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, Kolja Kauder (nominated), 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**: Andy Buckley (nominated), Torri Jeske.

Roles

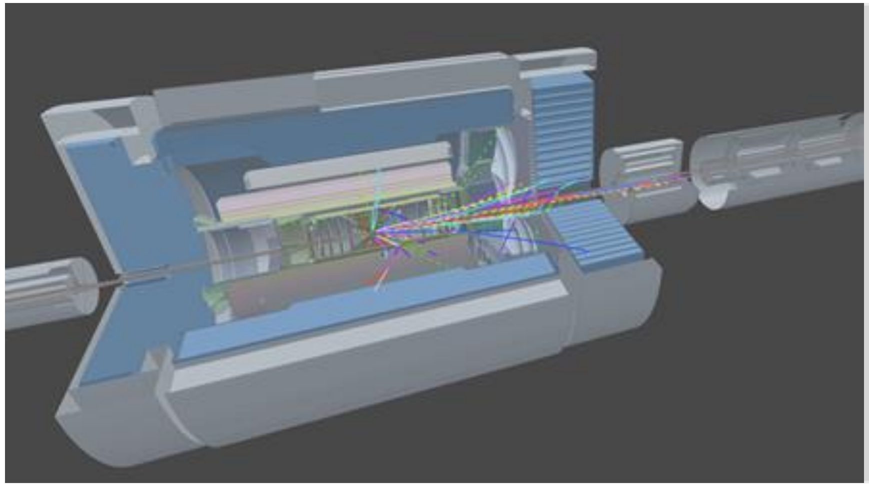

Enable Science

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

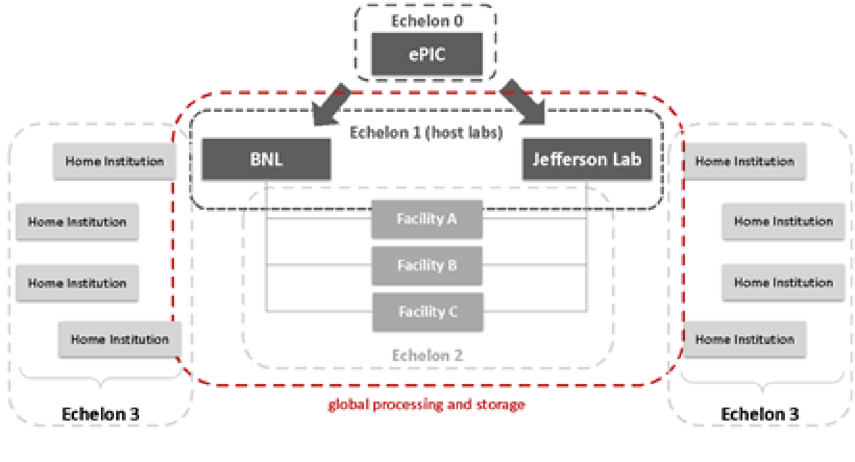

Compute–Detector Integration

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

The ePIC Streaming Computing Model

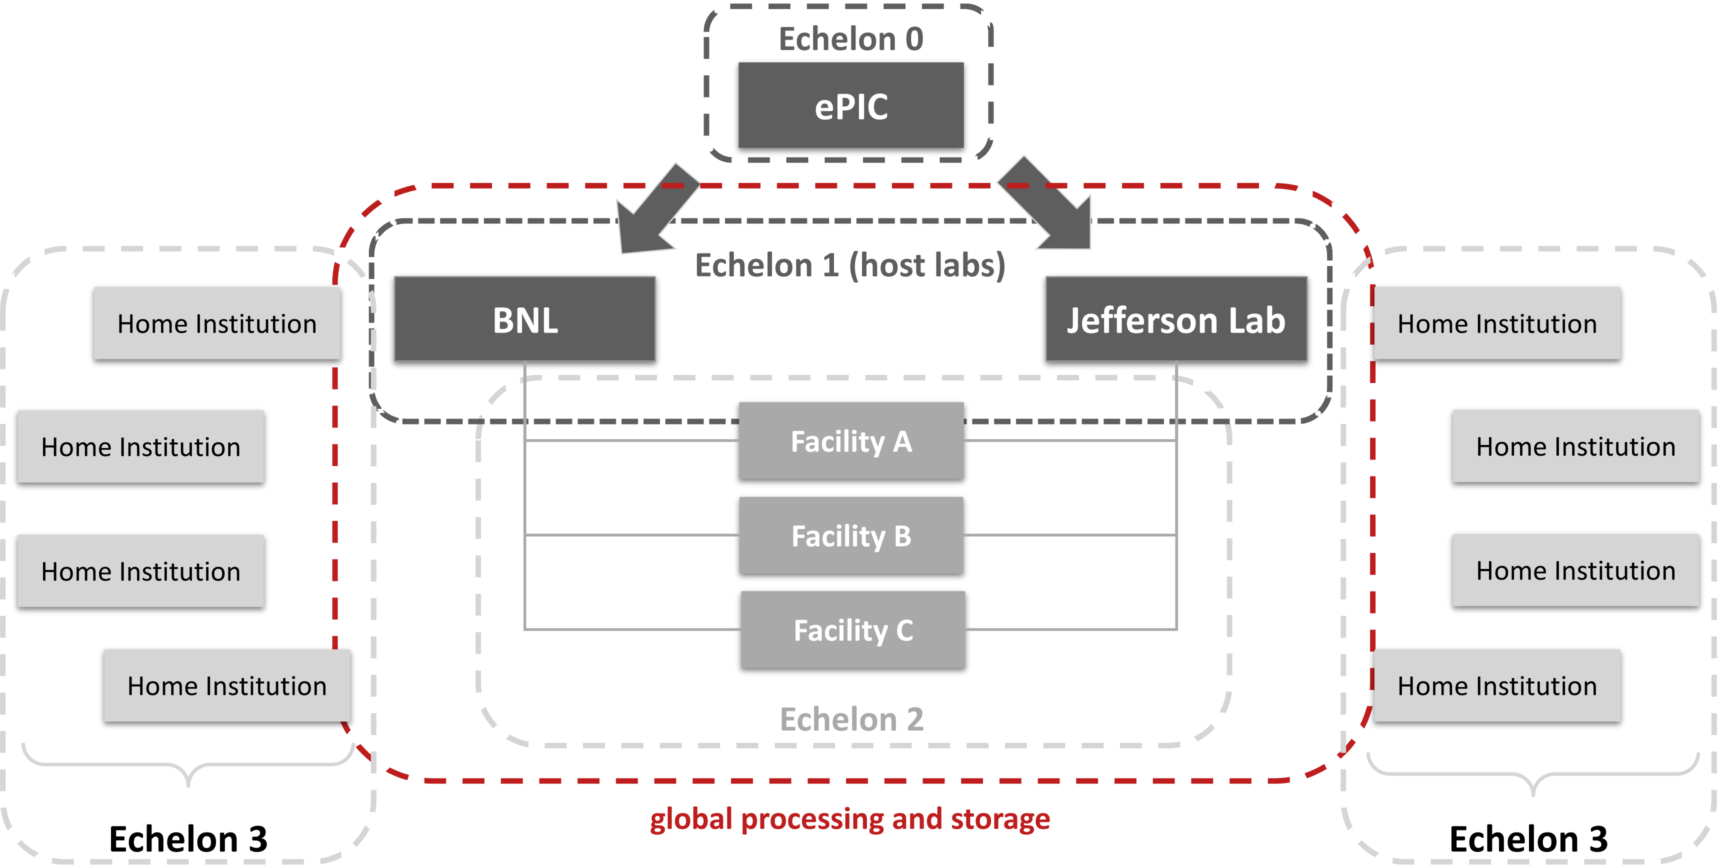

Successfully reviewed in 2023–2025 and documented

The ePIC Streaming Computing Model
Version 2, Fall 2024

Marco Battaglieri[1], Wouter Deconinck[2], Markus Diefenthaler[3], Jin Huang[4], Sylvester Joosten[5], Dmitry Kalinkin[6], Jeffery 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.
[6]University 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.**

https://doi.org/10.5281/zenodo.14675920

The ePIC Streaming Computing Model is **physics-driven**. It has been developed to **accelerate the pace of discovery** and **enhance scientific precision** through streaming readout, AI, and distributed computing.

ePIC in the Global Particle-Physics Computing Landscape

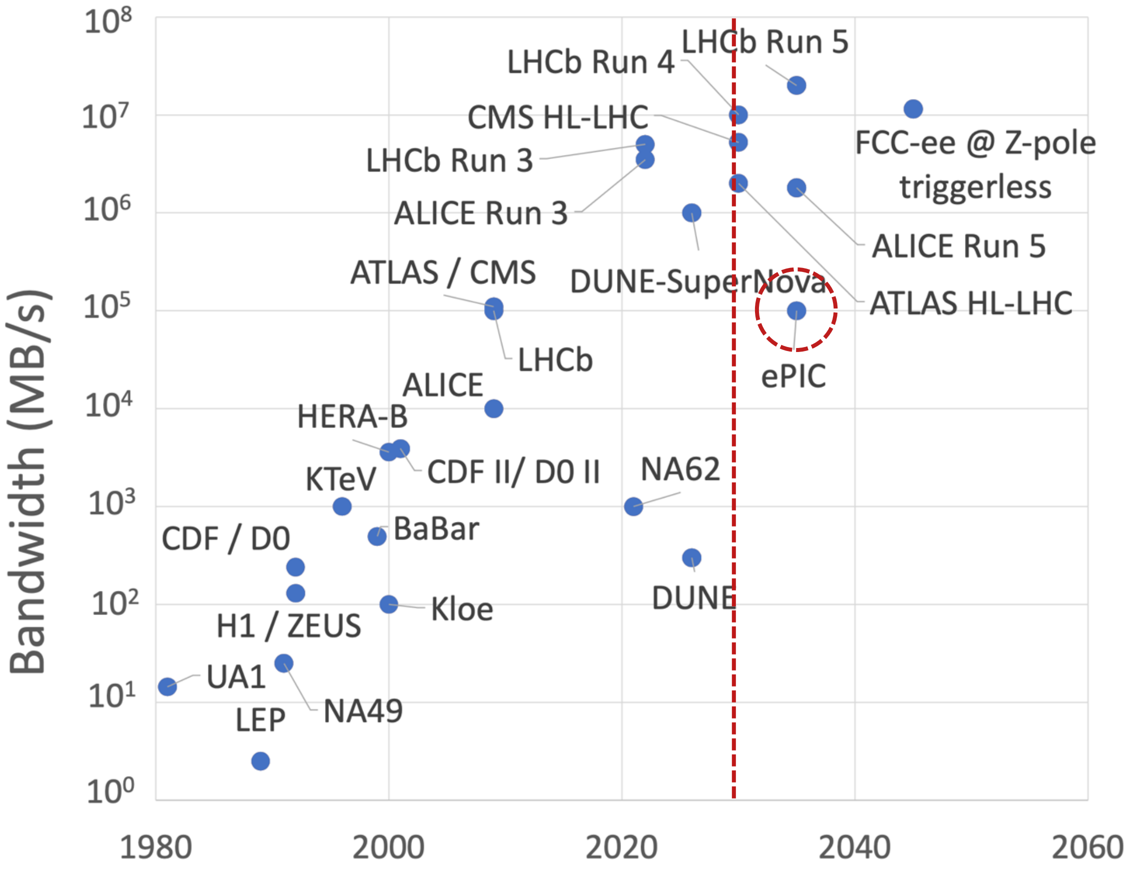

From Open Symposium on the European Strategy for Particle Physics

From CHEP 2026 keynote

The computing **model** is generating **strong interest** within the global particle physics community, with our **leadership recognized** through advisory roles in the HSF, HL-LHC Computing Review Panel, and the LHC Experiments Committee.

What the ePIC Streaming Computing Model Enables

The ePIC Streaming Computing Model integrates streaming readout, AI, and distributed computing.

1. **Streaming Readout**
   Maximizes scientific reach through unbiased event selection and continuous measurement of backgrounds and detector noise.

2. **AI**
   Accelerates science through automated alignment, calibration, and validation, shortening the path to physics results and publications.

3. **Distributed Computing**
   Enables flexible access to dedicated and opportunistic computing resources where they are needed.

Closed-loop experimental control

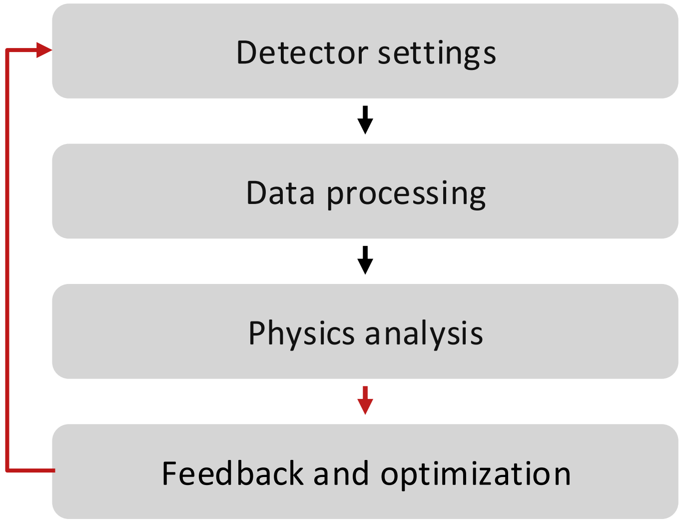

Investing now in the ePIC Streaming Computing Model ensures that the collaboration will have the infrastructure, resources, and workflows needed to **commission and operate the experiment**.

Building the ePIC Streaming Computing Model

- The **Streaming Computing WG** coordinates the development and implementation of the ePIC Streaming Computing Model, including the required computing infrastructure.

- The collaboration-wide survey showed strong interest in learning more about **streaming readout and computing**.

- We are addressing this request in this talk and at the **Electronics, Readout and DAQ and Streaming Computing WGs Joint Workfest**:
  - Wednesday session,
  - Thursday session.

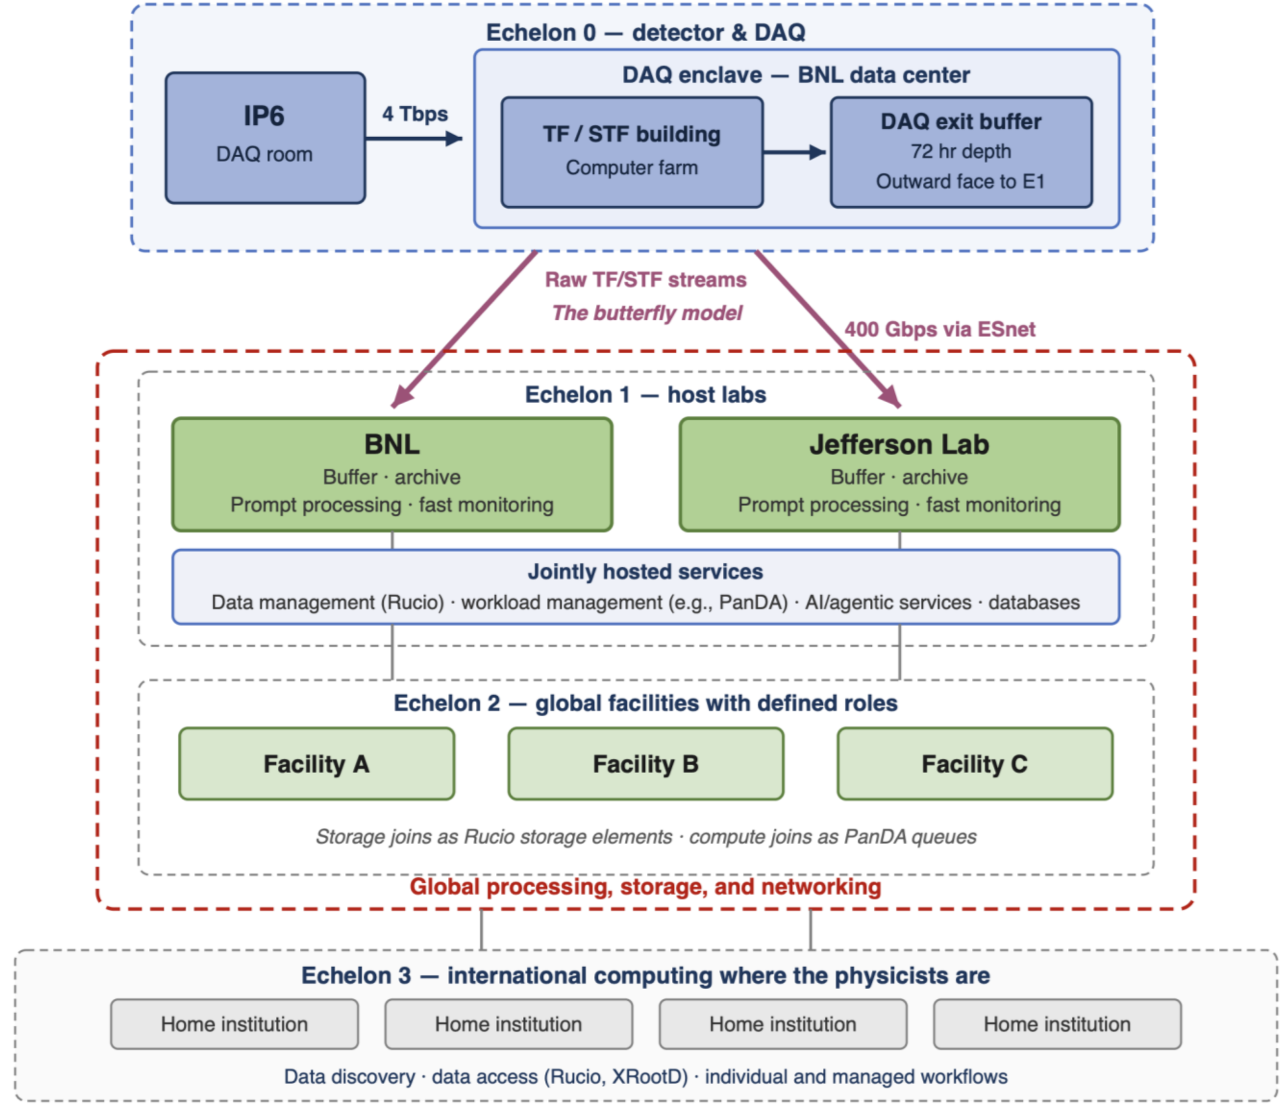

Streaming DAQ and Computing Milestones

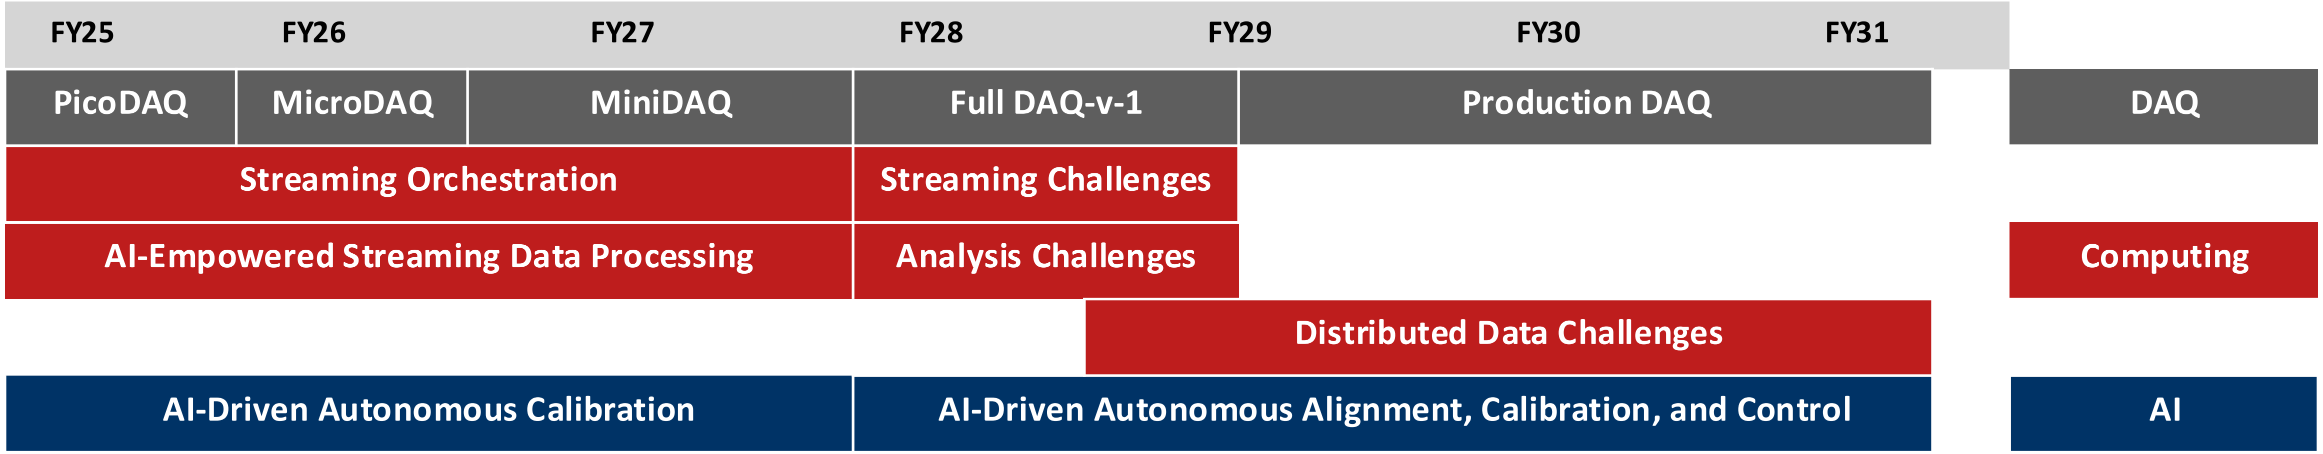

- **Compute-Detector Integration**:
  - Joint deliverables between **DAQ** and **computing** to develop an integrated system.
  - We define and test DAQ–computing interfaces to reduce integration risk and enable early system testing.
  - **FY28Q1 priorities**: Testbeds and prototypes for streaming orchestration and autonomous detector calibration for one detector system
- **Integrating AI with ePIC Detector**:
  - AI has been built into the ePIC computing model from the outset.
  - We are developing AI-driven workflows for alignment, calibration, validation, control, and diagnostics.
  - This enables seamless processing from detector readout to physics analysis.

Streaming Orchestration: BNL Testbed

Rapid Progress

- Build streaming orchestration with **PanDA** and **Rucio**
- Enhance PanDA with **MCP Tools** and **AI Services** for autonomous workflows

Prototype workflow

- Automated simulation campaigns
- **More details**: Slides 17–24

Components of the BNL testbed implementing streaming orchestration

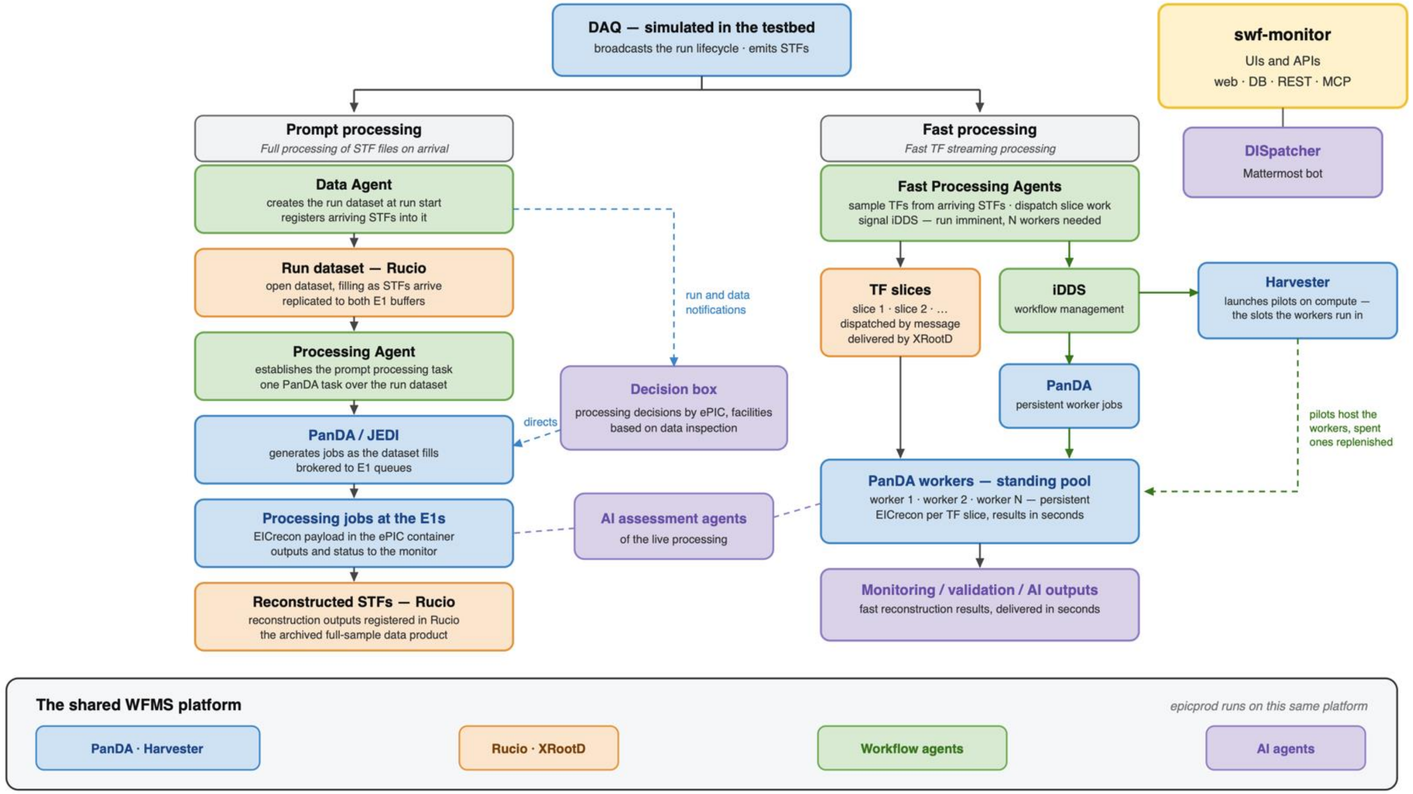

Near-term milestone of automated simulation campaigns will **exercise streaming orchestration with PanDA and Rucio** and clarify where ePIC streaming workflows differ from traditional offline processing.

Autonomous Calibration as a Data Lifecycle Demonstrator

HPDF will be a cutting-edge high-performance data platform that provides advanced storage, processing, and analytical services, empowering data-intensive research and AI-driven discoveries.

In collaboration with the **HPDF Project**, we are demonstrating how **AI-ready data lifecycle management** and FAIR principles can enhance **autonomous calibration workflows** and their **data pipelines**. This is key in the era of **AI and data**, and complementary to the BNL testbed.

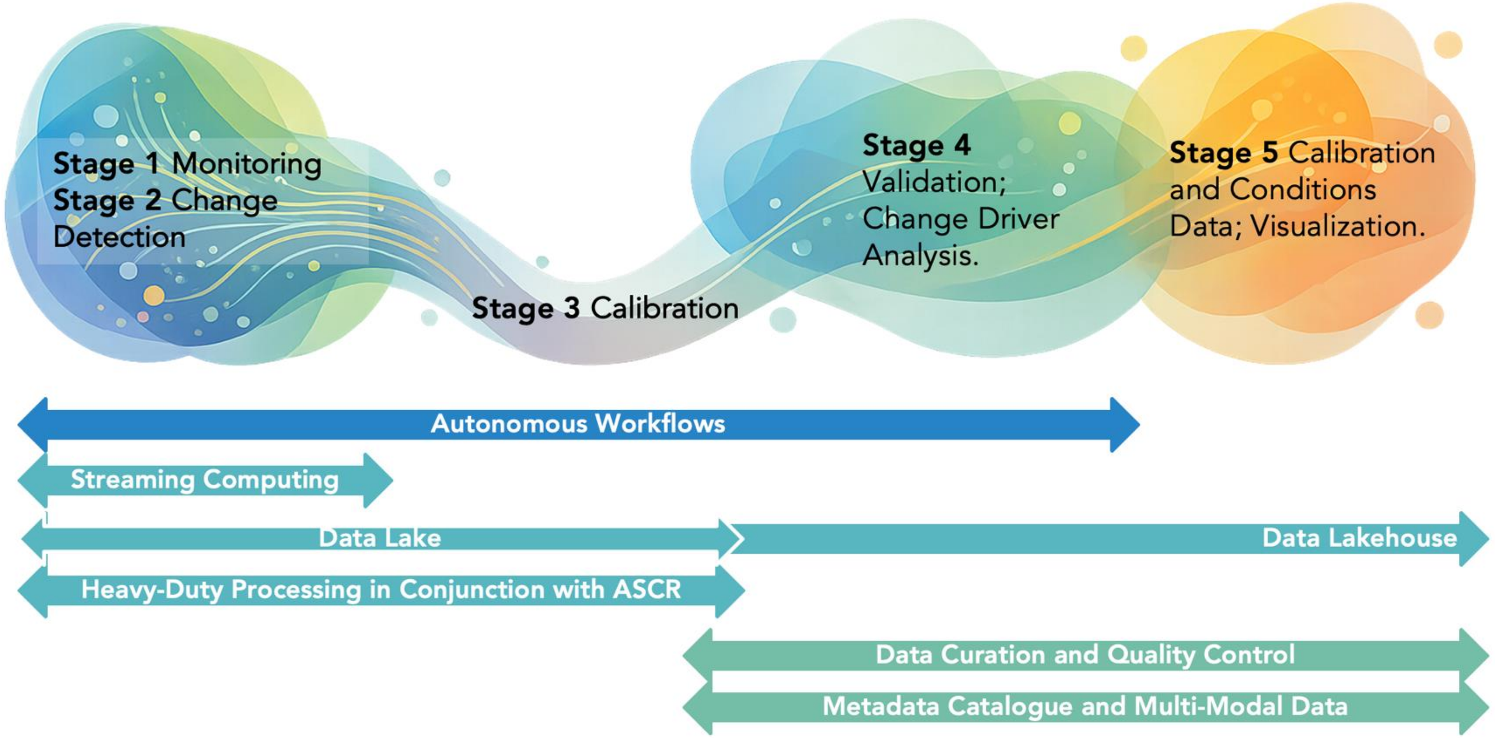

Embed FAIR principles across the data lifecycle

1. Data Capture
2. Initial Processing
3. Central Processing and Enrichment
4. Storage, Curation, and Archiving
5. Dissemination and Visualization

HPDF **Features**

Distributed Computing for ePIC

- **Echelon 1** sites uniquely perform the **low-latency streaming workflows**.
- Apart from low-latency, **Echelon 2** sites fully participate in use cases and **accelerate** them.

Strong international interest and engagement:

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		✓	✓	

Substantial role for Echelon 2 in preliminary resource requirements model.

Assumed Fraction of Use Case Done Outside Echelon 1	
Alignment and Calibration	50%
First Full Reconstruction	40%
Reprocessing	60%
Simulation	75%

Distributed computing provides **flexibility**, influencing the fraction of use cases handled outside Echelon 1.

ECSJI Role and Echelon 1 Status

The EIC Computing and Software Joint Institute (ECSJI), established by BNL and JLab, is the host laboratories' joint organization for EIC computing and software. It:

- provides a single interface to ePIC and the EIC Project, as well as DOE and international computing partners;
- coordinates responsibilities, resources, services, and personnel at the two host laboratories;
- manages agreements, service expectations, reviews, and resource planning;
- supports the development and operation of the distributed EIC computing infrastructure.


Alexei Klimentov (BNL) and David Lawrence (JLab) are the ECSJI directors.

Echelon 1 comprises the computing facilities operated by BNL and Jlab:

- Early deployment should begin in 2027.
- Initial resources will support detector studies, simulations, streaming tests, and joint Echelon 0 to Echelon 1 data challenges.

EIC International Computing Organization

Discussion Focus

- **International partner engagement**:

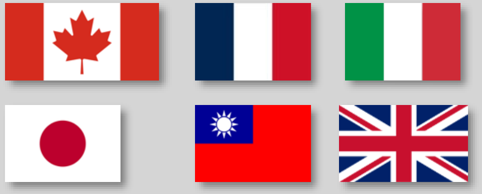

- **ePIC Streaming Computing Model**:
    - Coordination of data challenges with compute and data centers.
    - Echelon 2 roles and participation.
    - **Networking**: ESnet and LHCONE
- **WLCG**: Technical exchange and associated membership path.

ECSJI and EICO Annual Meeting (May 18–19, 2026, CERN)

ePIC is working closely with ECSJI and EICO to define **international participation** in the ePIC Streaming Computing Model and **advance its implementation**.

Distributed Computing and Storage: Status

Monthly simulation campaigns for preTDR and CD-2 preparations:

Compute Capacity

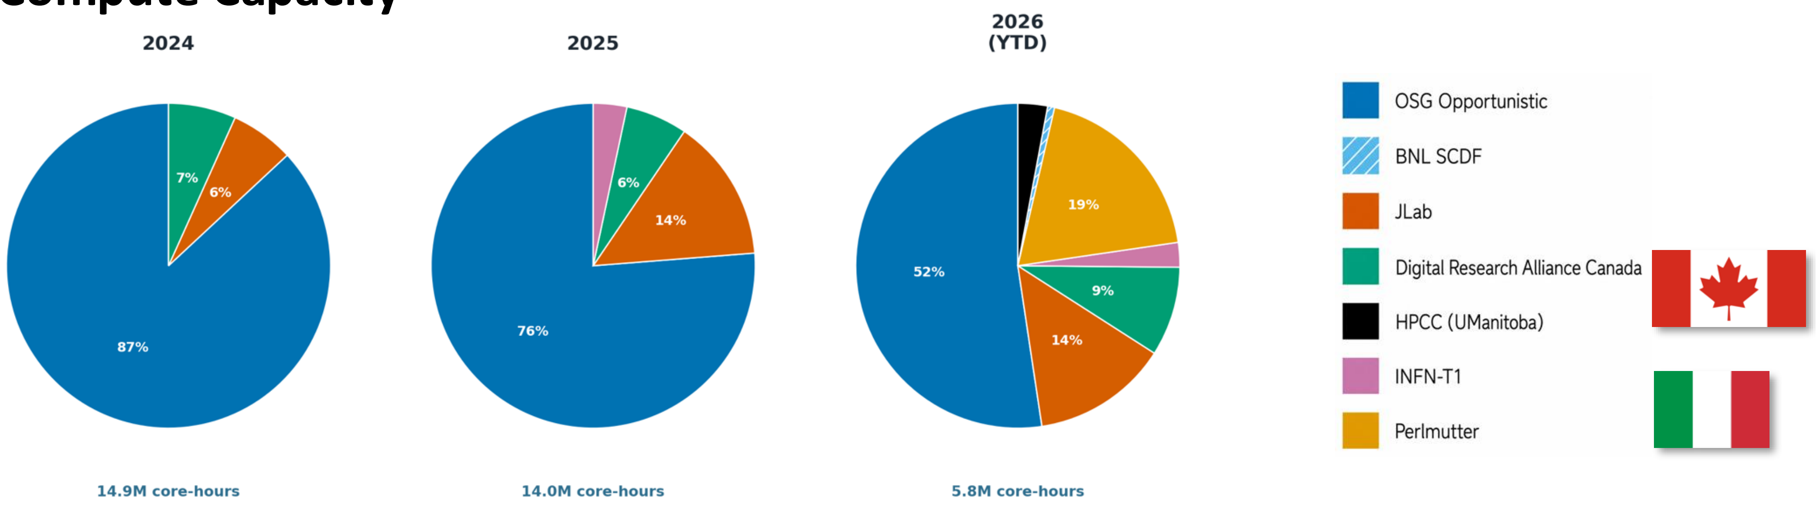

Storage Capacity

500 TB

500 TB

750 TB BNL

500 TB JLab

NEW + 500 TB **Academia Sinica Grid Computing Centre (ASGC)**

ePIC is building its distributed compute and storage fabric now, with ongoing discussions about **additional international contributions**.

ePIC Software Infrastructure and Collaborative Tools

Software Infrastructure

eic-web
Dedicated infrastructure at ANL for CI/CD. ePIC software development and releases depend on it.

MCP Server
LLM support for automated production workflows and AI assistants.

Collaborative Tools

GitHub: Collaborative software development.

eic.github.io: Documentation and tutorials.

Mattermost: Communicaton and HelpDesk support.

Zenodo.org: Document management

Roles

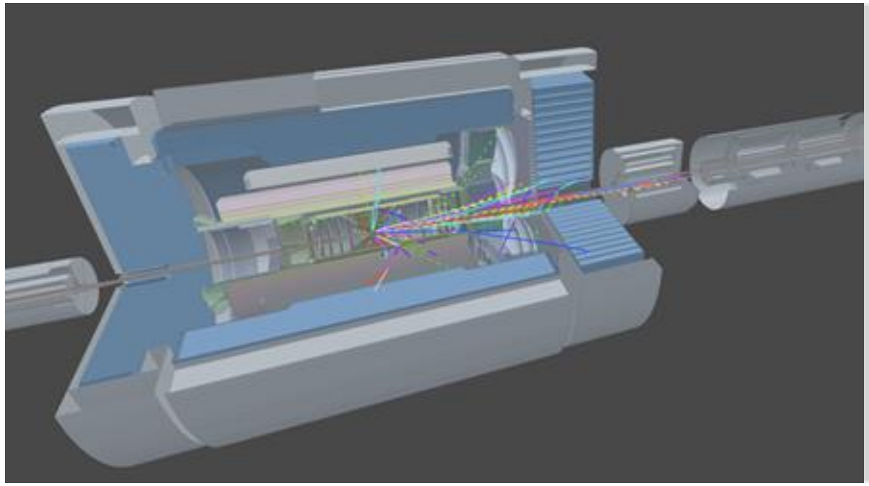

Enable Science

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

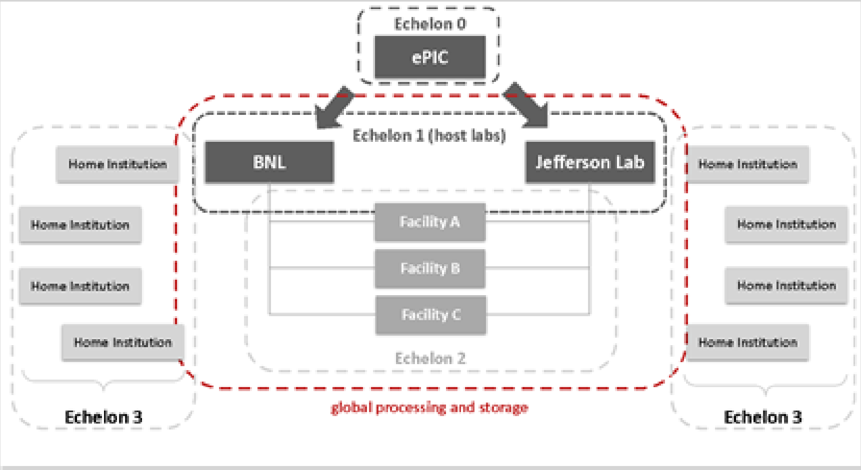

Compute–Detector Integration

- 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

Simulation Requests

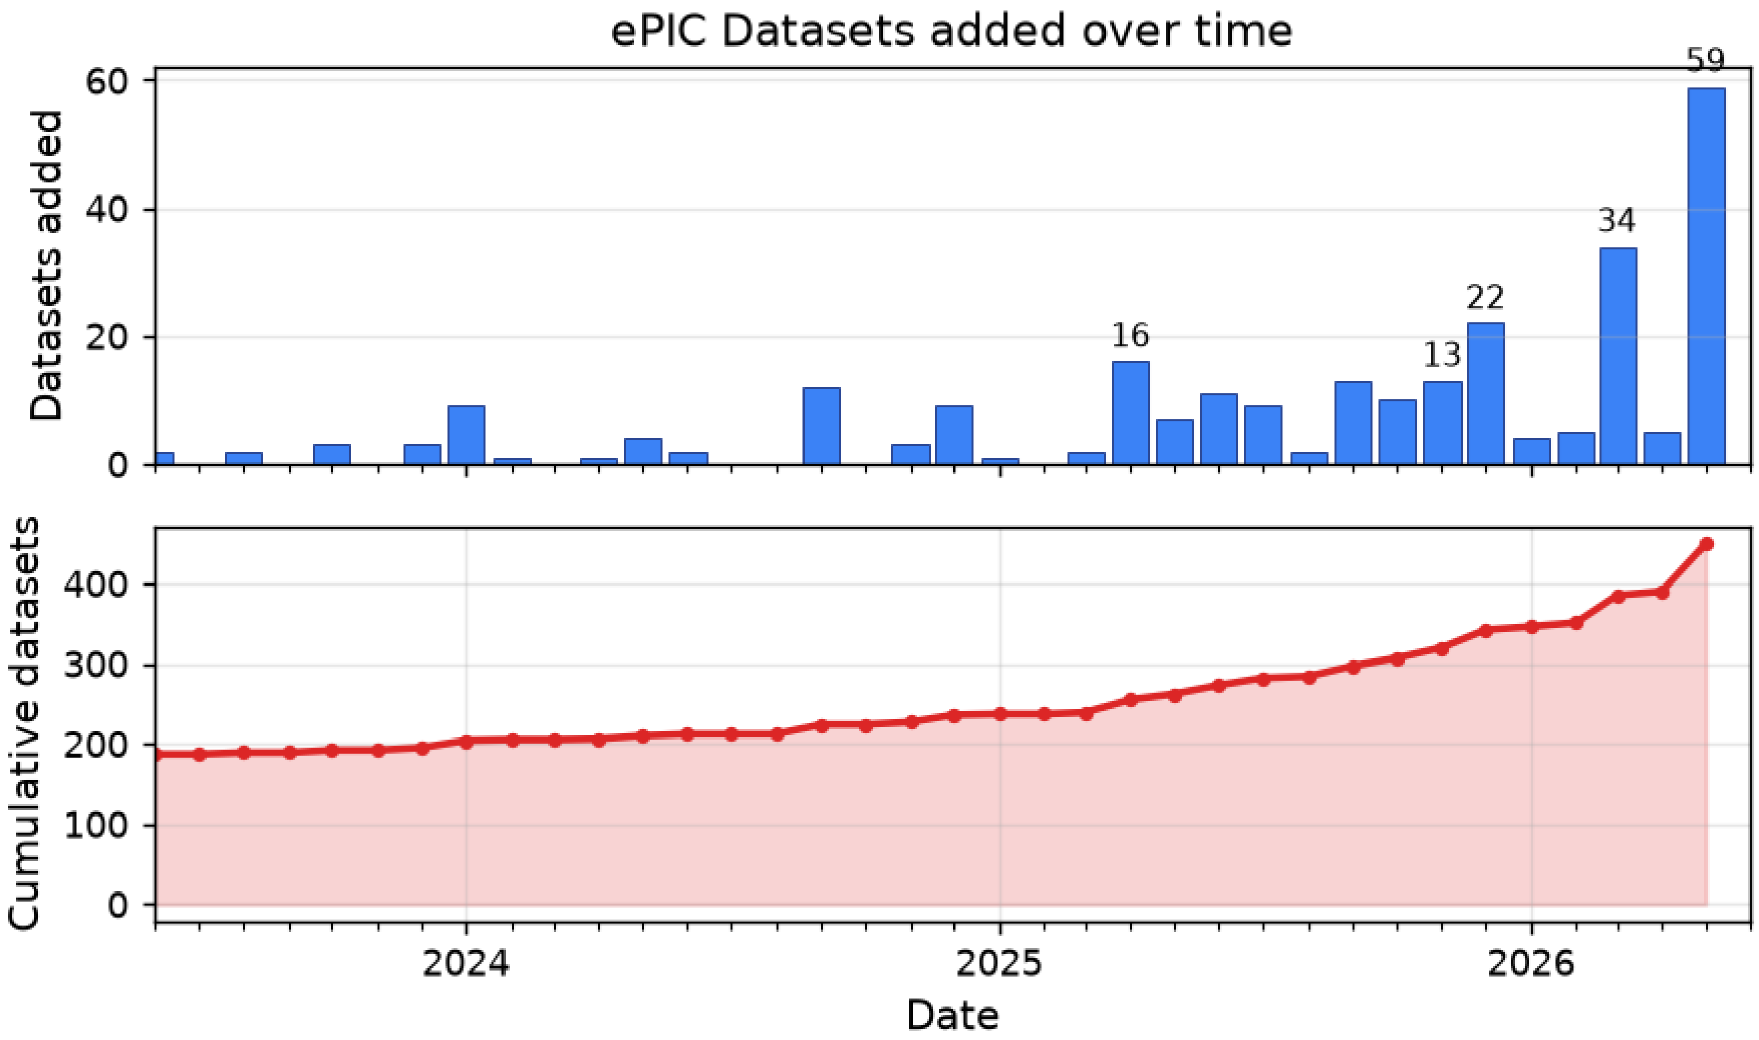

July Campaign Forecast

Approximately 350 production requests

Expected range: 342–364

epicprod: ePIC Automated Production Workflow

- Growing simulation demand made **automation a priority**.
- After evaluating options, we **selected PanDA**.
- In parallel, **PanDA-based agentic workflows** showed promising results in the BNL testbed.
- The charge: **build epicprod** on the PanDA agent-based streaming workflow testbed.

More details: [ePIC Software & Computing presentation]

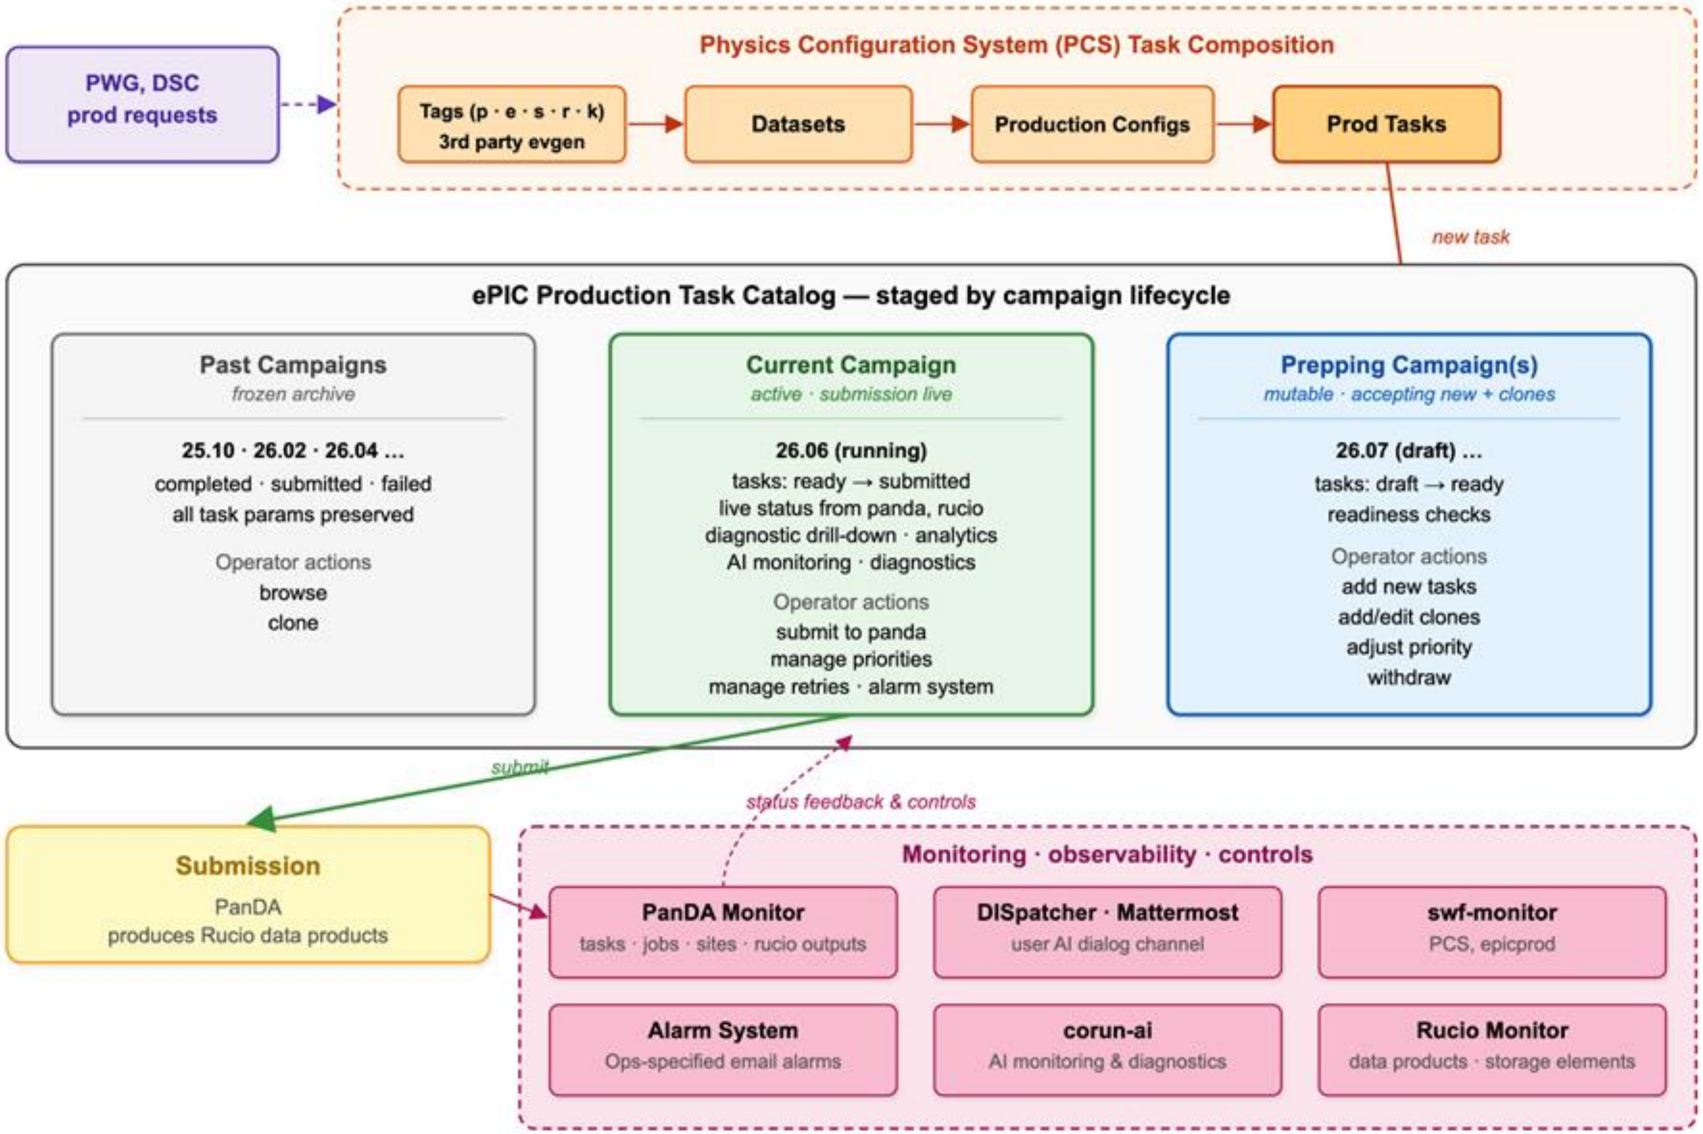

PanDA: Proven, Scalable, and Automation-Ready

PanDA is a mature, scalable platform for automating distributed computing.

- **Proven at scale:** 21 years of stable ATLAS operations, supporting 600k+ concurrent cores and about 1 million jobs per day. 1 Extabyte of data processed.
- **Automation across resources:** grid, cloud, HPC, and local clusters
- **Integrated operations:** workload orchestration, data management, monitoring, and error recovery
- **Built to evolve:** supports new workflows and AI-assisted operations

More details: PanDA Documentation

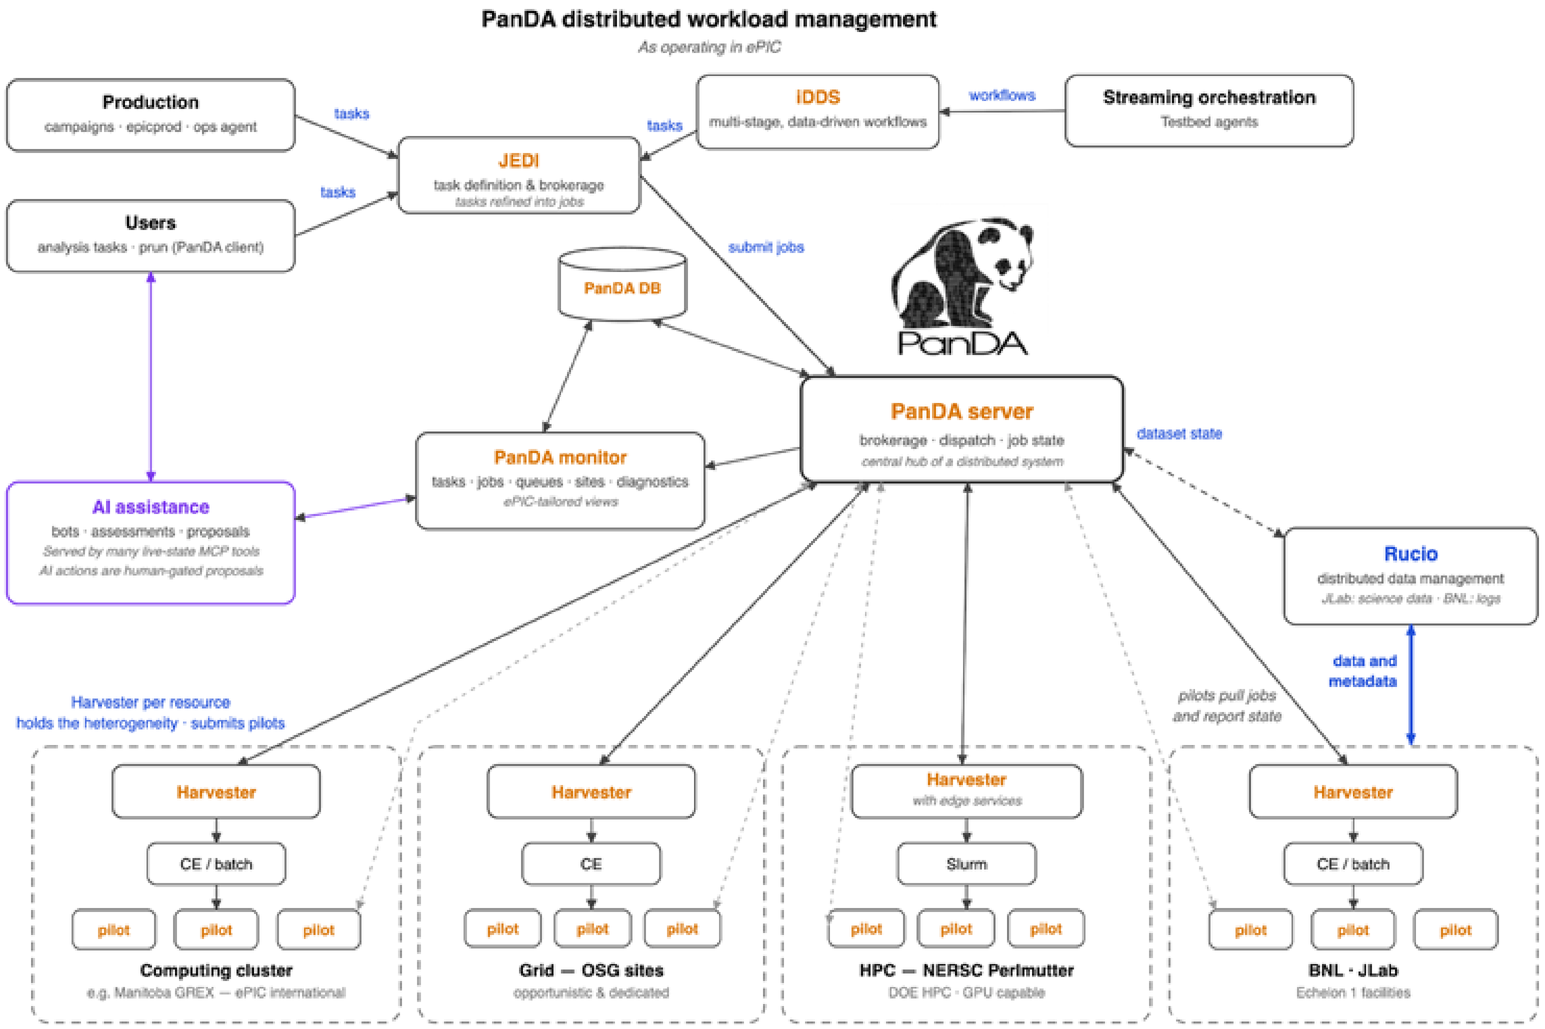

End-to-End Production Automation

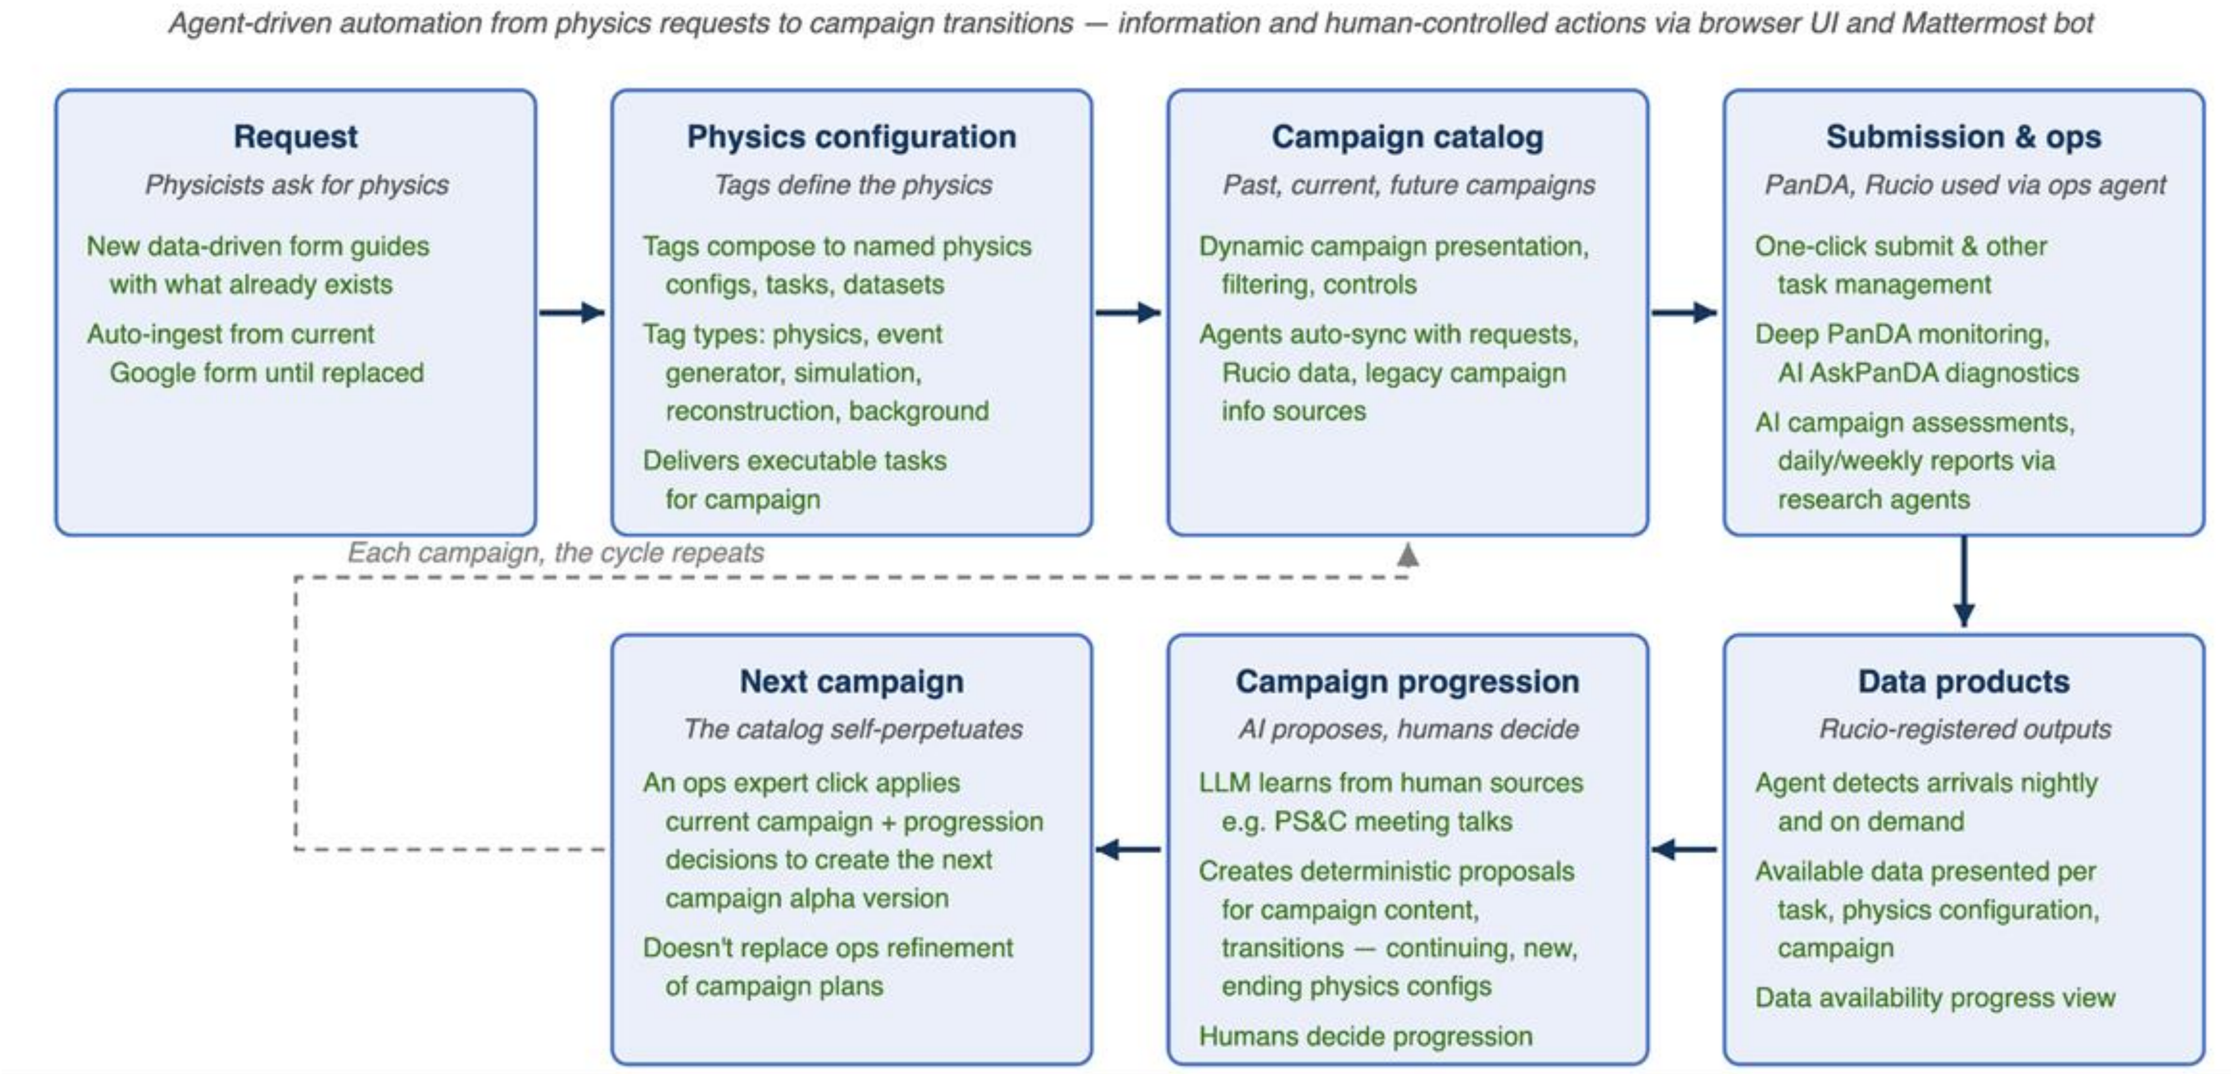

Explore epicprod In Your Browser

https://epic-devcloud.org/prod/

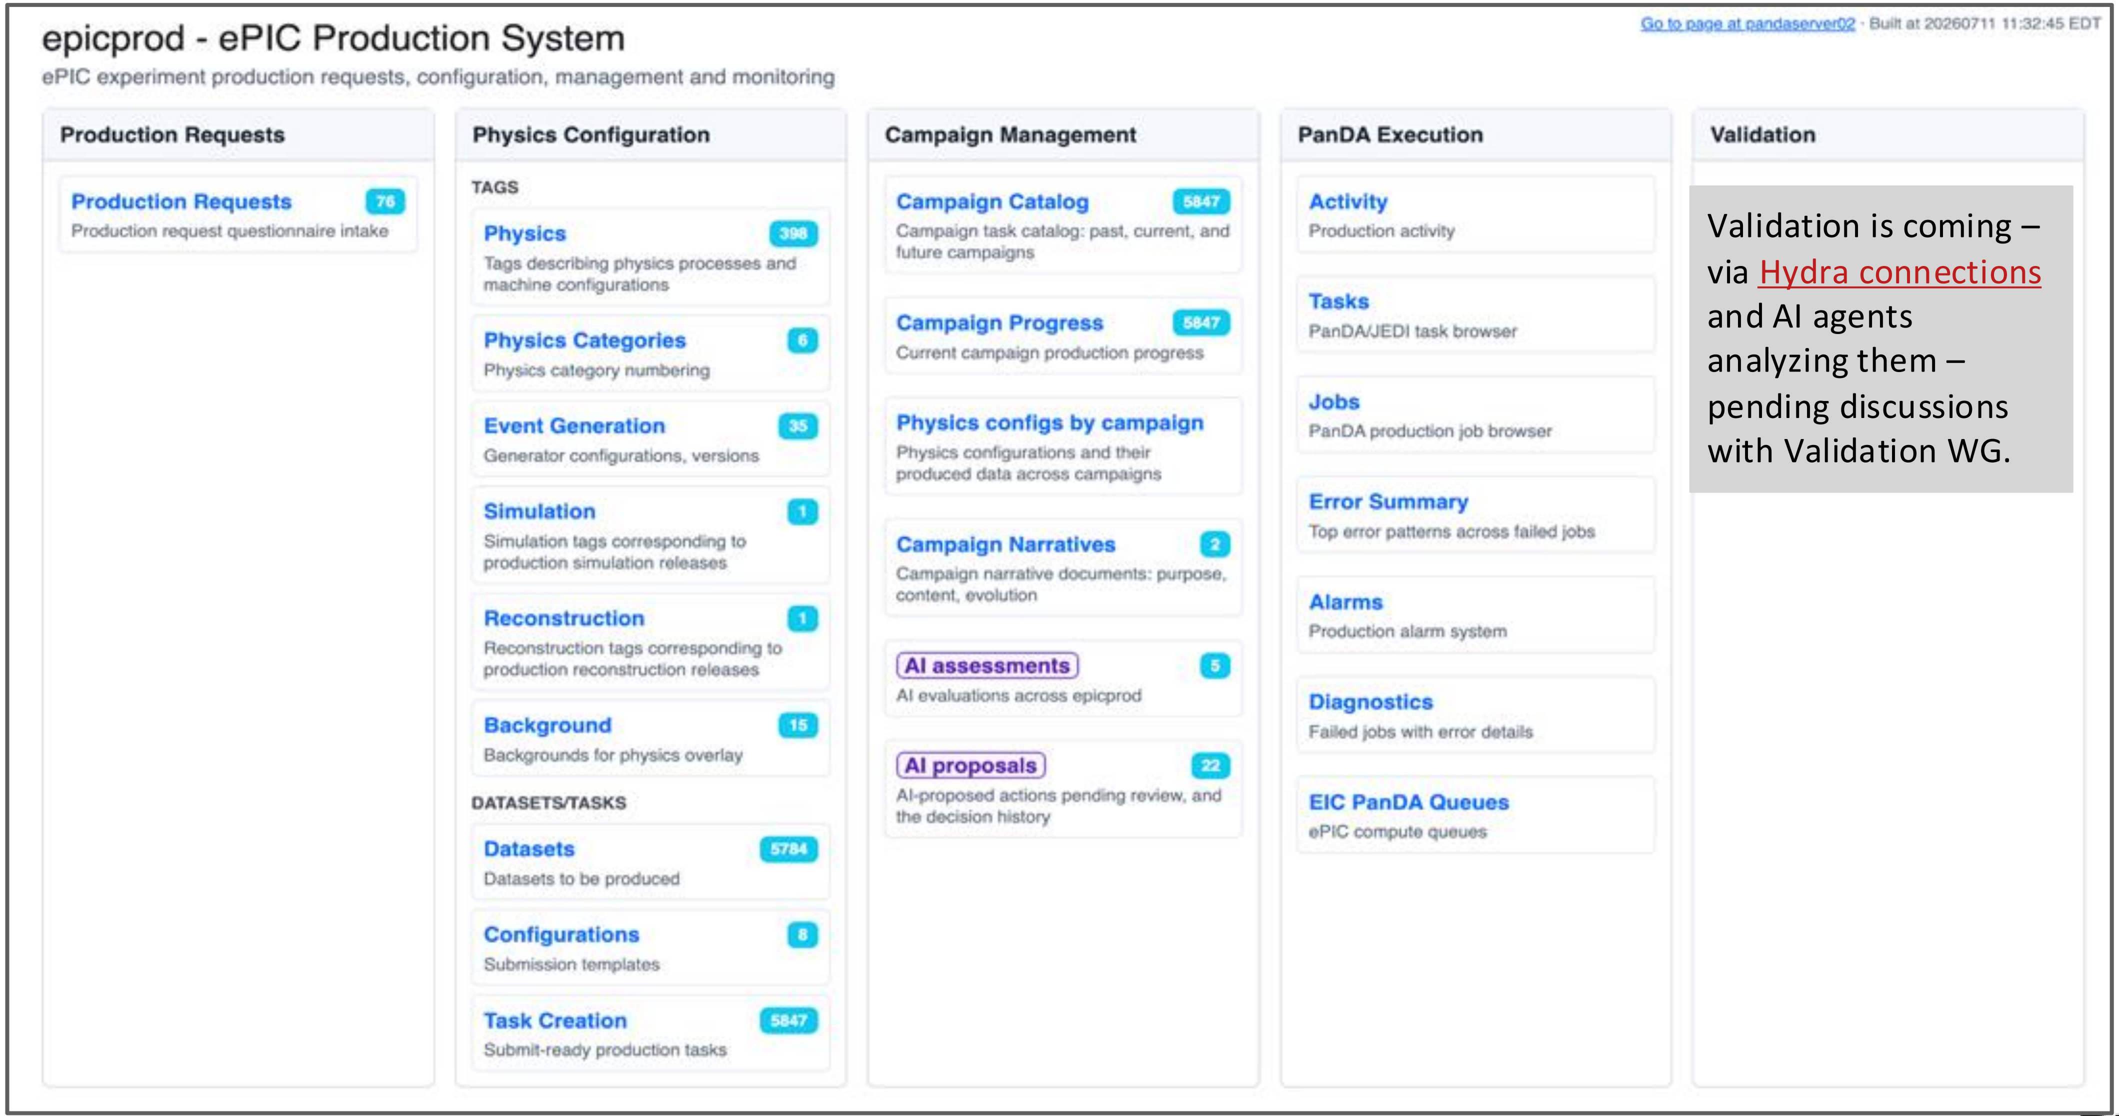

Validation is coming – via Hydra connections and AI agents analyzing them – pending discussions with Validation WG.

Validation via **HYDRA**

https://epsciweb.jlab.org/epic_hydra/allPlotsBrowser

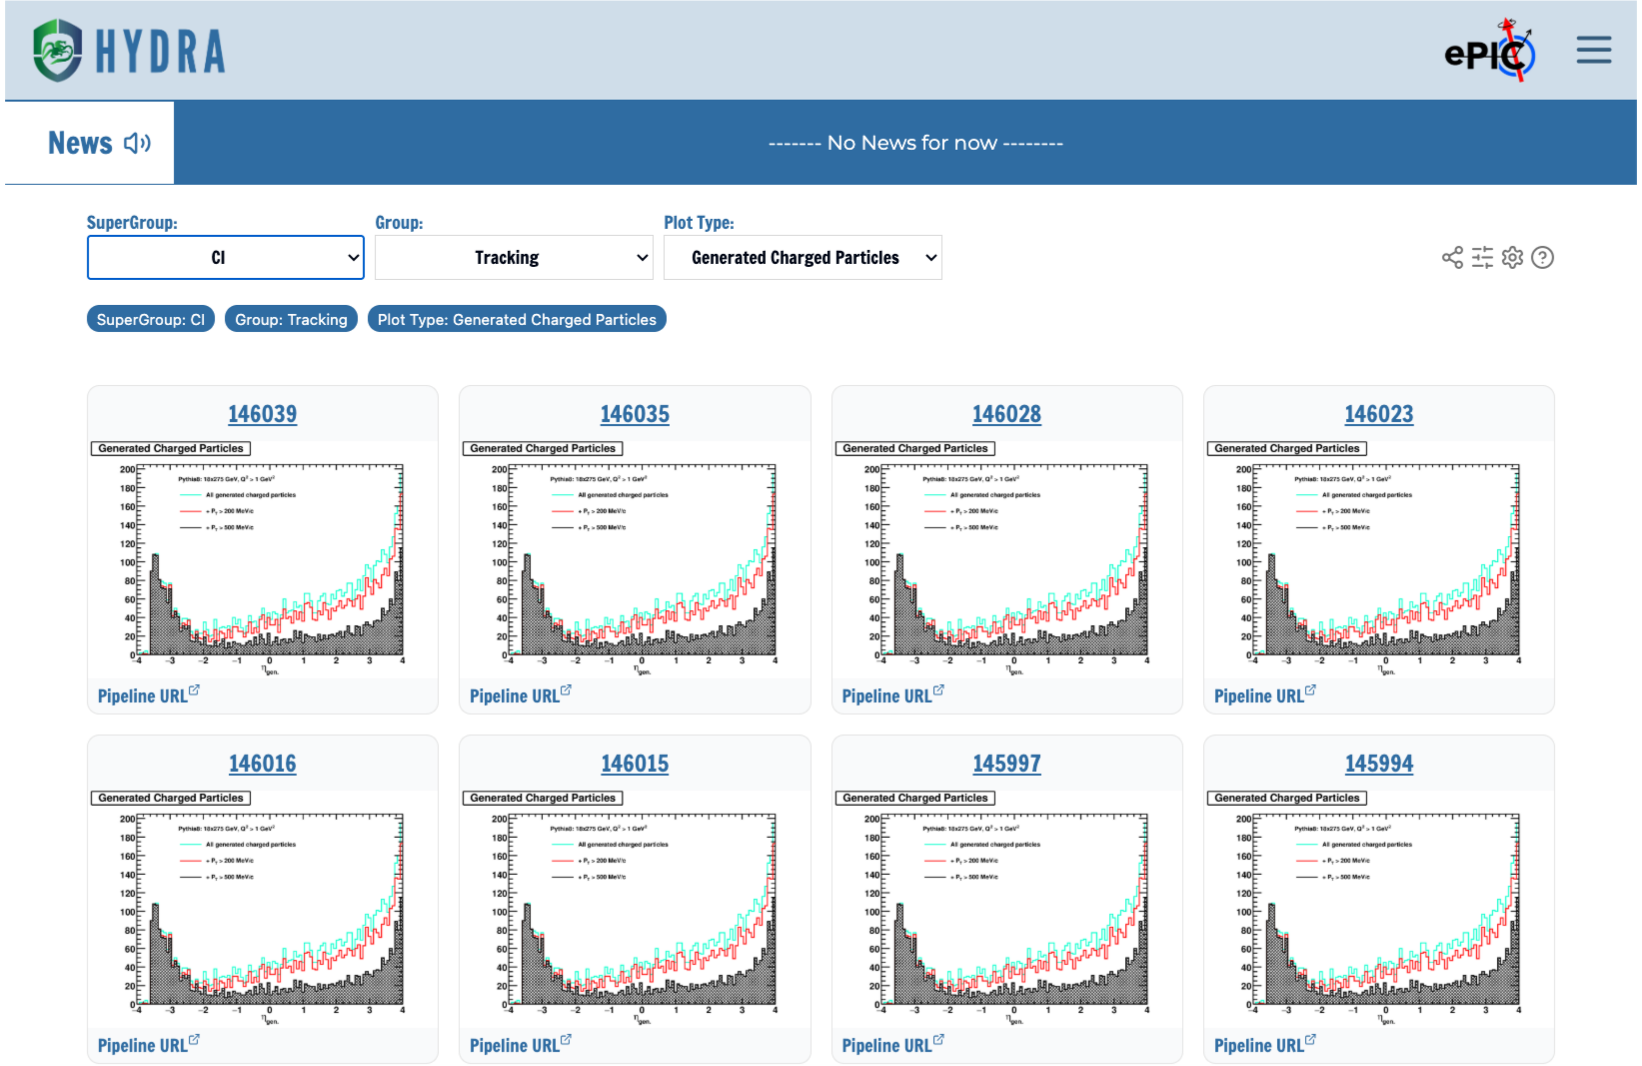

Tracking Physics Configurations Across Campaigns

Go to page at pandaserver02 · Built at 20260711 12:04:19 EDT

Physics Configurations

Physics first: one row per physics configuration — the campaign-invariant identity a request asks for — with its editions along the campaign axis, showing fulfillment through time. A presentation view: nothing is edited here; edition links lead to the catalog, where actions live. See docs/CAMPAIGN_CONTINUUM.md.

Campaigns: 2026 2025 all

Process: All ALP (44) BACKGROUND (18) DDIS (1) DDVCS (6) DEMP (6) DIFFRACTIVE_JPSI (5) DIFFRACTIVE_PHI (2) DIFFRACTIVE_RHO (1) DIS (24) DIS_CC (27) DIS_NC (97) DVCS (5) DVMP (6) MESON_SF (4) PHOTOPRODUCTION_JPSI (3) SIDIS (26) SIDIS_D0 (8) SIDIS_DIJET (1) SIDIS_Lc (3) SINGLE (253) UPSILON (13)

Generator: All BeAGLE (28) DEMPgen (6) DJANGOH (8) EIC_ESR_Xsuite (9) EIC_SR_Geant4 (2) EpIC (14) GETaLM (3) eSTARlight (3) eicMesonSFGen (4) lAger (8) particle_gun (253) pythia6 (22) pythia8 (34) rapgap (1) sartre (2)

Beam: All 10x100 (49) 10x110 (47) 10x115 (4) 10x130 (19) 10x166 (3) 10x250 (14) 10x275 (23) 18x110 (3) 18x275 (63) 5x100 (22) 5x41 (26) 9x130 (3) 9x275 (3)

Species: All eAu (10) eCu (2) eH2 (4) eHe3 (12) eRu (2) ep (74)

Q²: All minQ2=1 (15) minQ2=10 (15) minQ2=100 (30) minQ2=1000 (24) q2_0_10 (6) q2_0to1 (10) q2_1 (4) q2_100 (1) q2_1000to10000 (4) q2_1000to100000 (4) q2_100to1000 (16) q2_100to10000 (4) q2_10_20 (2) q2_10to100 (22) q2_1_100 (5) q2_1_1000 (3) q2_1to10 (18) q2_1to1000 (7) q2_1to10000 (6) q2_1to20 (1) q2_1to20000 (1) q2_1to50 (3) q2_20_35 (2) q2_2to10 (4) q2_3_10 (2) q2_nocut (2)

Sample: All 130to177deg (80) 3to50deg (87) 45to135deg (80) etaScan (6)

Request: All matched (124) unmatched (429)

Has data: All 26.06.0 (116) 26.05.0 (154) 26.04.1 (200) 26.03.1 (79) 26.03.0 (126) 26.02.0 (144)

Search: physics / generator / edition name... Filter Clear

553 of 553 configurations

Physics & #requests	Process	Beam	Q²	Generator	Sample	26.06.0	26.05.0	26.04.1	26.03.1	26.03.0	26.02.0
					produced, over the rows below:	116443 files 49.7 TB 437 datasets / 116 of 415 configs	529331 files 172.8 TB 677 datasets / 154 of 447 configs	400590 files 141.6 TB 381 datasets / 200 of 200 configs	184463 files 27.3 TB 79 datasets / 79 of 79 configs	415211 files 80.9 TB 310 datasets / 126 of 126 configs	256972 files 64.0 TB 323 datasets / 144 of 144 configs
p1015 req#406 #191	DDIS	10x100		rapgap 3.310-1.0 ep noRad		○ task	○ task	✓ 4370 f 1.4 TB data · task	✓ 4370 f 735.7 GB data · task	·	✓ 3505 f 740.8 GB data · task
p1057 req#561	SINGLE			particle_gun	130to177deg	○ task	○ task	·	·	·	·
p1057 req#562	SINGLE			particle_gun	3to50deg	✓ 17 f 8.6 GB data · task	✓ 59 f 29.9 GB data · task	✓ 68 f 15.8 GB data · task	·	✓ 54 f 12.8 GB data · task	✓ 63 f 13.4 GB data · task
p1057 req#563	SINGLE			particle_gun	45to135deg	○ task	○ task	·	·	·	·
p1058 req#564	SINGLE			particle_gun	130to177deg	○ task	○ task	·	·	·	·
p1058 req#565	SINGLE			particle_gun	3to50deg	✓ 202 f 213.4 GB data · task	✓ 695 f 771.3 GB data · task	✓ 730 f 667.5 GB data · task	·	✓ 764 f 242.6 GB data · task	✓ 1025 f 245.6 GB data · task
p1058 req#566	SINGLE			particle_gun	45to135deg	○ task	○ task	·	·	·	·
p1059 req#567	SINGLE									·	·
p1059 req#568	SINGLE									✓ 164 f 66.7 GB data · task	✓ 136 f 67.2 GB data · task
p1059 req#569	SINGLE									·	·
p1060 req#570	SINGLE									·	·
p1060 req#571	SINGLE									✓ 51 f 17.7 GB data · task	✓ 105 f 18.3 GB data · task

AI Proposals

AI-proposed actions pending human review, and the decision history. Approval executes the frozen payload deterministically th

Status: All proposed (22) executed (1)

Action: All campaign propagation (25)

Change: All continue → final (16) continue → hold (9)

Decision: All wenauseic (3)

Proposer: All claude-fable-5 (1) fable (22) verify-undo (2)

Batch: All 26.07.0-dispositions-20260709 (22) verify-ai-app (1) verify-undo-20260708 (2)

0 selected Approve Deny quality... Quality choice is recorded with the decision. Undo Delete

Ref	Proposed	Action	Change	Subject	Comment	Status	Decision
cp-26	2026-07-09 06:10	campaign propagation	continue → hold	group.EIC.26.02.0.epic_craterlake.p2224.e38.s1.r1	Legacy DJANGOH NC DIS 18x275 is reference and comparison only, not part of 26.07.0 production, per S. Rahman, "Campaign 26.07 Preparations", PSC meeting 2026-07-08 (indico [illegible])	proposed	—
cp-25	2026-07-09 06:10	campaign propagation	continue → hold	group.EIC.26.02.0.epic_craterlake.p2224.e37.s1.r1	Legacy DJANGOH NC DIS 18x275 is reference and comparison only, not part of 26.07.0 production, per S. Rahman, "Campaign 26.07 Preparations", PSC meeting 2026-07-08 (indico [illegible])	proposed	—
cp-24	2026-07-09 06:10	campaign propagation	continue → hold	group.EIC.26.02.0.epic_craterlake.p2223.e38.s1.r1	Legacy DJANGOH NC DIS 18x275 is reference and comparison only, not part of 26.07.0 production, per S. Rahman, "Campaign 26.07 Preparations", PSC meeting 2026-07-08 (indico [illegible])	proposed	—
cp-23	2026-07-09 06:10	campaign propagation	continue → hold	group.EIC.26.02.0.epic_craterlake.p2223.e37.s1.r1	Legacy DJANGOH NC DIS 18x275 is reference and comparison only, not part of 26.07.0 production, per S. Rahman, "Campaign 26.07 Preparations", PSC meeting 2026-07-08 (indico [illegible])	proposed	—

AI proposes campaign to campaign continuation (or not) based on analyzing, e.g., Physics, Software, and Computing Discussion talks (These are product of July 8 meeting)

ePIC

From Requests to Data

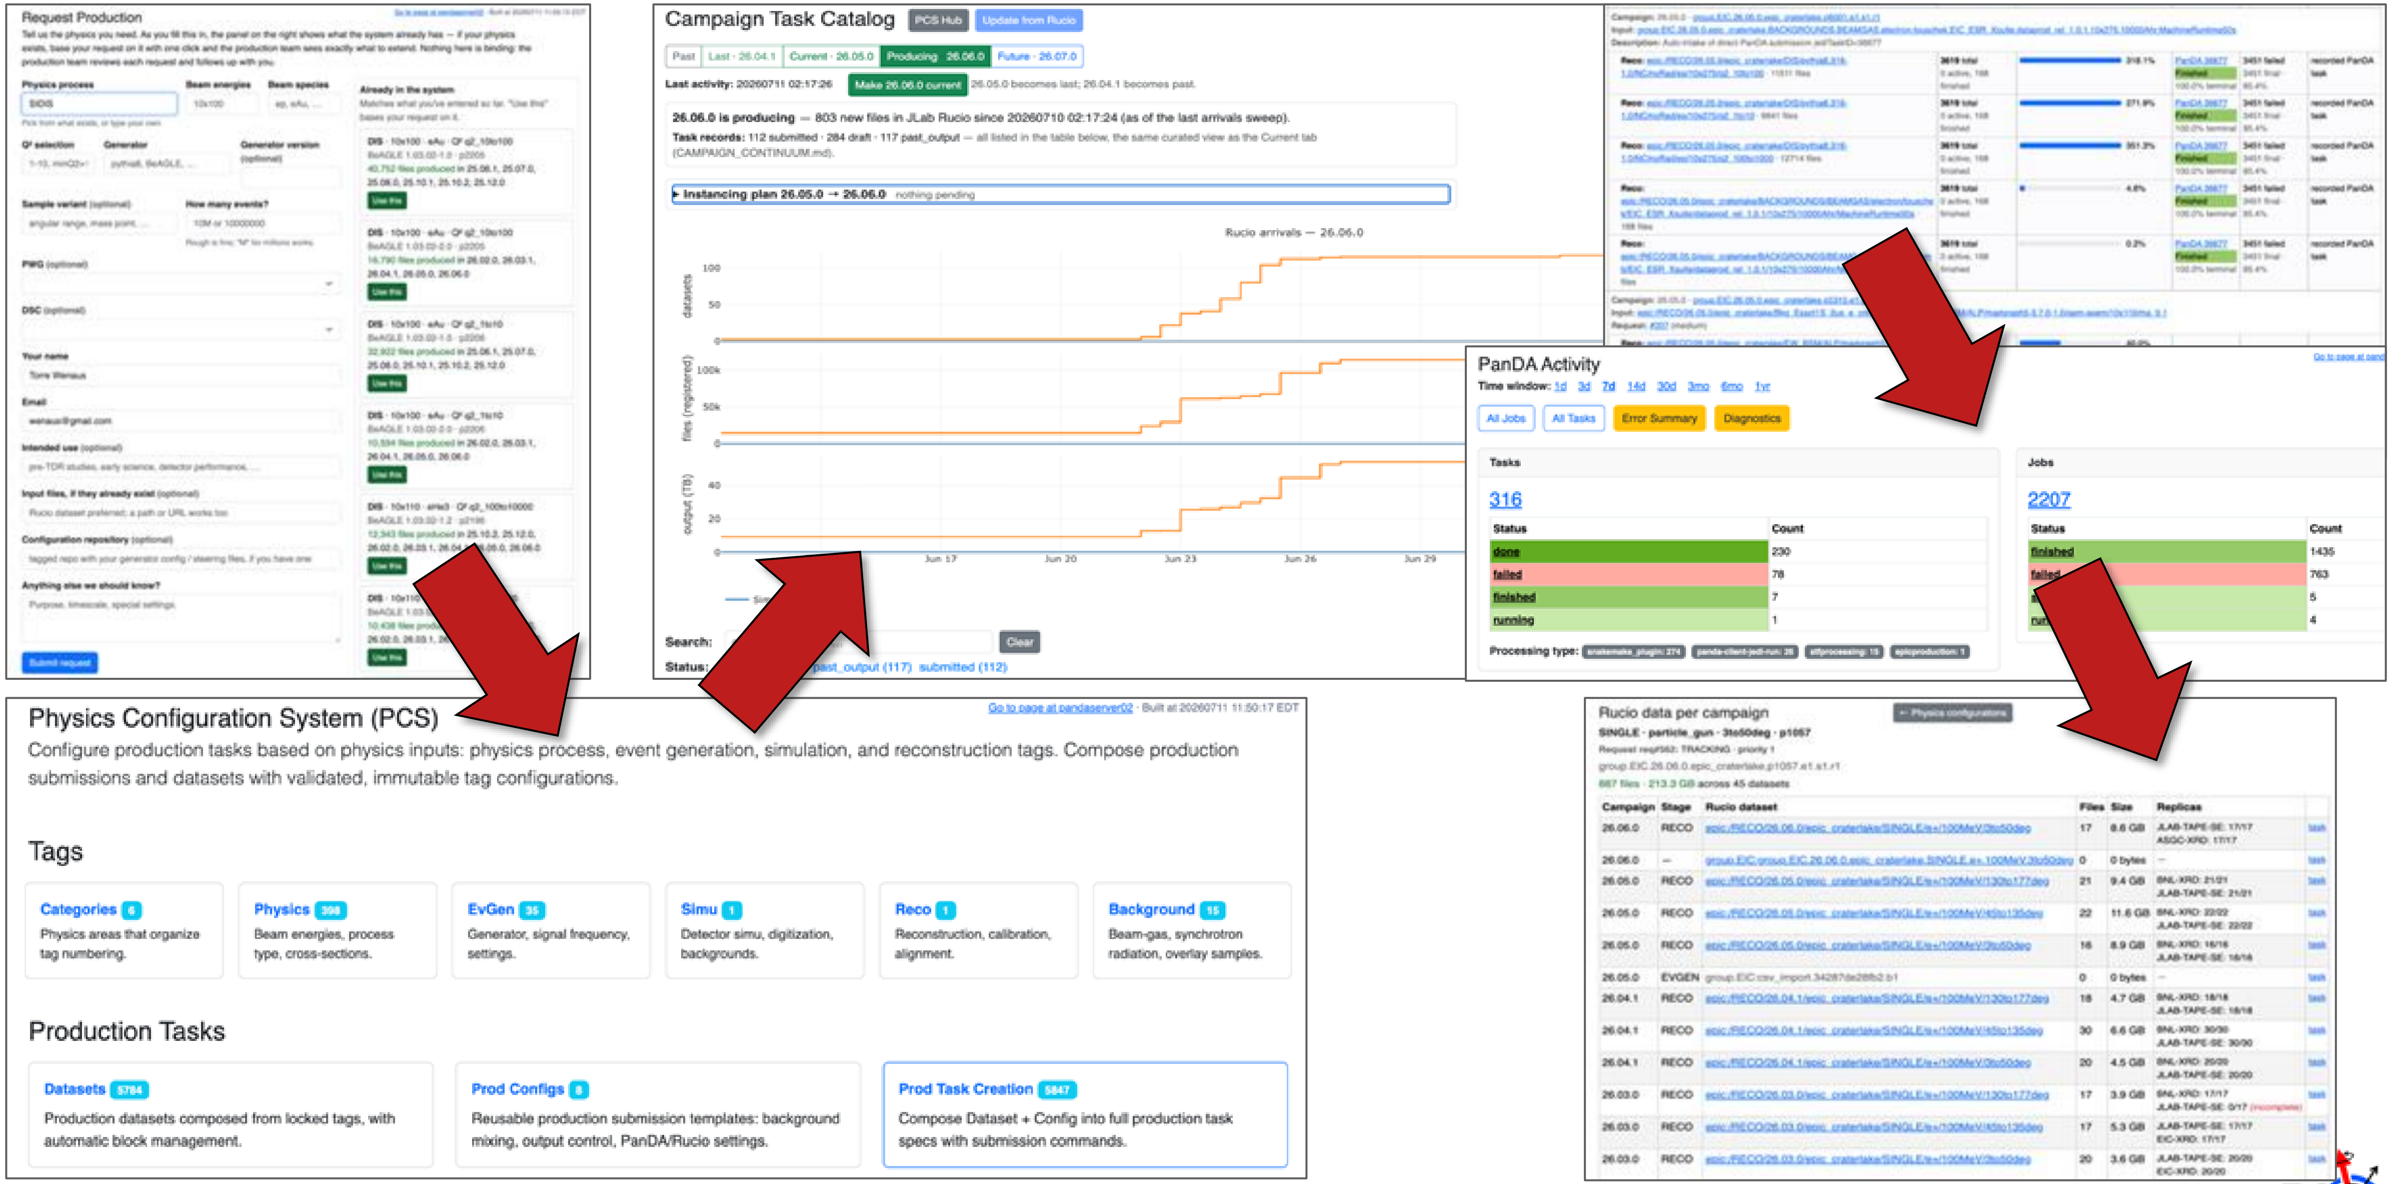

preTDR Priorities and Progress for 2026

At our annual collaboration meetings at ANL in 2024, in Frascati in 2025, and at BNL in 2026, we defined the priorities for each year. At the 2026 meeting, we agreed on the following priorities.

Production Provide the simulation campaigns with background to finalize the ePIC detector design and validate its physics performance ✅, roll out Rucio as the default system for finding and accessing simulation data ✅, and automate production workflows ✅.	**Slides 17–24**
Reconstruction Coordinate the effort to address gaps in reconstruction ✅. Work toward a holistic reconstruction approach, such as particle identification that integrates information from calorimeters, Cherenkov detectors, and time-of-flight systems.	**Slides 26–28**
Simulation Implement and operate a workflow between detector and simulation experts to track the status of the comparison between the simulation design and the engineering design, and to resolve any discrepancies in a timely and systematic manner ✅.	**Slide 29**
Streaming Orchestration Define and test the interface between DAQ and computing to mitigate risks in the integrated DAQ-computing system. As part of the integration, test develop and deliver a functional testbed that validates a workflow management system for the autonomous calibration of a detector subsystem.	**Slides 4–11**
Cross-cutting priority: Workforce Foster a developer community within ePIC and support the careers of scientists focused on software & computing.	**Slides 30–32**

Reconstruction: Coordination with PWGs

Current Shared Priorities with PWGs

Priority	Contacts	Description
Hadron ID	Chandra	PID algorithms for HF physics and SIDIS studies.
Particle Flow	Derek	Improving the jet reconstruction using PF information.
Sec. Vertexing	Xin, Shaym, Rongrong	Enhancing the vertexing capabilities to study HF physics.
Event Kinematics	Charlotte, Rachel	Establishing an initial dictionary of kinematics for physics process and supporting physics analyses.

Current Reco WG Priorities

Priority	Contacts	Description
Tracking	Shujie	Enhancing tracking capabilities, supporting detector studies
Lepton ID	Win, Stephen M, Stephen K	Improving and extending lepton ID capabilities
Jet/Calo Dev.	Derek	Enhancing general jet and calorimeter reconstruction

Reconstruction Progress

Tracking

- fTOF/bTOF enabled with charge sharing
- Migration to Acts Seeding2 underway

Secondary Vertexing

- AMVF secondary vertexing merged in addition to Helix method
- A new dedicated SecondaryVertex data structure to be added to EDM

Far-detectors

- Lambda reconstruction combining ZDC, LFHCAL, fEMCal, forward insert

Event Kinematics

- Work started on utility to create list of subsystems used by a reconstructed particle (see 06.08.2026 Reco WG meeting)
- Ongoing discussions on best approach for merging FF and central reconstructed particles

Particle ID

- **Lepton ID**: See last update at 05.27.2026 PAC/S&C Discussion
- Excellent progress on IRT2 integration since January Collaborative efforts between groups at BNL and at INFN
- hpDIRC showing results obtained using official ePIC simulation stack, integration in EICrecon still TODO

Jet/Calorimeter Development

- EICrecon jet reconstruction now outputs edm4eic::Jet objects (PR#2674) as of v1.38.0
- A cool optimization of the Centauro algorithm! Contributed in the latest FastJet Contrib

Particle Flow

- **PFA1a** (track-cluster subtraction), **PFA0be** (track-protocluster match promotion) in EICrecon
- **PFA2** (calo.-remnant combination), **PFA3** (candidate particle conversion) targeting EICrecon **v1.40.0**
- PFA0, PFA1 tuning **in-progress**

Firebird Event Display

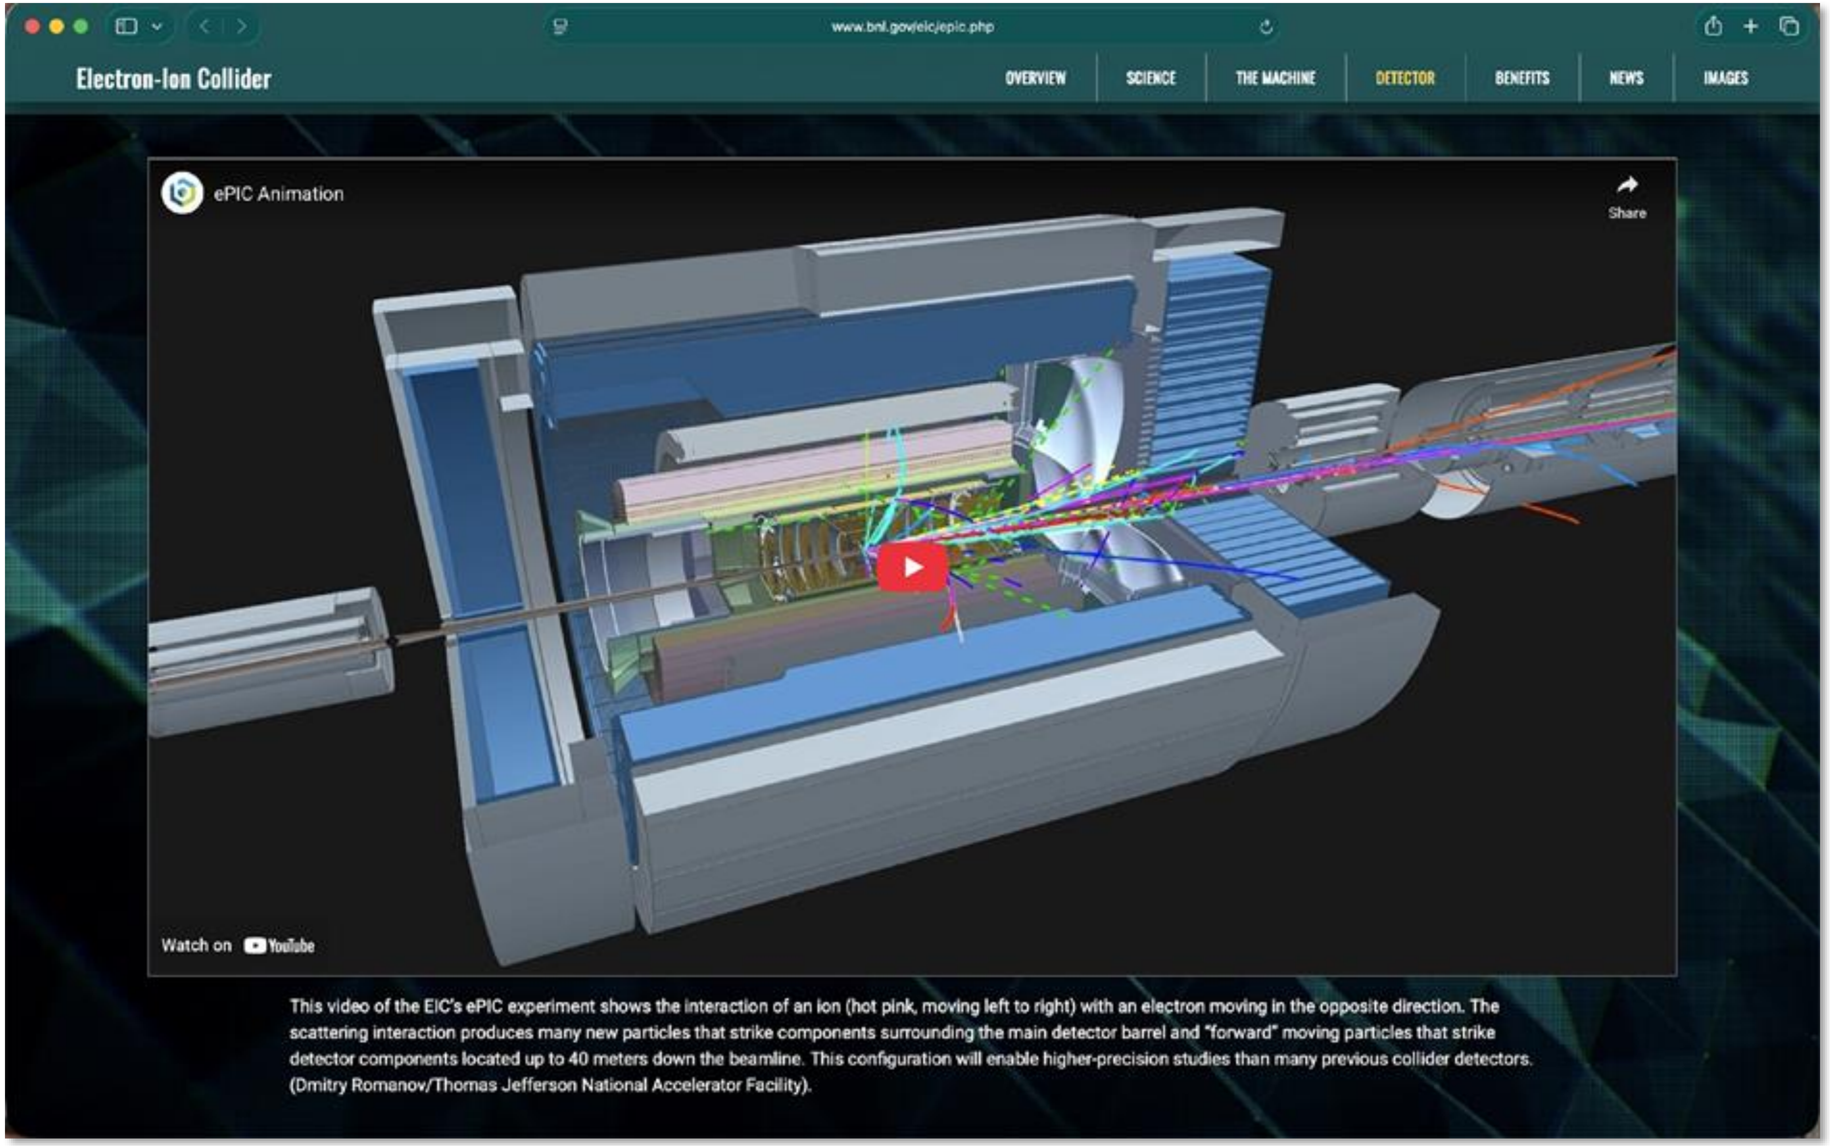

Firebird supports debugging, quality control, and outreach through intuitive visualizations of detector responses, particle interactions, and reconstruction outputs.

A recording of the ePIC event display is available on the EIC webpage.

Physics and Detector Simulation Progress

Simulation Geometry Validation

- **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

Simulation Speed Optimization

- Systematic studies for bottlenecks
- Conventional approaches (cuts, settings, geometry fixes)
- R&D project (G4HepEm, GPU acceleration)

Background Mixing in Simulation

- Work ongoing on background merging after Geant4 simulation step to improve performance
- Background event sampling was revised in July campaign to reduce cross-file correlations
- Synchrotron radiation samples to be updated after ESR lattice is available

Metadata Propagation

- Making metadata accessible for physics analysis:
  - **Event generator information.** Limitation: external generator code.
  - **Beamline parameters** Perfect chance to incorporate in the revised 9 GeV Monte Carlo samples
  - **Simulation configuration**

Software Contribution Guidelines

In response to demand for clear Contribution Guidelines we've formed ePIC Review Team with goal of reviewing existing workflows and formulating guidelines.

Full write up is in CONTRIBUTING.md. The PR templates were revised to reflect our approach (including AI use reporting).

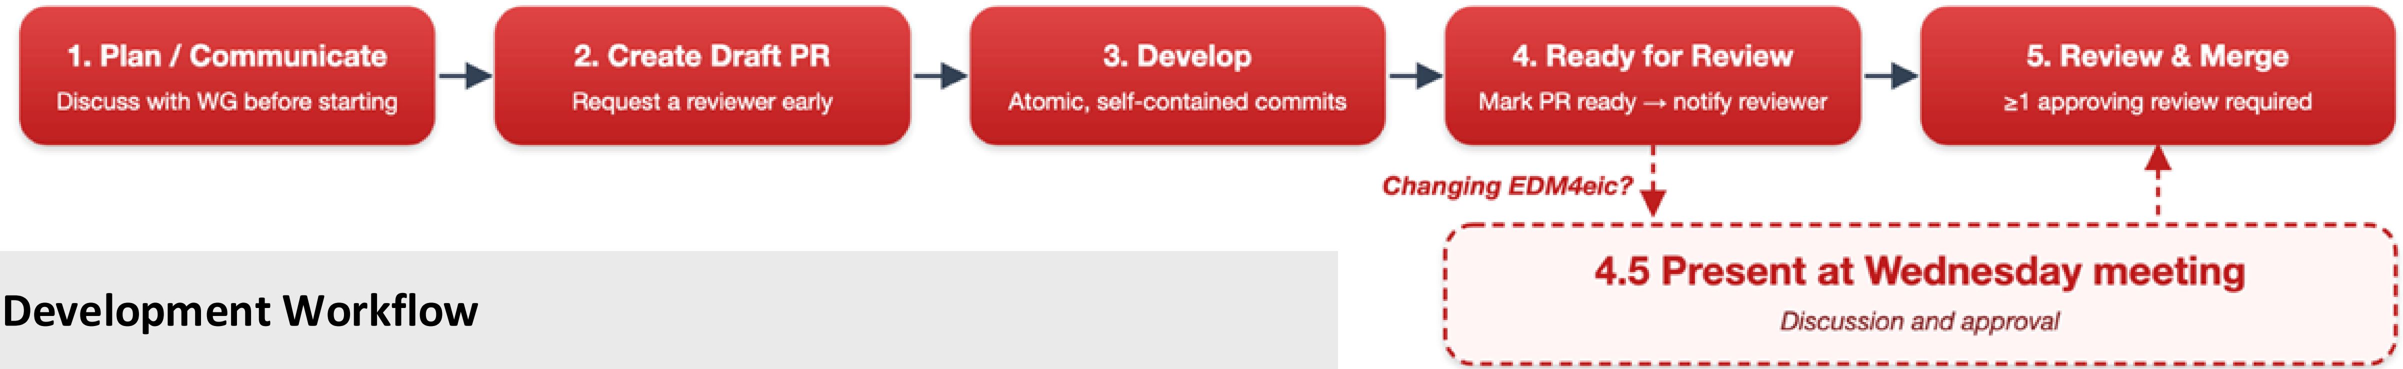

Development Workflow

- Start each development in a **new branch**
- Commit atomically and frequently — **one change per commit**
- Write **descriptive commit messages**
- Keep PRs **small and focused** on a single concept
- Avoid mixing logic changes with cleanup/typos
- Reviewer states a **timescale** for review promptly
- **Review Team** is the default reviewer

Questions? Contact Team at epic-software-review@lists.bnl.gov!

PR Checklist — Ready for Review

✓ Detailed PR description provided

✓ All CI jobs pass

✓ Legible, well-structured commit history

✓ Code is appropriately documented

Updated Landing Page

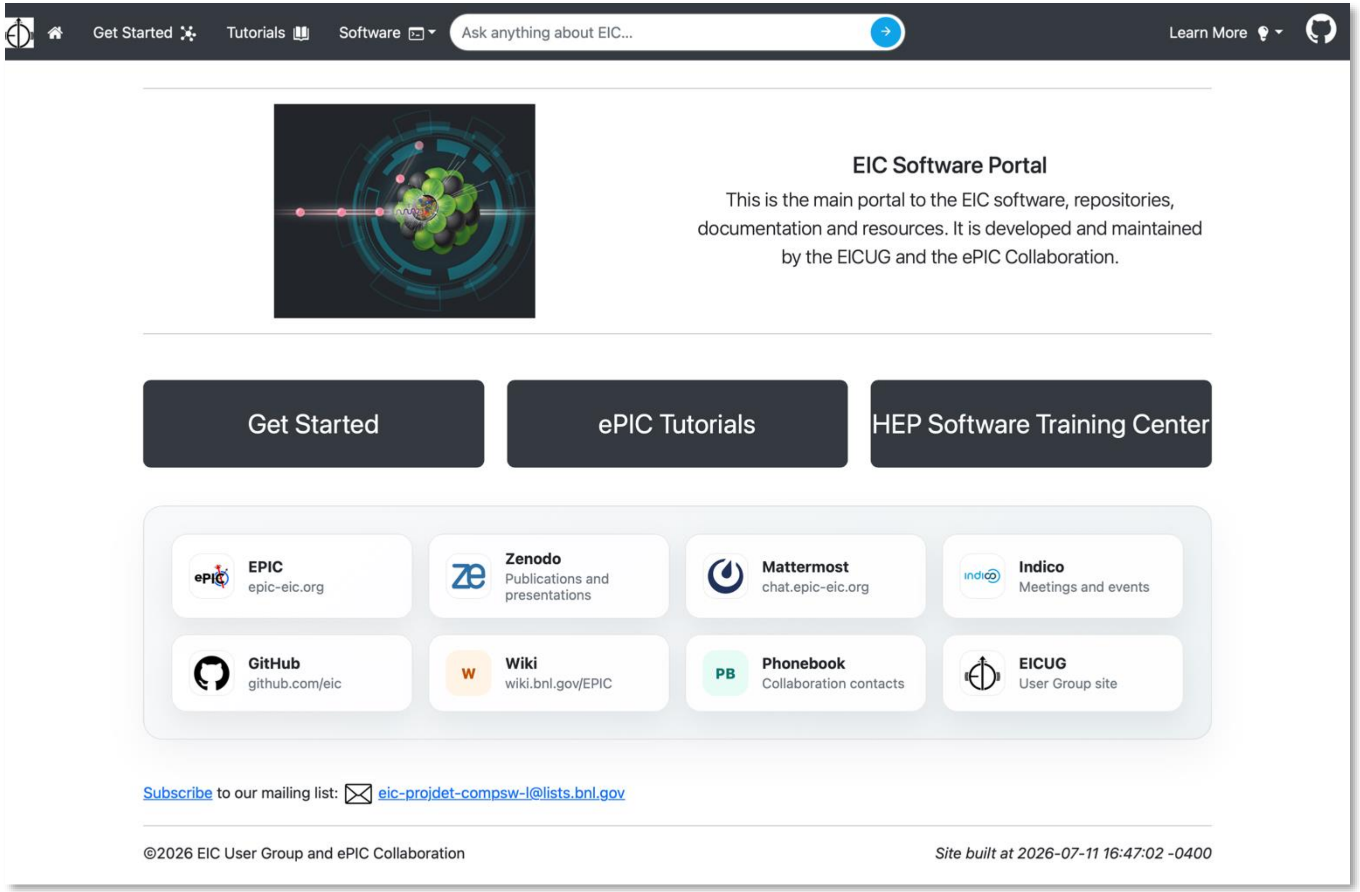

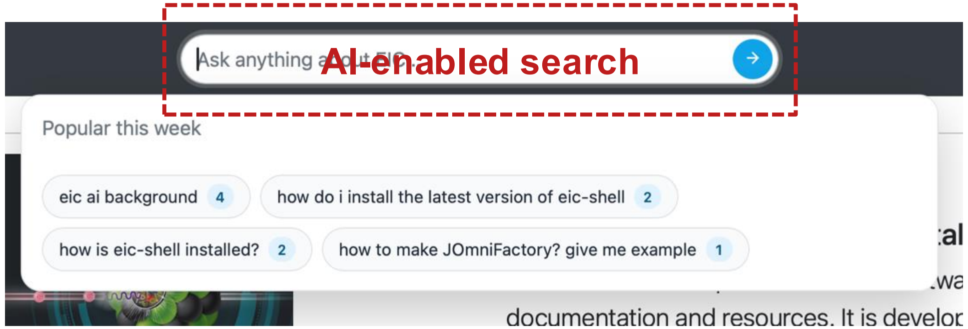

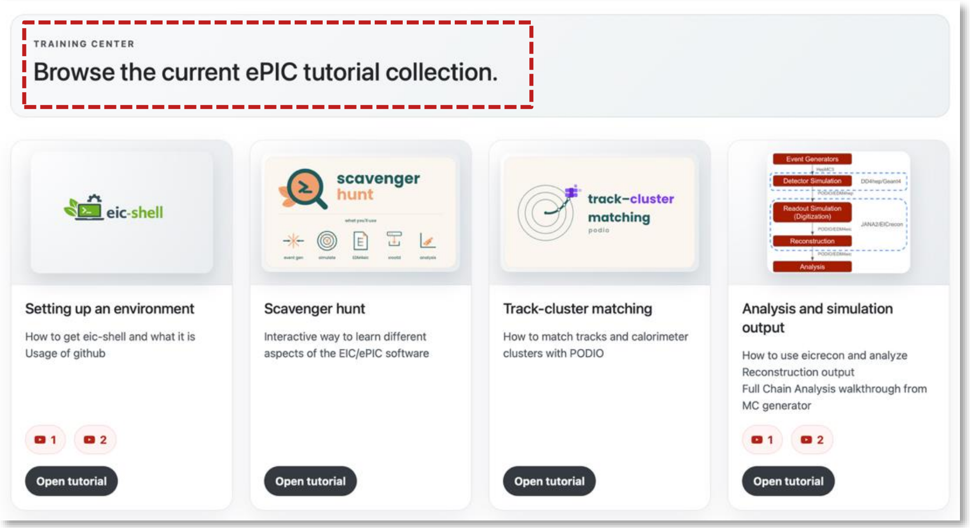

Reorganization and new features, including **mobile-device support** and a **search bar**, ensure that the landing page remains an excellent resource for **onboarding new collaboration members** and **supporting training**.

User Learning Activities

The ePIC/EIC Software Scavenger Hunt

- Recent onboarding and software-discovery initiative;
- Interactive way for users to explore and engage with the software;
- Excellent starting point for new users;
- Tips, tricks, and new material for experienced users;
- Series of tasks covering key software components;
- Final challenge to complete and submit;
- Prizes for participation and correct solutions (sticker).

A huge thank you to Gema Villegas for improving and testing the alpha version!

Additional 2026 Learning Activities

- Tutorials updated and converted to the Software Carpentry Workbench format
- **Early Career Workshop**: Track-cluster matching and PODIO by Derek
- **Collaboration Meeting**: AI MCP tutorial by Sasha

Additional Priorities for 2026

Motivation: With substantial progress in analysis prototyping, it is timely to discuss analysis requirements for Software & Computing.

Analysis Tools Review analysis workflows and requirements with the Physics WGs to improve reconstruction and simulation outputs and define priorities for the ePIC analysis model and tools.

2026 Action: Jointly with the PACs, host a workshop analysis model, tools, and workflows.

Motivation: AI prototypes are ready to move from demonstrations into software, simulation, reconstruction, and analysis workflows.

AI at Scale Coordinate and support AI development within ePIC. Use integration into production workflows as the primary success metric and provide ML operations infrastructure and support.

2026 Action: Agentic workflows for production and DISpatcher AI bot. Advance this topic through the AI Workfest and MCP tutorial.

Recommendation: Develop a data management and lifecycle plan with the ePIC Collaboration, host laboratories, and funding agencies.

Data Develop a data management and lifecycle plan that supports AI readiness, multimodal data and metadata, and reproducible, reusable data and analysis workflows.

Status: Initial work on reusability has started through data cards, validation workflows, and visualization in Hydra.

Three-Year Software & Computing Plan

Planning approach: We are shifting from year-to-year planning to a three-year planning horizon..

Bottom-up • Collaborative

1. **Work with WGs**
   Build on existing three-year outlooks.
   DONE

2. **Define 3-Year Goals**
   Define three-year goals for each WG.
   DONE

3. **WG Plans**
   Each WG develops its plan using an example plan as guidance.
   DONE

4. **Draft Plan**
   Draft plan includes infrastructure plans and resource estimates.
   DONE

5. **Discuss & Refine**
   Share draft with collaboration and refine through discussion.
   NEXT

6. **Final Plan**
   Prepare final plan for October review.
   OCT REVIEW

Next step: Discuss the draft plan with the collaboration by End of July and refine it.

Three-Year Plan Example: User Learning

Pillars: Engagement • Training, including train-the-trainer support • Documentation

	PRIORITIES	METRICS
2026	**Documentation:** prioritize documentation of reconstructed objects and simulation outputs; improve software findability with AI tools and better landing-page search. **Training:** create Jupyter notebooks as entry points; ensure every tutorial includes a working example; launch a pilot tutorial on AI/MCP tools. **Engagement:** launch a software scavenger hunt; add user profiles to personalize onboarding.	• landing-page visits and search usage • % of tutorials with a working example • scavenger-hunt participation and user-profile engagement
2027	**Documentation:** expand data cards and AI-assisted documentation/findability tools to more software products. **Training:** build a curriculum for real detector/physics problems; expand workflow notebooks; launch a train-the-trainer initiative. **Engagement:** reassess onboarding using 2026 results; expand profiles with community/outreach features; establish tutorial review and update process.	• number of documented software products and completed data cards • number of trained contributors and new users completing advanced workflows • tutorial review rate and % updated within the previous year
2028	• Evaluate software findability, reusability, training effectiveness, and user engagement across the collaboration. • Review coverage and quality of data cards and reconstructed-object documentation. • Use findings to define priorities for the next phase.	• collaboration-wide survey results • documentation coverage and training outcomes • workshop recommendations / next-phase priorities

NIM A Special Issue: Software & Computing Contributions

The ePIC Streaming Computing Model
Comprehensive description of the ePIC Streaming Computing Model, including its Echelon architecture and the streaming, AI, and distributed computing technologies that enable its implementation. Full draft available.

ePIC Software and Simulation
TIC recommended: A paper documenting the ePIC software stack and monthly simulation campaigns used for the NIM A Special Issue. Outline reviewed by the TIC.

ePIC Streaming Reconstruction
Description of event building and reconstruction for streaming data, including a comparison of two early prototypes: one based on conventional methods and one using AI. Building up on work by EIC Japan.

A Streaming Reconstruction Workshop will be held August 19–21 to advance the publication.

AI in ePIC and Community Collaboration

AI is embedded across the ePIC computing model

- **AI-driven workflows:** alignment, calibration, validation, control, diagnostics, and production automation.
- **End-to-end processing:** from detector readout to physics analysis, including closed-loop approaches.
- **Integration:** following AI activities across ePIC, including the **AI Town Hall on Thursday**, and bringing relevant developments into the ePIC software stack.

Next steps with AI4EIC

- **Joint AI training:** Launch EIC/ePIC-specific training in early 2027 in response to community requests.
- **Continuous knowledge transfer:** Establish quarterly AI4EIC–ePIC Software & Computing meetings to share relevant methods, tools, and results.

Strengthening Event Generator Efforts in ePIC

Launch a coordinated validation effort that connects physics needs, generator development, and simulation readiness.

The goal is to compare the physics event generators used in the TDR with data and, through these comparisons, validate their underlying models for EIC science.

Physics generators | **Validation against data** | **Reliable inputs For studies**

How we organize the effort

COORDINATION

Validation WG

Andy Buckley has been nominated as co-convener of the Validation WG and will work with event-generator liaisons from the PWGs.

PWG INPUT

Identify liaisons

Ask each PWG to discuss internally and nominate a liaison for generator validation.

WORKFORCE

Open participation

Several people have expressed interest. We welcome more contributors across the collaboration.

COORDINATION

Simulation WG

Kolja Kauder has been nominated as co-convener of the Simulation WG to coordinate and strengthen these efforts alongside Andy.

This effort aligns with MC4EIC priorities and enables close collaboration with the MC4EIC community.

Summary

Computing turns detector data into physics insight through calibration, reconstruction, validation, and analysis.

The **ePIC Computing Model** defines a global fabric required for data processing and analysis. It is driven by physics requirements and built on streaming readout, AI-enabled workflows, and distributed computing.

ePIC is building its **distributed compute and storage fabric** now. International computing resources are expected to play a substantial role.

Priorities

Compute-Detector Integration: Define DAQ-computing interfaces and deliver a testbed for streaming orchestration, with autonomous calibration workflows as test cases.

Enable Science: Monthly software releases, simulation campaigns, production automation, and progress in simulation, reconstruction, and user support have enabled the preTDR and Early Science Report and are enabling the NIM-A Special Issue and future milestones.