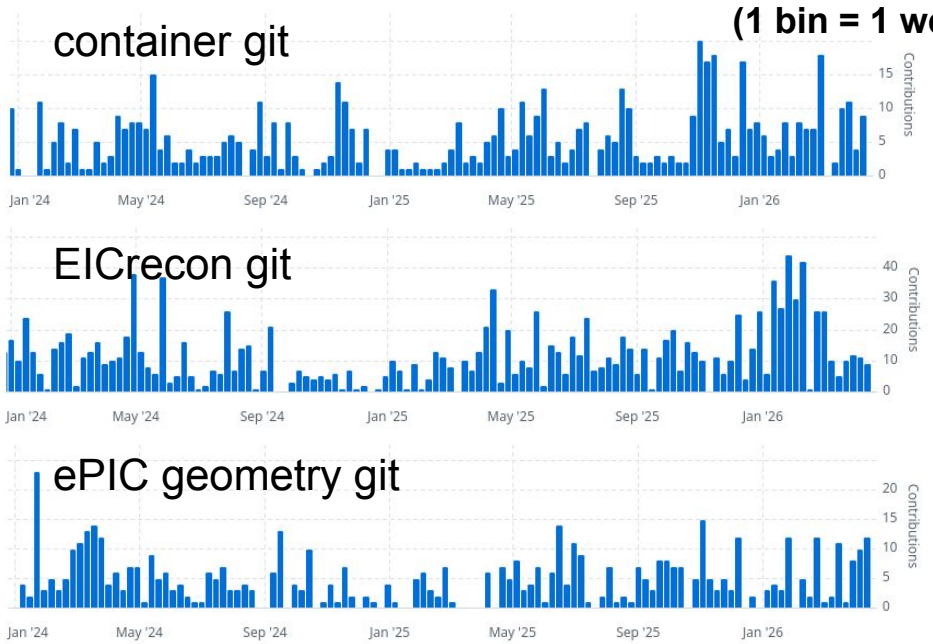

Recent highlights

- Geometry updates to detectors and services (see the next slide)
- Countless bugfixes/workarounds to issues reported by users and the production
- Evolution of tracking (performance optimizations, 10 um gold coating, updated material map, 4 hit track cut)
- Installed new codes: IRT2, eic-opticks, G4OCCT, g4adept
- Various fixes to MC truth associations (central tracking, FF, calorimetry)
- Container delivery from DockerHub to CVMFS using write node at BNL – reduced worst case delay from weeks to hours
- Infrastructure updates towards future developments: RNTuple in PODIO, Acts Gen3 geometry, MCP servers for AI

DD4hep updates

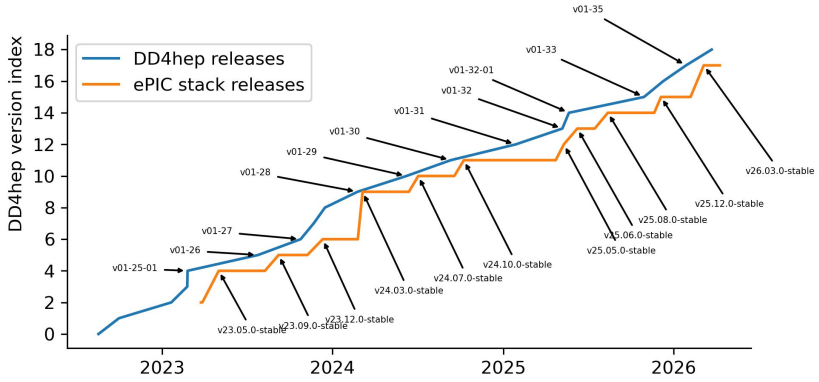

Acts updates

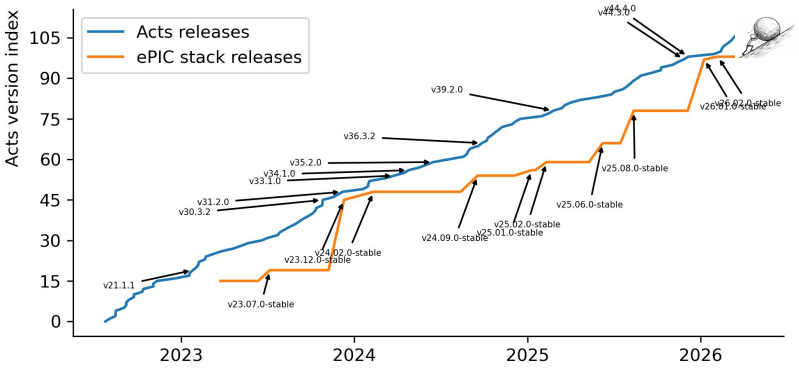

WG Highlights: Physics and Detector Simulation

Simulation Geometry Validation

- End of May, 2026 deadline for validation of ePIC simulation geometry against September, 2025 EIC project CAD envelopes.
- **Simulation geometry is now updated to be compliant with reference envelopes.**
- Finalizing documentation of understood discrepancies
- **Next Steps with DSCs:** Working towards validation of material budget for supports and service
- **Future Work:** Developing tools for automated validation

Background Mixing in Simulation

- Inner beampipe gold coating thickness increased => Updating synchrotron background samples. Only 10x100 and 10x275 configs available so far
- Work ongoing on background merging after Geant4 simulation step to improve performance

Metadata Propagation

- Making metadata accessible for physics analysis:
  - **Event generator information. Limitation: external generator code.**
  - **Beamline parameters**
  - **Simulation configuration**

WG Highlights: Reconstruction Framework and Algorithms WG

- **EIC Physics Readiness Workshop:** very strong joint PC-SC session joint PC-SC session on last day
  - Waypoint between January CM and July UGM
  - Included updates on key reconstruction priorities

- **Highlights from Workshop**
  - **PID:** strong progress towards IRT2 integration in EICrecon (PR 2515). All PID DSCs working towards big updates at July UGM.
  - **Particle Flow:** all code available, actively being reviewed (links here). Several new institutions joined effort.
  - **Lepton ID:** progress being made towards improving eID in EICrecon (links in slides). Effort towards integrated muID in full swing.
  - **Track Reco:** progress on TOF charge-sharing/clustering (PR 2616) and on SVT noise injection/alignment (PR 2603). Strong progress on understanding background impact on tracking.

- **Other Highlights**
  - **Event Reconstruction:** settled on policy for *ReconstructedParticles* frames, identified next steps (03.09 Reco WG)

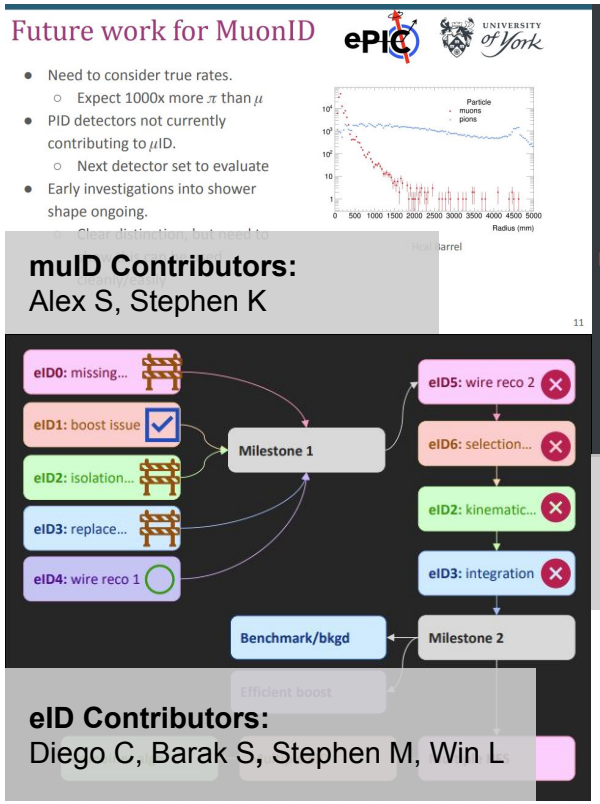

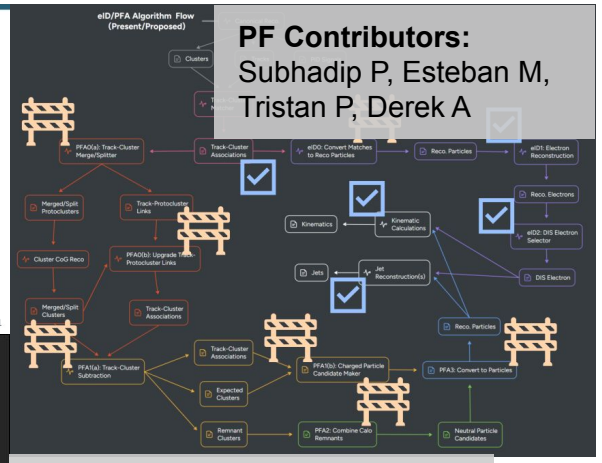

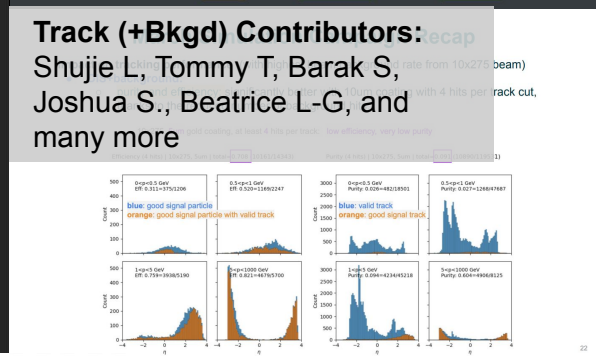

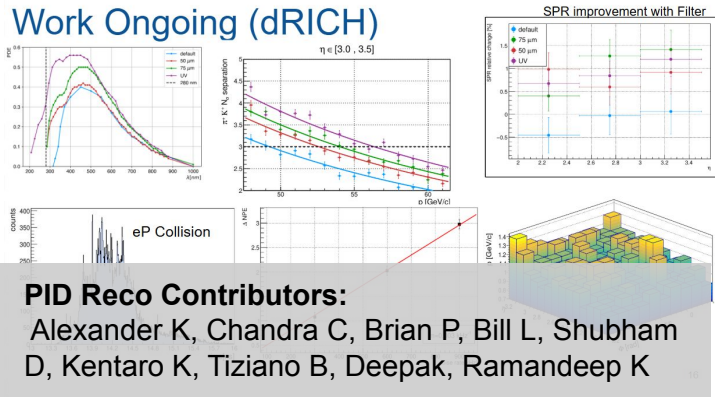

Production WG Highlights

Compute resources are sufficient for campaign demands till the end of 2026.

Compute Resource	Core Hours Per Year
Opportunistic Open Science Grid	~11.5 million
Allocated Resource Accessed via OSG (JLAB, DRAC, INFN)	~2.5 million
Perlmutter CPU Only	3.2 million (added in 2026)
GREX (Canada)	8.8 million
BNL	Opportunistic (added in 2026). 24k CH used so far.

Rucio server at JLab for data management

Storage Resource	Amount (in TB)
JLab	500
BNL	750 (added in 2026)
JLAB Tape	Scales to need

Sufficient storage to hold last 4 regular campaigns and 2 milestone campaigns: pre-TDR (25.10.*) early science (26.03.*).

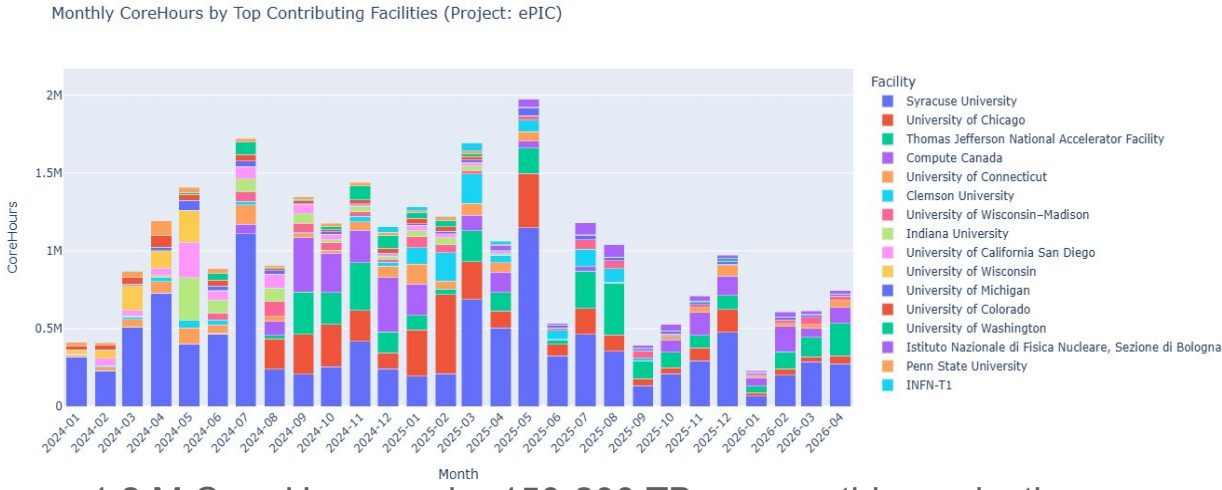

~ 1.2 M Core Hours and ~ 150-200 TB per monthly production campaign

Production WG Highlights

Datasets are processed in terms of priority from the Overview Tracking Document.

Must use Request form for submitting requests.

Requests must follow input processing guidelines and arrive before monthly release.

Current Campaigns In Flight: March 2026 and April 2026

Background mixed files only available for 10x100 and 10x275 beam energies.

Rucio Metadata Support

Simulation files are now findable using metadata tags like energy, event generator and requester PWG. This done via Rucio: Tutorial Link

Next steps:

- Work with PWGs and DSCs to identify common needs and minimal statistical requirements to avoid duplicate entries in overview tracking document.
- Streamline dataset request process.

Production WG Highlights

PanDA Workload Management System instance and Physics Configuration System being developed at BNL for automated productions in FY 26-27

PanDA WMS with iDDS

- PanDA WMS handles:
  - Task brokerage
  - Scalability
  - Adaptive Scheduling
- iDDS (Intelligent Data Delivery System) defines:
  - Data transformation
  - Task dependency chain

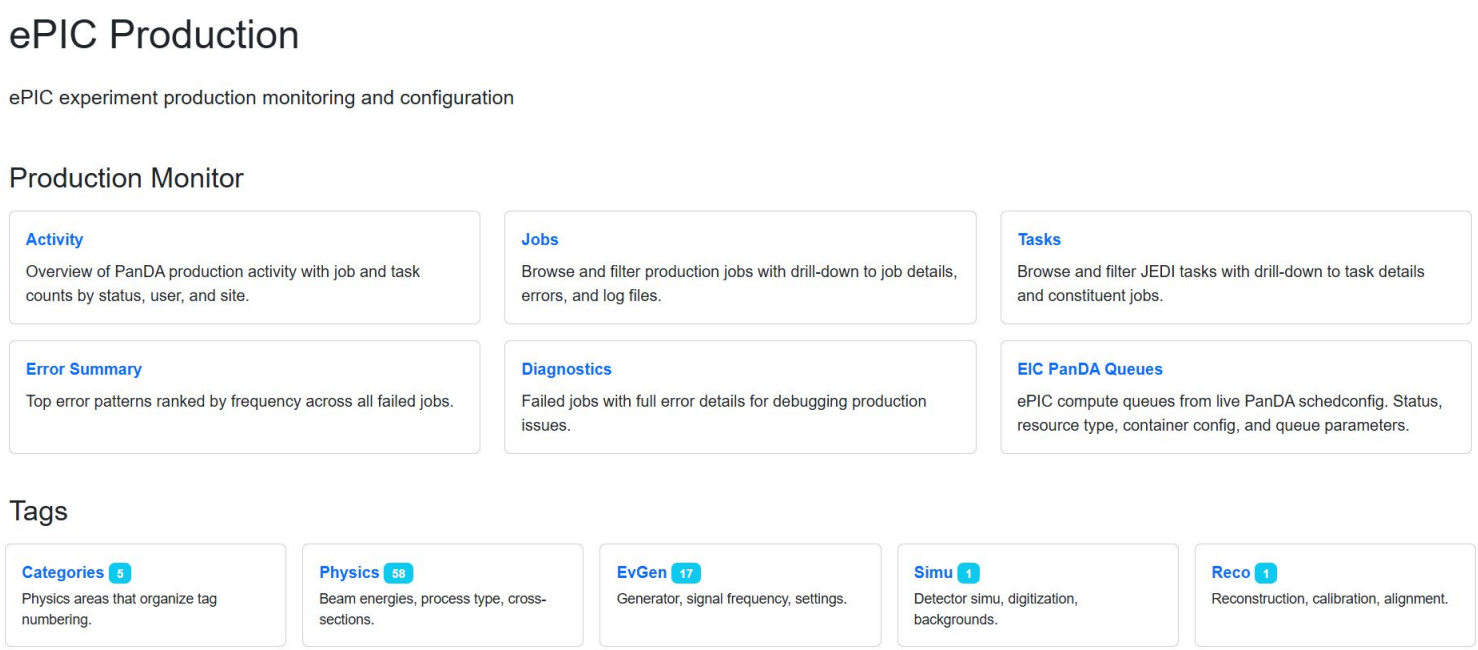

Physics Configuration System

- Assists in generating and tracking exactly what submission parameters were used in monthly campaign tasks.
- Direct job injection into PanDA planned for long term after testing period.

WG Highlights: Validation

Hydra

- Monitoring QA plots produced from simulation campaigns is critical for communicating the physics reach of the ePIC detector and for identifying potential issues in simulation.
- These plots provide essential insight into detector performance, reconstruction quality, and analysis readiness.
- Hydra allows collaborations to find, view, compare, and classify image data
- Sample BIC TDR plots are available in Hydra
  - More will be added

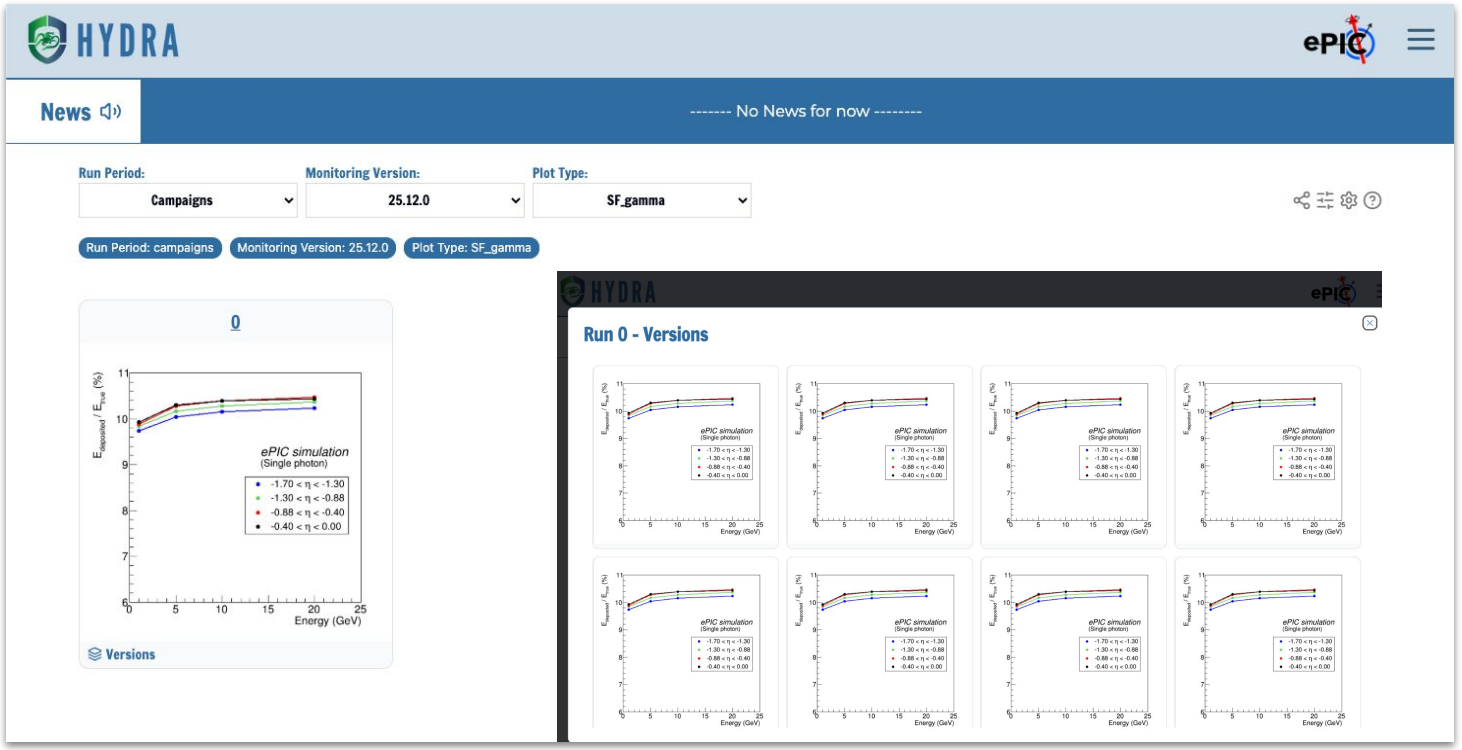

User Learning: Discoverability and Reusability Workshop on April 29th

Register at https://indico.bnl.gov/event/31760/overview

ePIC User Learning Workshop: Discoverability and Reusability for the preTDR and Beyond

April 29, 2026
Online
US/Eastern timezone

Enter your search term

- Overview
- Timetable
- Contribution List
- My Conference
  - My Contributions
- Registration

The ability to easily discover data, locate analysis code, and reuse existing workflows is critical for the success of the ePIC experiment. This workshop is a dedicated forum for users and developers to bridge the gap between software infrastructure and physics analysis.

We will present the current strategies being deployed to make ePIC software more "discoverable"—thereby bridging the gap between user needs and system capabilities, thus ensuring the right tools are easy to find and adopt. As our software is user-centered, we need community engagement. This workshop is designed as an open dialogue to gather feedback on what is working, what isn't, and to collectively refine the priorities that will carry us through the preTDR and beyond.

Starts Apr 29, 2026, 8:30 AM
Ends Apr 29, 2026, 1:00 PM

Online

User Learning: Updated Landing Page at https://eic.github.io/

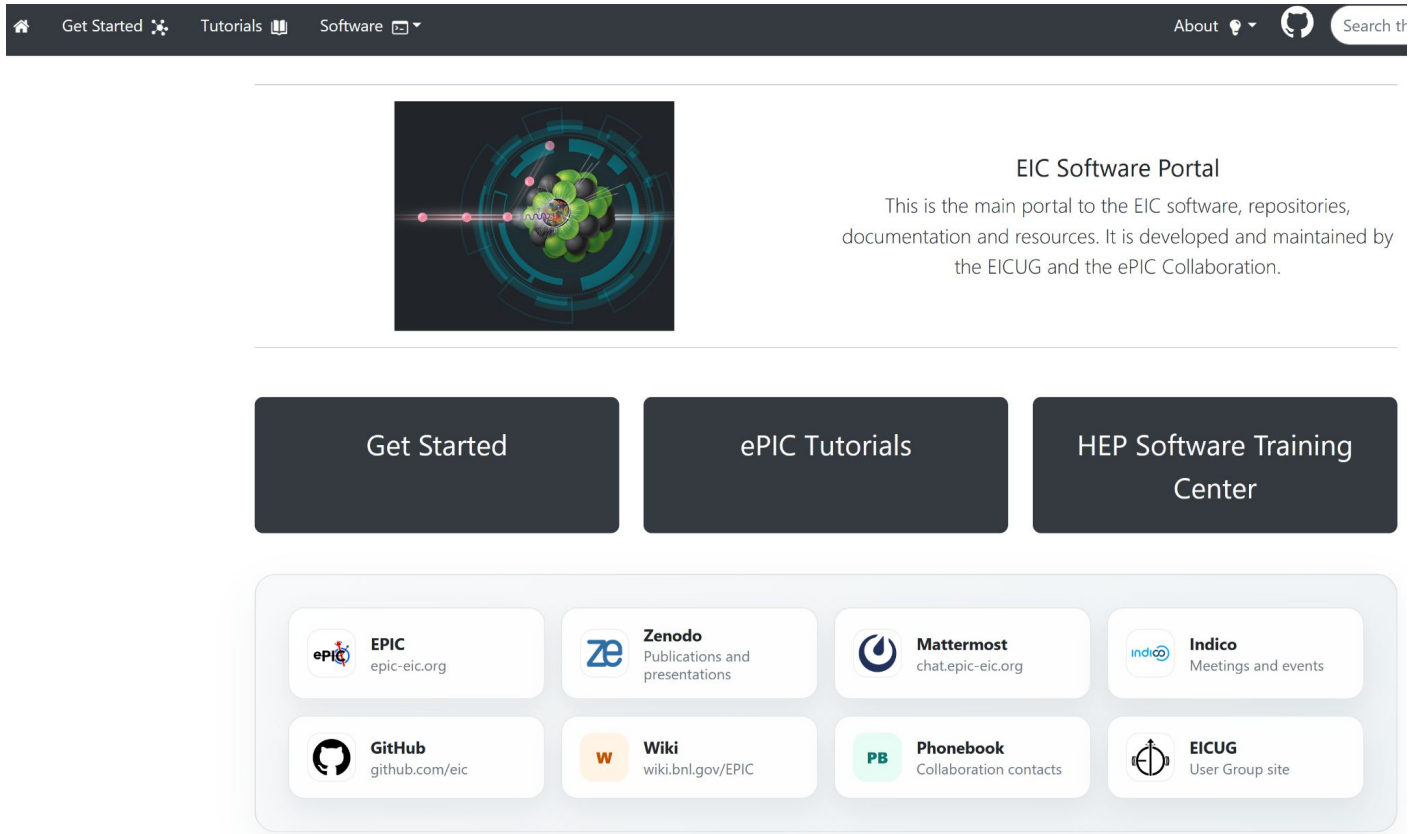

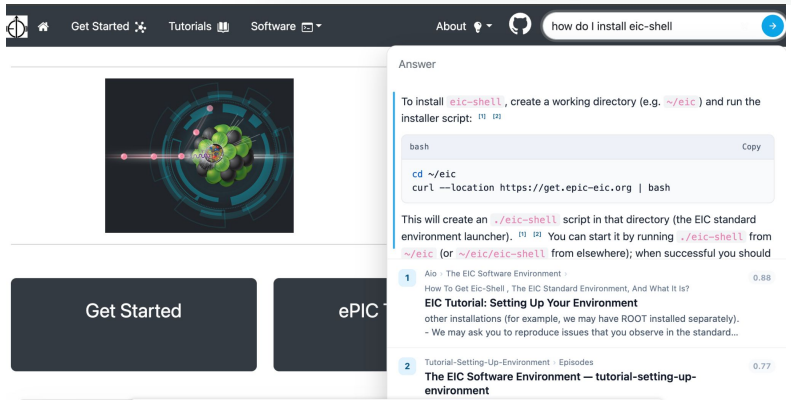

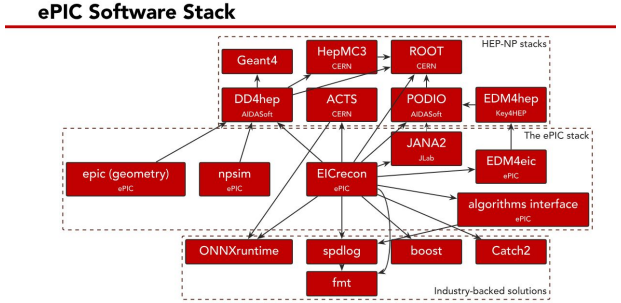

Key features:

- Renders nicely on mobile devices
- Clear buttons to direct and onboard new users
- One-stop page with links to other ePIC resources for computing and involvement
- Searchable sidebar!
- Tutorial page re-organized for improved access:

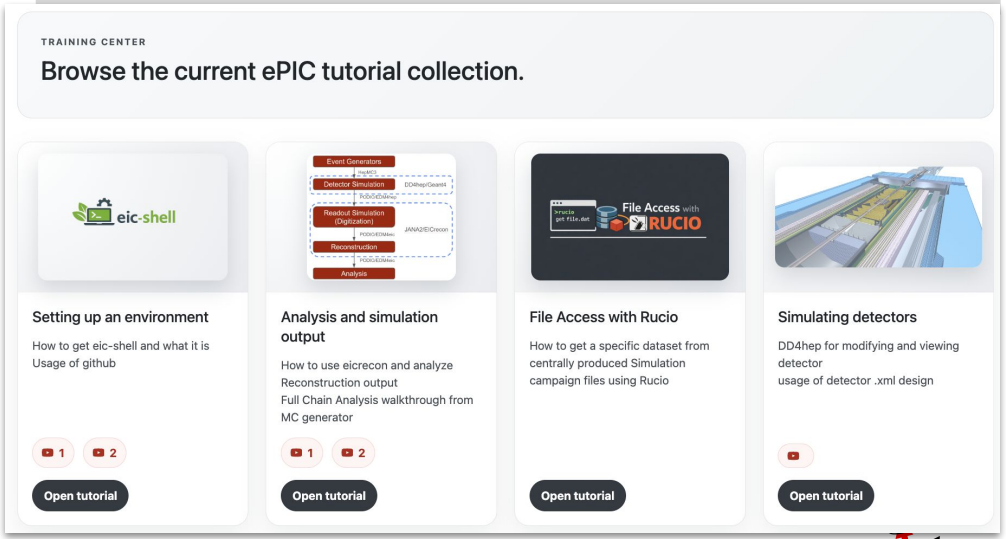

User Learning: Training Opportunities!

Upcoming tutorial on File Access with Rucio:

- Use of Rucio to find ePIC simulation data, including the usage of metadata tags. Explores specific cases.
- 05/05/26, 0930-1130 EDT (GMT-4) https://indico.bnl.gov/event/32515/
- Supports file storage locations (both JLab and BNL)

EIC Tutorial: File Access

Welcome to the EIC/ePIC tutorial on file access. This tutorial will explore how to find and access EIC/ePIC simulation data using Rucio. Usage of metadata tags to quickly find data will be explored as well as some example use cases.

Schedule

	Setup	Download files required for the lesson
00:00	1. Introduction	How are EIC/ePIC simulation outputs organised?
00:10	2. Rucio Usage	How can I use Rucio?
00:40	3. Use Cases	How do users interact with EIC/ePIC data?
01:30	Finish	

Future tutorials planned:

- Scavenger Hunt tutorial soon
- 2026 CFNS Summer School on the Physics of the Electron -on Collider (EIC) hosted by the Center for Frontiers in Nuclear Science (CFNS) at Stony Brook University, USA, June 1-12, 2026 https://indico.cfnssbu.physics.sunysb.edu/event/604/

- EIC Early Career Workshop in Glasgow, July 2026

Roles

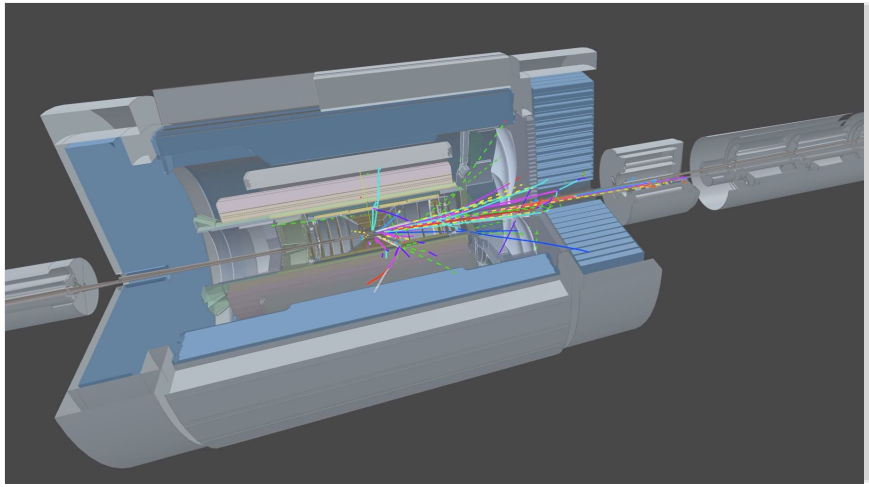

Enable Science

- Development, operations, and infrastructure for preTDR and future milestones.
- **Activity:** Monthly software releases and simulation campaigns tailored to collaboration needs.

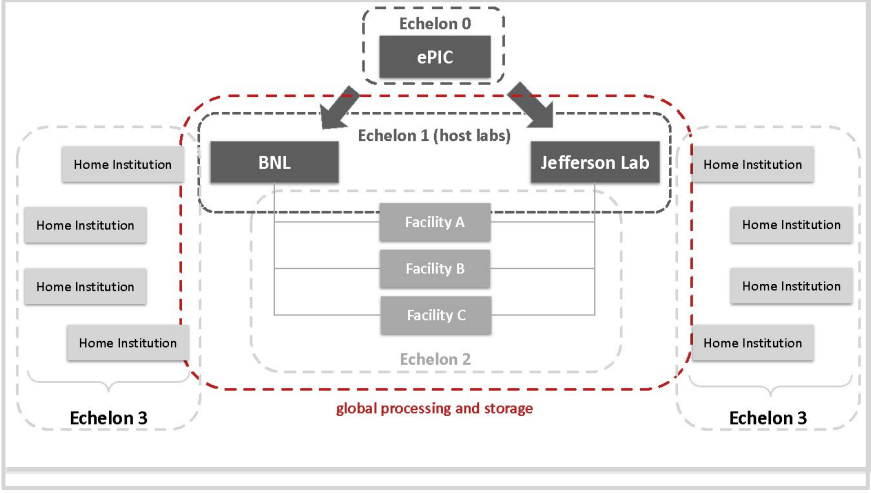

Compute–Detector Integration

- Maximize and accelerate science with the ePIC computing model
- **Activity:** Define and test DAQ–computing interfaces to mitigate integration risks. Develop and deliver a functional testbed to validate workflow management for autonomous detector calibration

Streaming DAQ and Computing Milestones

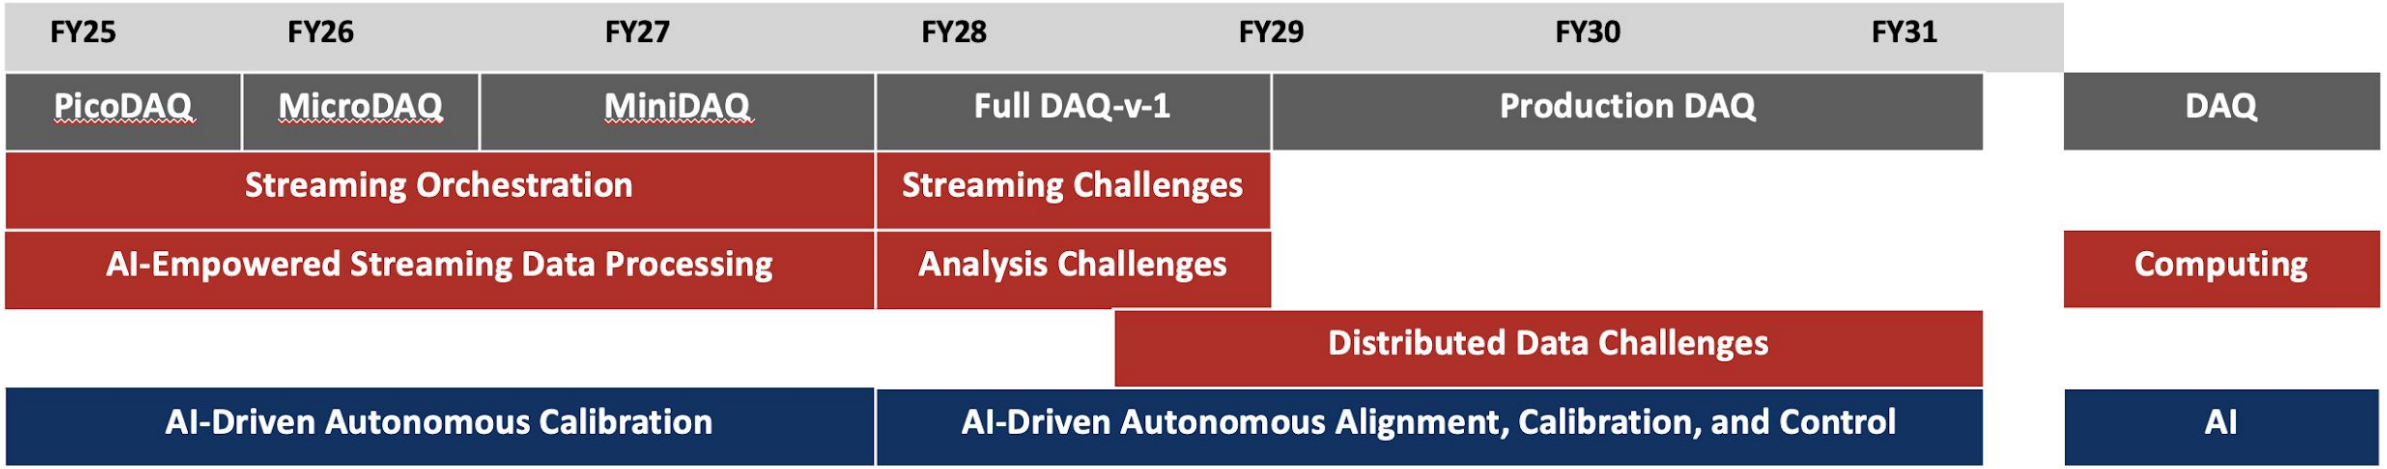

- **Compute-Detector Integration**:
  - Joint deliverables between **DAQ** and **computing** to develop integrated systems for detector readout, data processing, and ultimately physics analysis.
  - **Key role of AI(/ML)**: Empowering data processing and enabling autonomous experimentation and control.

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

WG Highlights: Streaming Computing: Workflow Orchestration

Quick mention of two current projects (among others previously presented)

Fast TF streaming workflow: The reco payload

Fast TF streaming workflow

- Fast Processing Agent
  - Run_imminent
    - Send config to iDDS to set up PanDA workers (e.g. how many workers)
    - iDDS creates the PanDA workflow
  - TF processing
    - Fast proc agent slices STF samples down to smaller TF ranges scaled for quick parallel worker processing: deliver result for control room fast
  - Streaming slices to workers
    - A PanDA worker loop
      - Fetch one slice message
      - Process the payload
      - Send back the result in a message
      - Ack the message to indicate processing complete

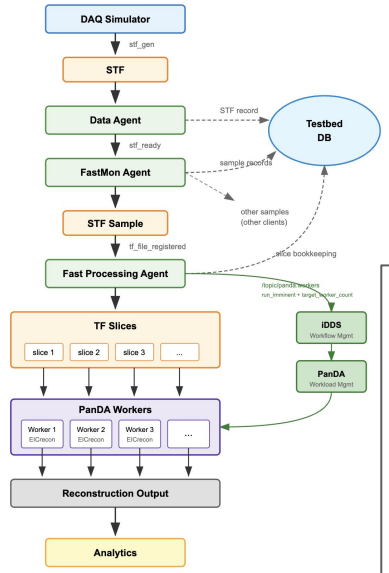

Fast processing TF-based reco payload

Current work: persistent reconstruction process needed to provide **low-latency** reconstruction, receives fine-grained fast processing tasks (TF ranges)

- EICrecon + testbed work

Design document for payload with key decisions

- Define responsibilities of payload and the persistent reconstruction process
- Bi-directional communication
- Input TF slices are queued for reconstruction
- Ready output files are reported back to the payload
- Transport: ZeroMQ over UNIX socket
- JSON messages

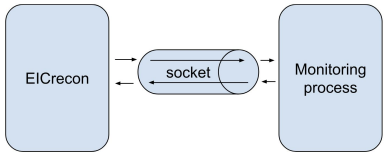

Calibration workflows at the testbed

- Build calibration workflow prototypes in the testbed
- EEEMCal detector is first target
- In progress: Support for standard Snakemake workflows definitions in PanDA, as basis for calib
  - Also a stepping stone to PanDA-based CI
- Soon: Example calibration/monitoring payload with a proto-interface to calibration DB

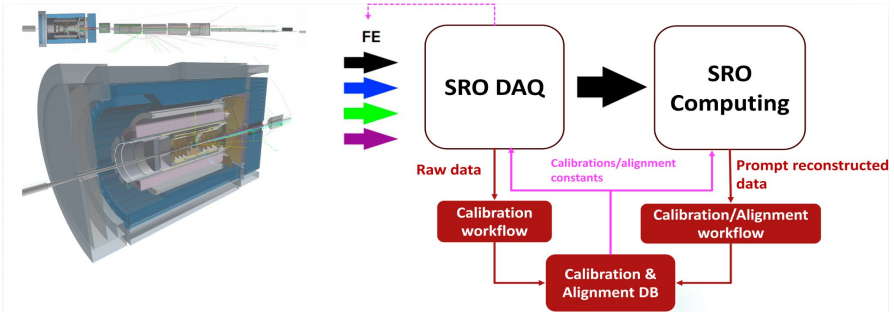

Infrastructure News: MCP Tools and AI Services

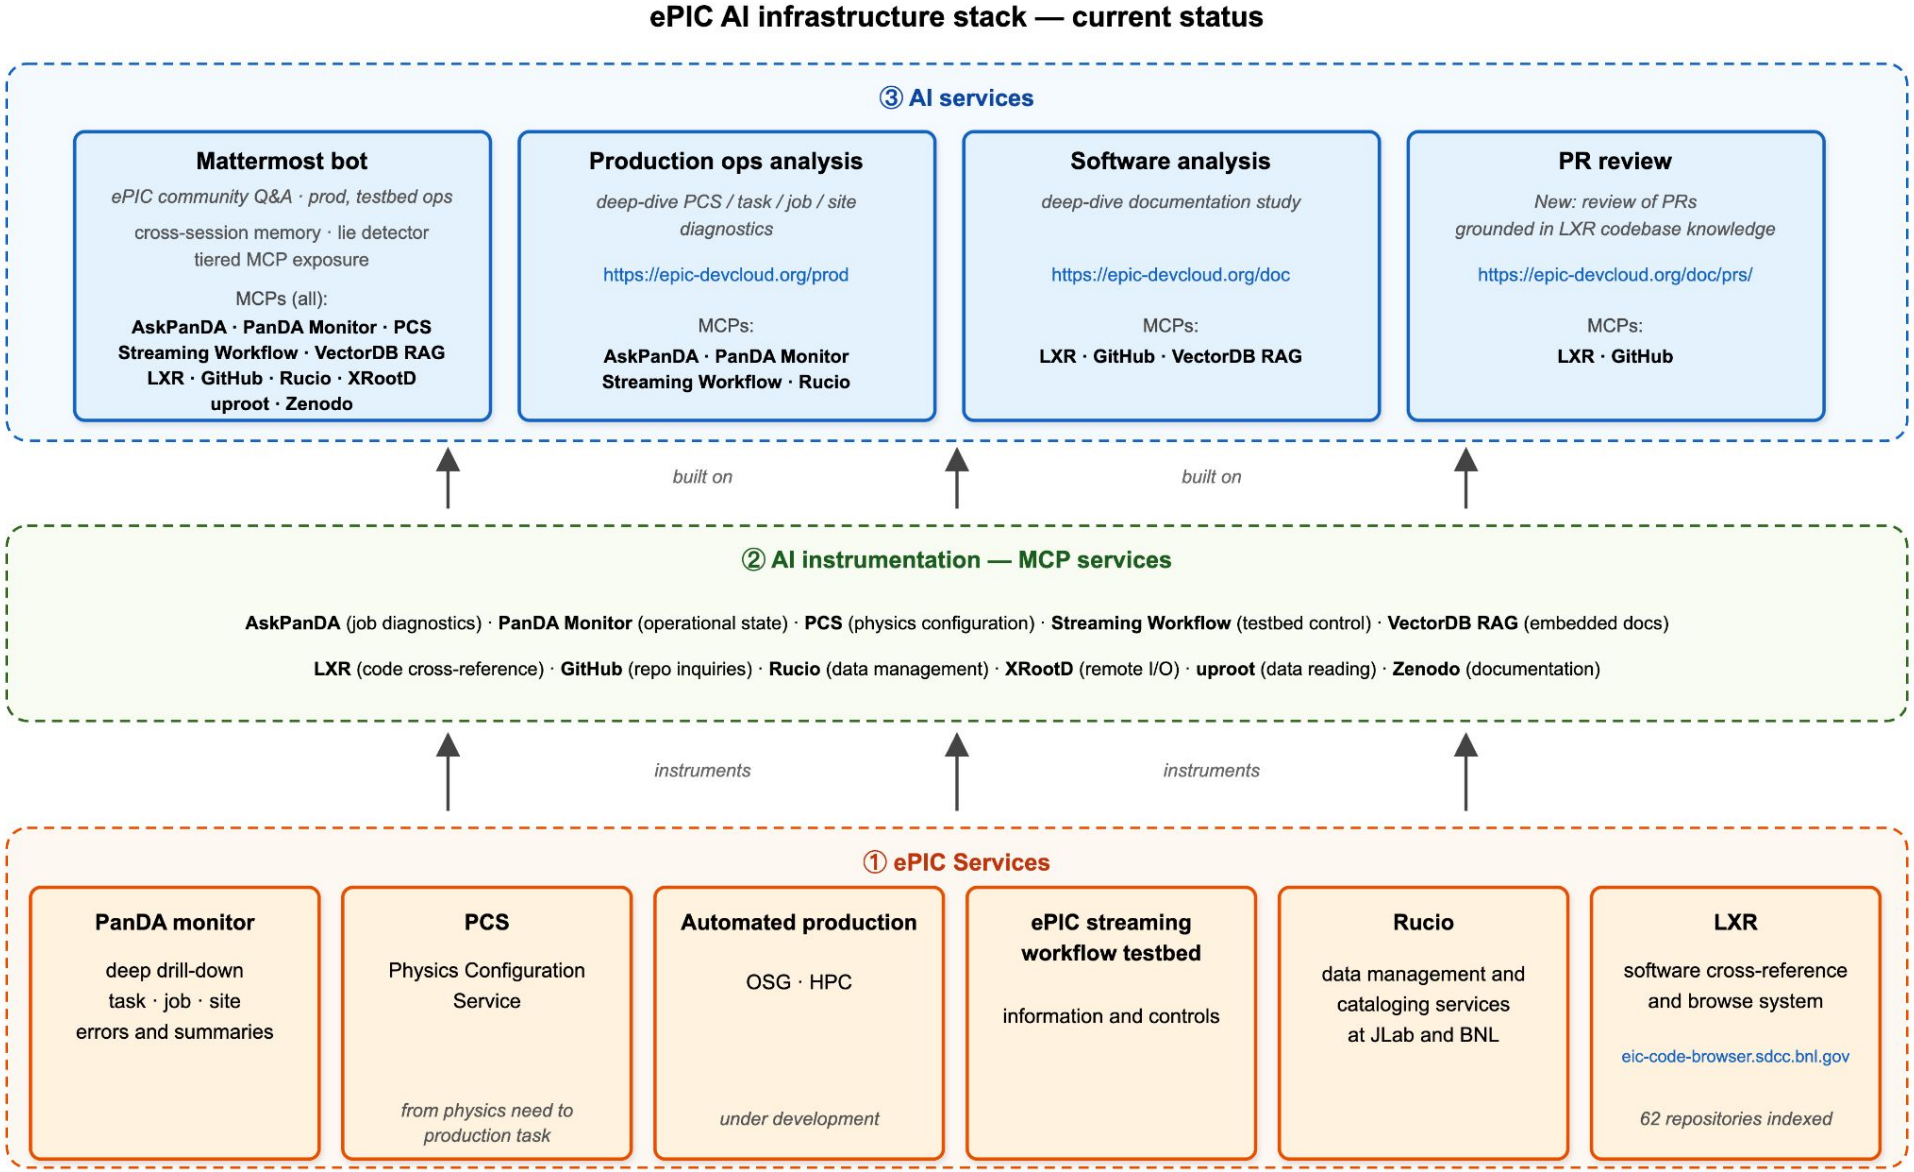

EIC International Computing Organization

- **EIC International Computing Organization (EICO)** coordinates international computing resources.

Activities:

- **Echelon 2**: Roles and resource requirements for the June EIC RRB meeting.
- **Networking**
  - **ESnet**: Update EIC Case Study (this year).
  - Evaluate reuse of LHCONE vs. dedicated network (May).

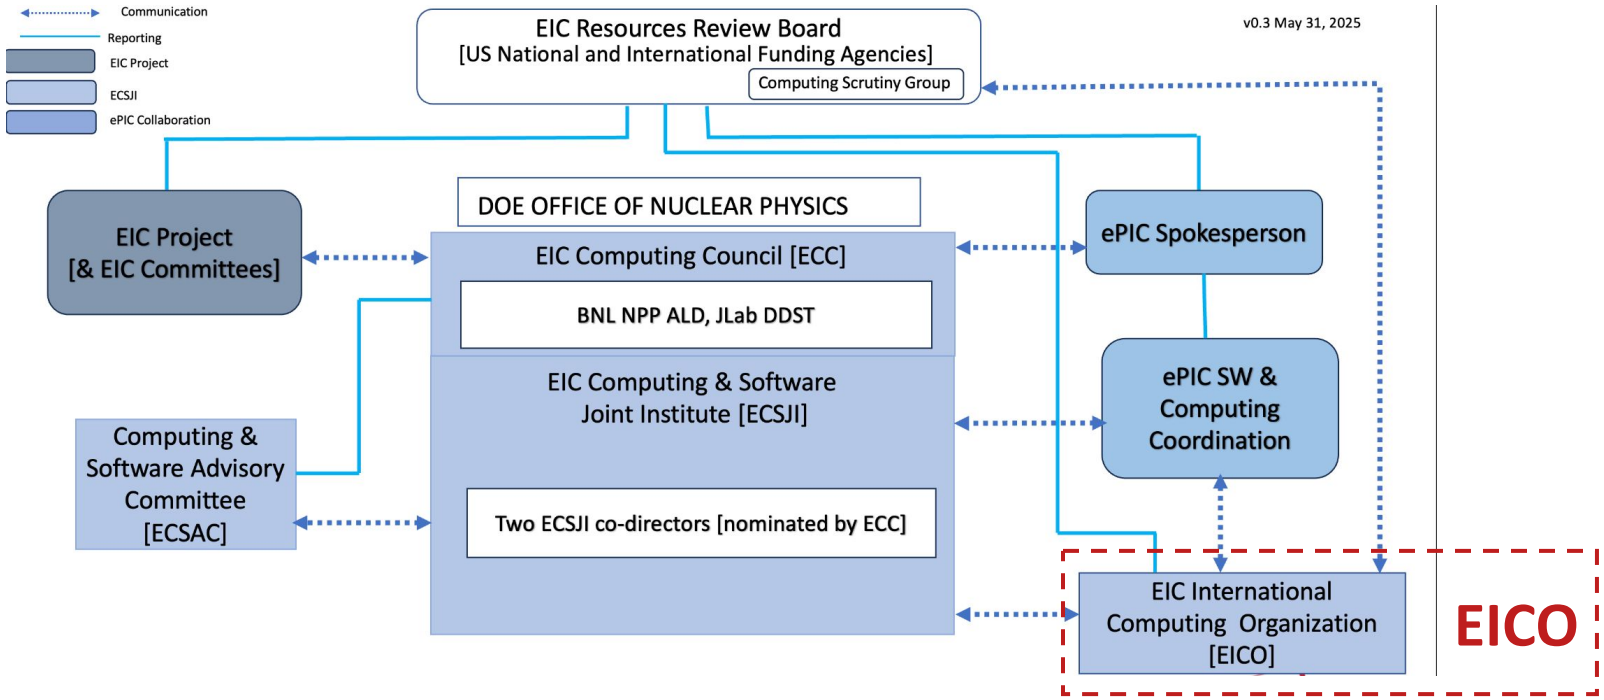

Great Software for Great Science

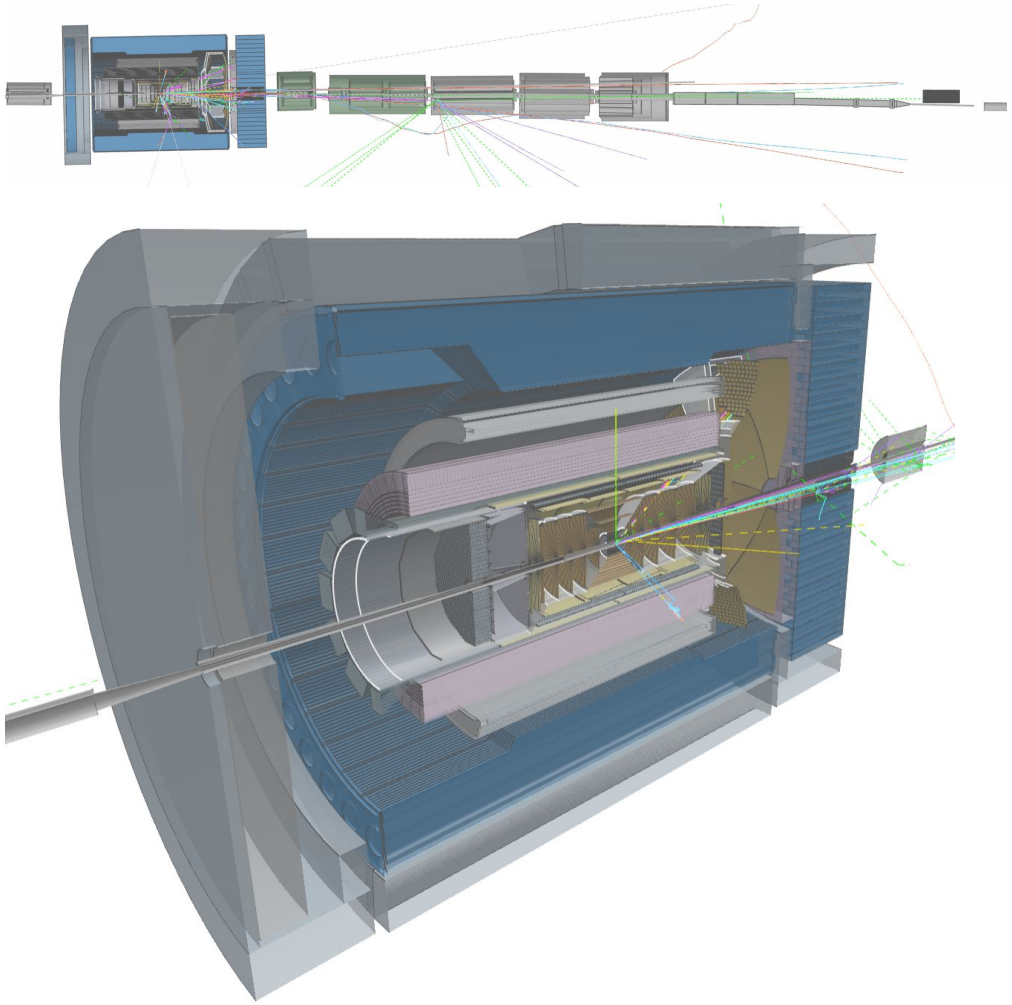

Captured from the ePIC Event Display

- **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**.

Evolution of ePIC Software & Computing

We aim for **collaborative software development across DSCs and WGs**, building a vibrant developer community in ePIC that advances our software stack.

We are improving the review process to make contributions easier and are establishing a **formal review team with clear guidelines for PRs and reviews**.

Generative AI presents an opportunity to lower barriers to contribution, while introducing **ethical, legal, and technical risks**, which we will address through a lightweight policy.

Following the formation of the review team, we will introduce a **formal release manager role**.

Three-Year Software & Computing Plan

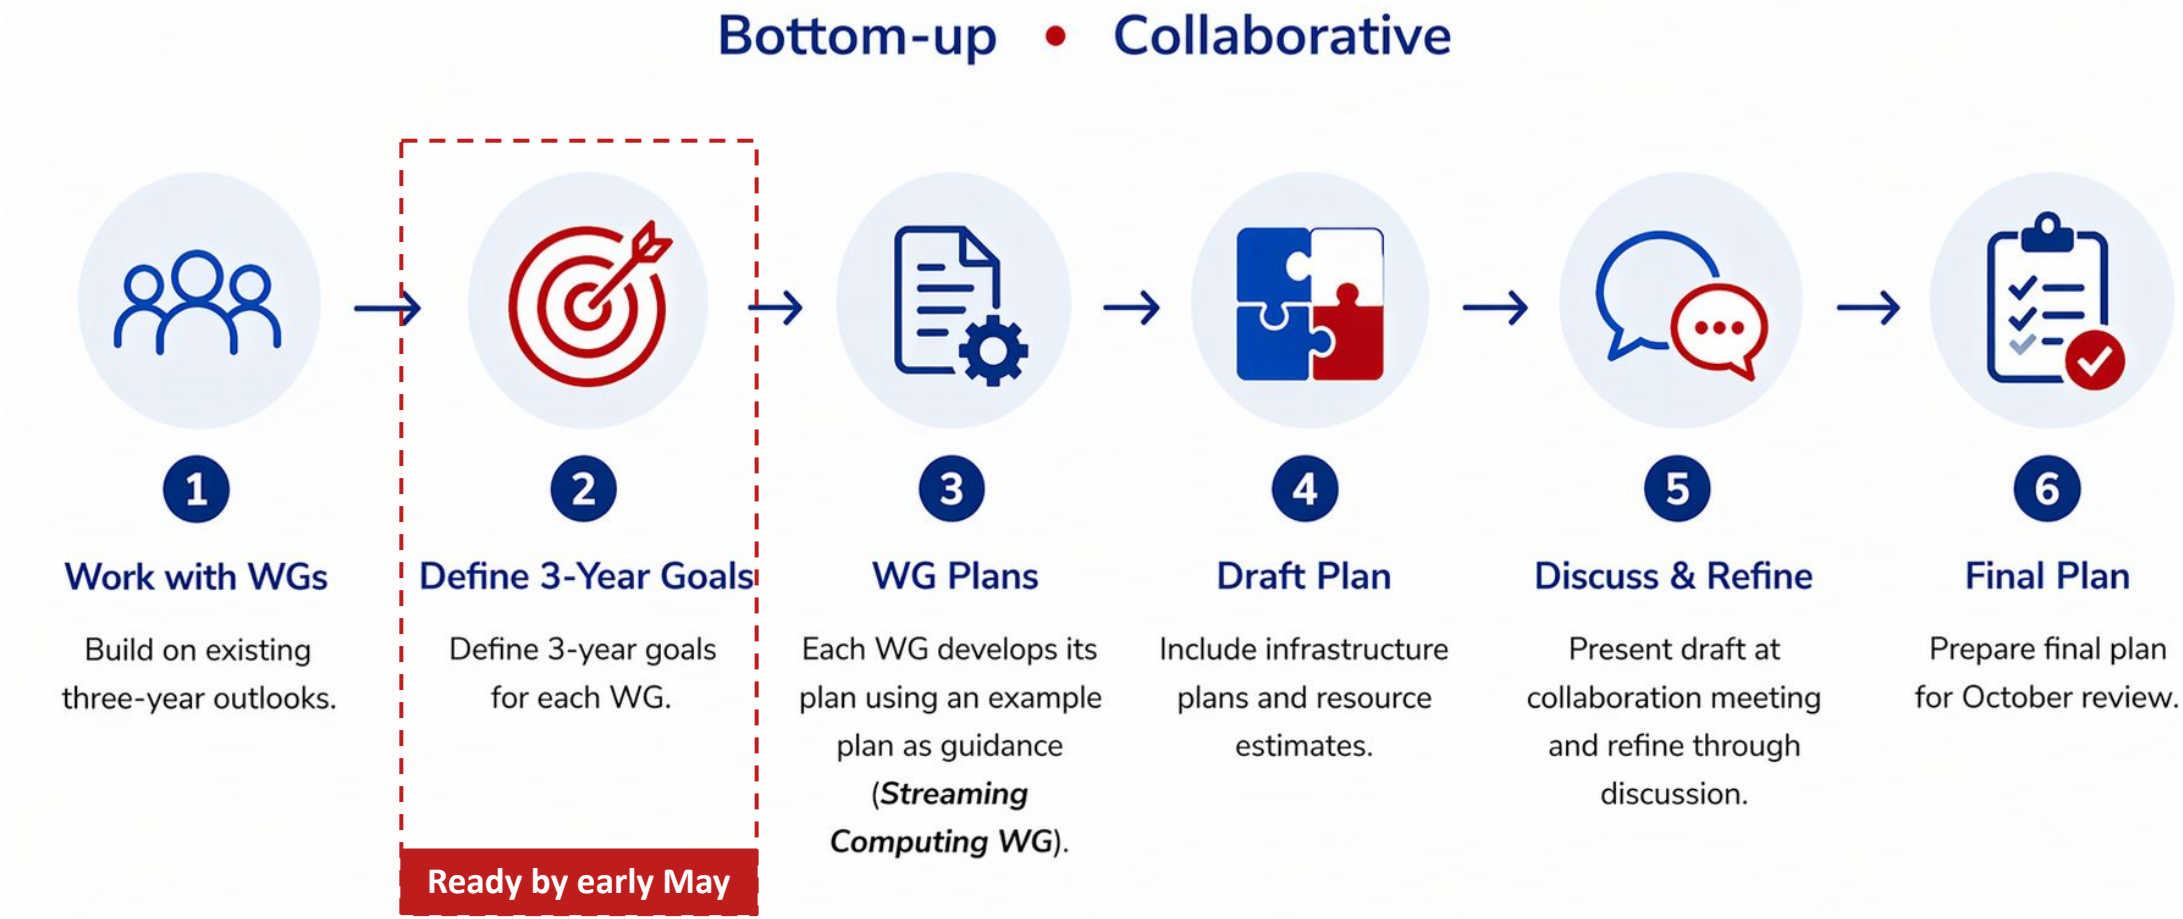

Schedule

ePIC Software & Computing Meetings

04/29 ePIC User Learning Workshop: Discoverability and Reusability for the preTDR and Beyond

05/06 Software News

05/13 ePIC Workshop on Heterogeneous Computing**:** First of series.

05/20 Software News

05/27 Joint Meeting with PWGs

06/10 Root: RNTuple, RDataframe

Other Meetings

05/18–05/19 ECSJI and EICO Annual Meeting

10/13–10/14 ECSAC Review

For additional details on our meeting schedule, including meetings of various WGs, please follow our Indico category: https://indico.bnl.gov/category/410/