ePIC Software and Simulation Campaigns

Dmitry Kalinkin

January 20, 2026

“Software is the Soul of the Detector”

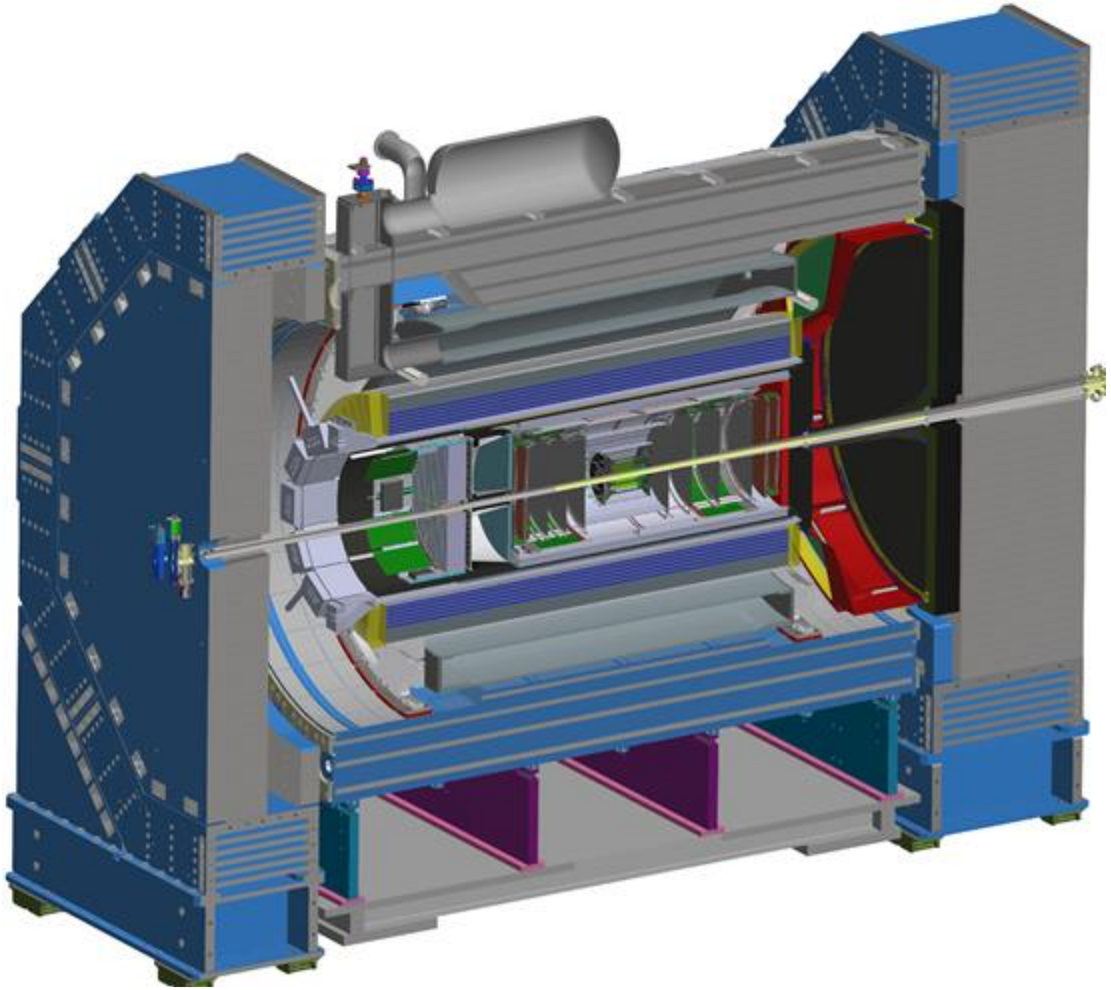

Great Software for Great Science:

- **Design and Construction:** Integrated and validated simulations are essential for evaluating detector performance and determining physics reach.
- **Operation:** Rapid processing of streamed data 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**.

User-Centered Software Design at EIC

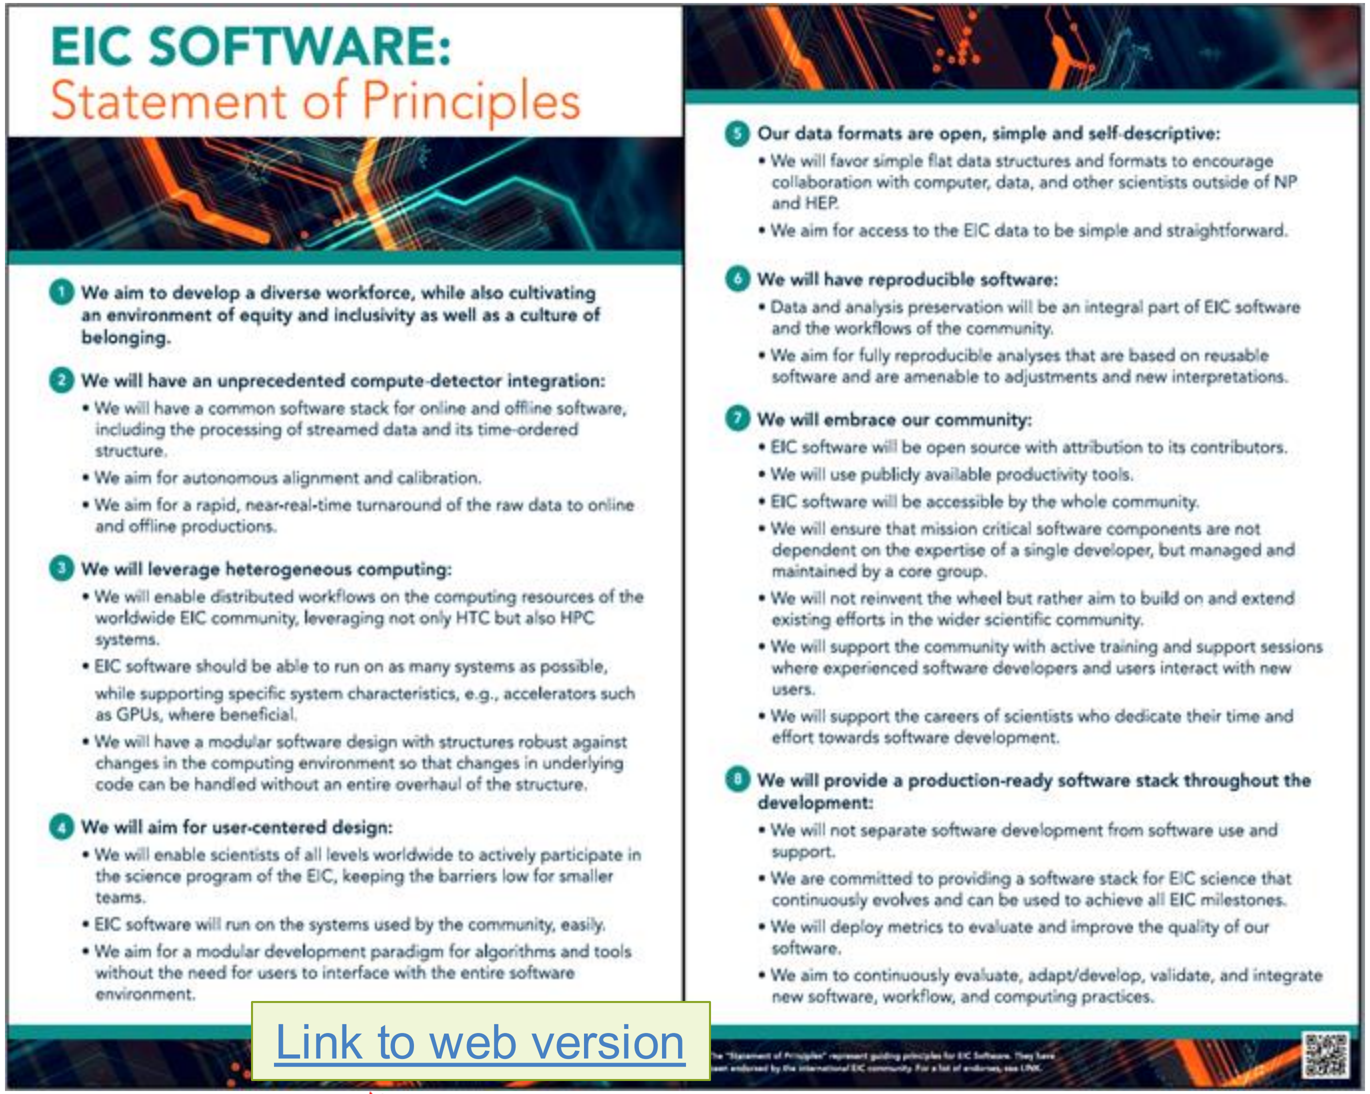

EIC SOFTWARE: Statement of Principles

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.
5. **Our data formats are open, simple and self-descriptive:**
   - We will favor simple flat data structures and formats to encourage collaboration with computer, data, and other scientists outside of NP and HEP.
   - We aim for access to the EIC data to be simple and straightforward.
6. **We will have reproducible software:**
   - Data and analysis preservation will be an integral part of EIC software and the workflows of the community.
   - We aim for fully reproducible analyses that are based on reusable software and are amenable to adjustments and new interpretations.
7. **We will embrace our community:**
   - EIC software will be open source with attribution to its contributors.
   - We will use publicly available productivity tools.
   - EIC software will be accessible by the whole community.
   - We will ensure that mission critical software components are not dependent on the expertise of a single developer, but managed and maintained by a core group.
   - We will not reinvent the wheel but rather aim to build on and extend existing efforts in the wider scientific community.
   - We will support the community with active training and support sessions where experienced software developers and users interact with new users.
   - We will support the careers of scientists who dedicate their time and effort towards software development.
8. **We will provide a production-ready software stack throughout the development:**
   - We will not separate software development from software use and support.
   - We are committed to providing a software stack for EIC science that continuously evolves and can be used to achieve all EIC milestones.
   - We will deploy metrics to evaluate and improve the quality of our software.
   - We aim to continuously evaluate, adapt/develop, validate, and integrate new software, workflow, and computing practices.

[Link to web version]

Important aspects:

- agile development
- production-ready software stack
- meeting near-term needs of ePIC
- timeline-based prioritization
- user-centered design

Seen in how we structure our work:

- continuous testing via automation and with community
- exercising full-scale productions with realistic conditions and interfaces
- deliverables and timelines developed with collaboration
- focus on improving accessibility of our software with deliberate design and development of learning materials

ePIC Simulation and Reconstruction Framework

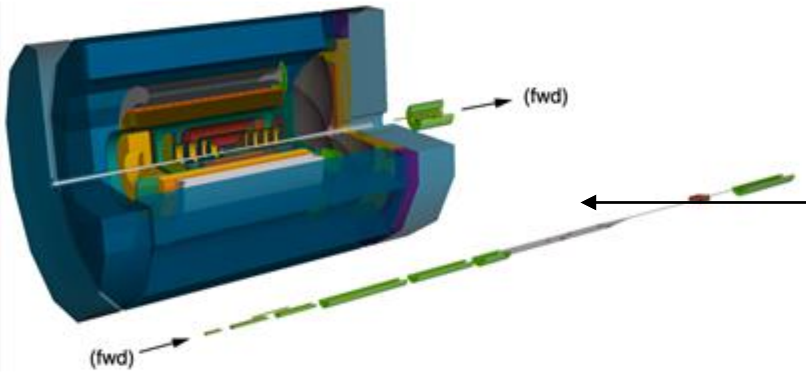

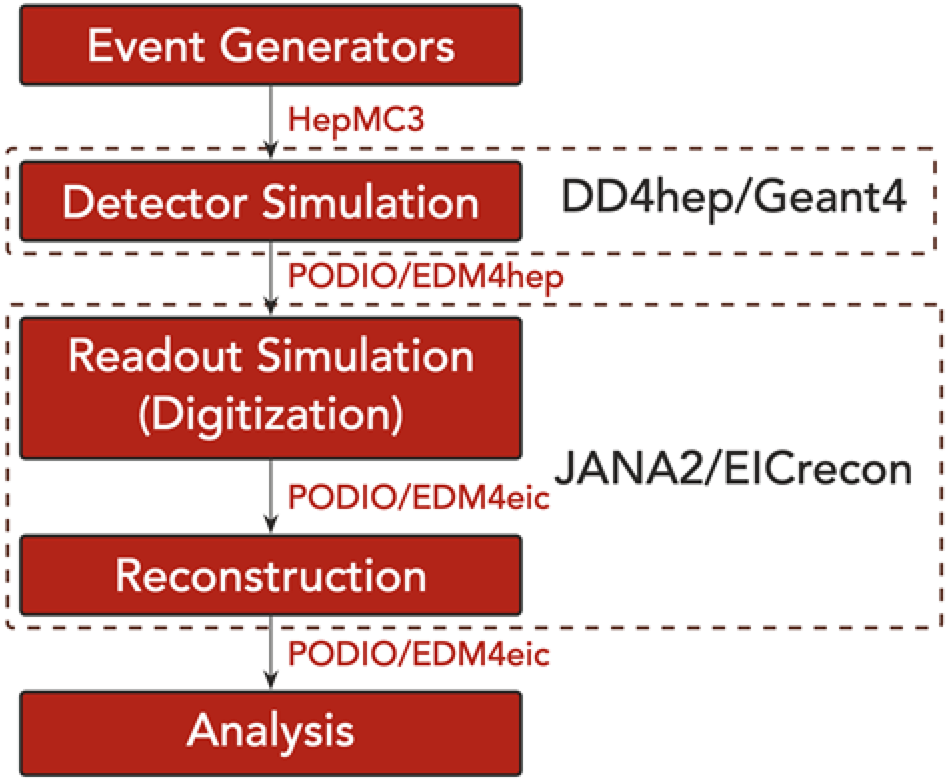

- Common geometry description in DD4hep
- Data model is EDM4eic - an extension of EDM4hep
- Framework requirements include
  - Multithreading support (now in EICrecon)
  - Heterogeneous computing
  - Streaming workflow support
- Analysis capability requirements
  - Holistic reconstruction
  - Test beam data analysis

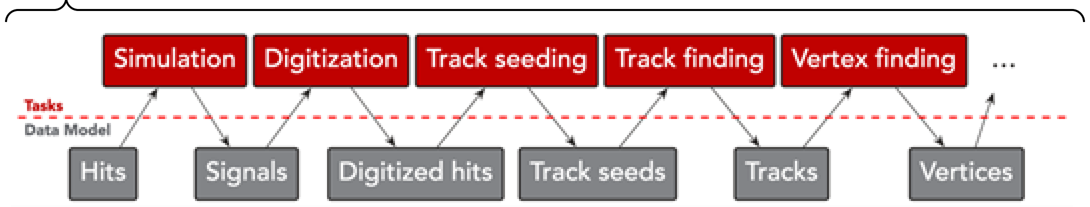

ePIC Software Stack

Software and Computing organized around principles of openness and collaboration with other communities in NP, HEP and CS

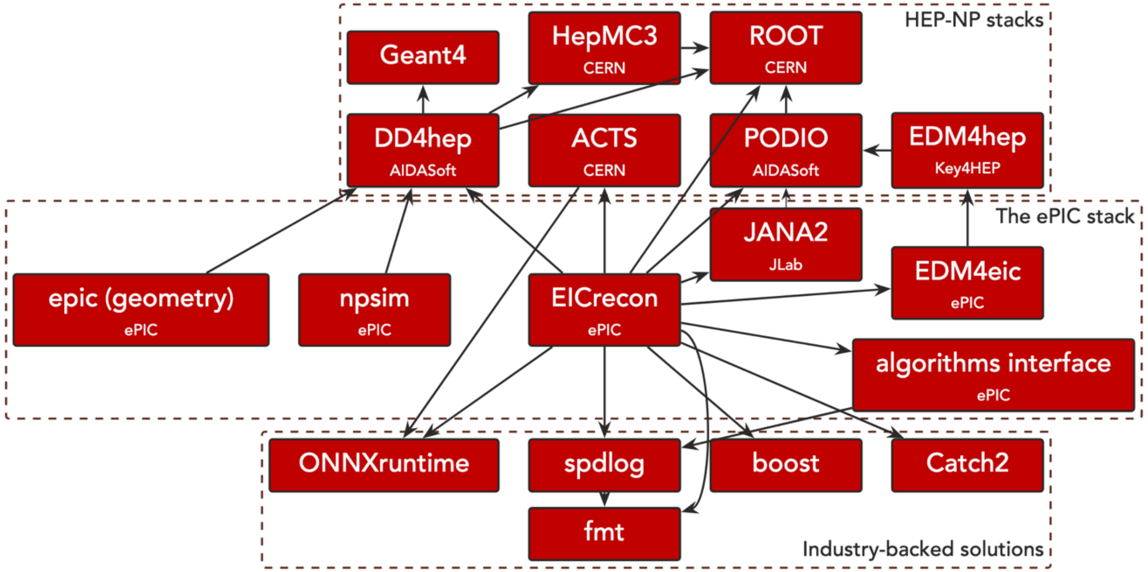

ePIC Software stack built with community components

We develop with the community:

- Strong ePIC representation at Acts4NP
- We participate and contribute to Key4hep discussions
- ePIC participates (remotely) in Acts hackathon
- We upstream our bugfixes and new features to Acts, DD4hep, EDM4hep, JANA2, PODIO, Geant4, HepMC3, ROOT and keep up with latest upstream developements

Highlights from 2025

- **Improved Communication**
  - **Weekly Software News**: Regular reports highlighting major changes to the software stack, notable code merges, and updates from the WGs.
  - **Meeting Notes**: Summaries of WG meetings, including outcomes and next steps, are provided to enable asynchronous participation.
- **Updated WG Charges and Priorities**
  - Reflected outcomes of the collaboration meeting in Frascati, Italy. (for background see in John's slides)
  - Example priorities included background integration in simulations and resolving discrepancies between engineering and simulation designs.
- **Collaboration with Physics WGs**
  - **Coordinated** on **simulation targets**, **reconstruction**, and **analysis tools**.
- **Coordination with DSCs**
  - Requested **updates on software priorities** to facilitate the integration of DSCs simulation efforts into ePIC software and simulation campaigns.
  - Coordinated **improved shared development**, e.g., by aligning four DSCs under a common digitization work plan.

Physics and Detector Simulation

- **Charge:**
  - Development of accurate MC simulations using a suite of physics and background generators and detector simulation based on Geant4 and DD4hep
- **Priorities for 2025:**
  - Continue to support the **detector design** and integration with services.
  - Collaborate with the EIC Project to evaluate the **differences between the engineering and simulation designs**, and lead discussions with the DSCs on how to address these differences.
  - Continue to support the development of **background modeling** and implement its timing structure in physics and detector simulations, together with the Background TF.
  - Enable **simulation of streaming readout** by providing the option to switch between streaming data and event data modes.
  - Coordinate the **development of digitization and noise models** with the DSCs and the Electronics and DAQ WG.

Reconstruction Framework and Algorithms

- **Charge:**
  - Development of a holistic and modular reconstruction for the integrated ePIC detector.
- **Priorities for 2025:**
  - Drive the **development of the reconstruction framework to meet ePIC needs**, e.g., on modularity or streaming data processing.
  - Host collaboration-wide discussions on all aspects of reconstruction, driving the **work toward holistic reconstruction.**
  - Enable reconstruction algorithms to **handle physics events with background.**
  - Collaborate with **PWGs** on **shared reconstruction priorities**, which currently include:
    - Secondary vertexing
    - Hadron identification
    - Particle flow algorithms for jet reconstruction
    - Event kinematics
  - Integrate continued development of **web-based event display** in reconstruction efforts.

ePIC pre-TDR and Publication Plan

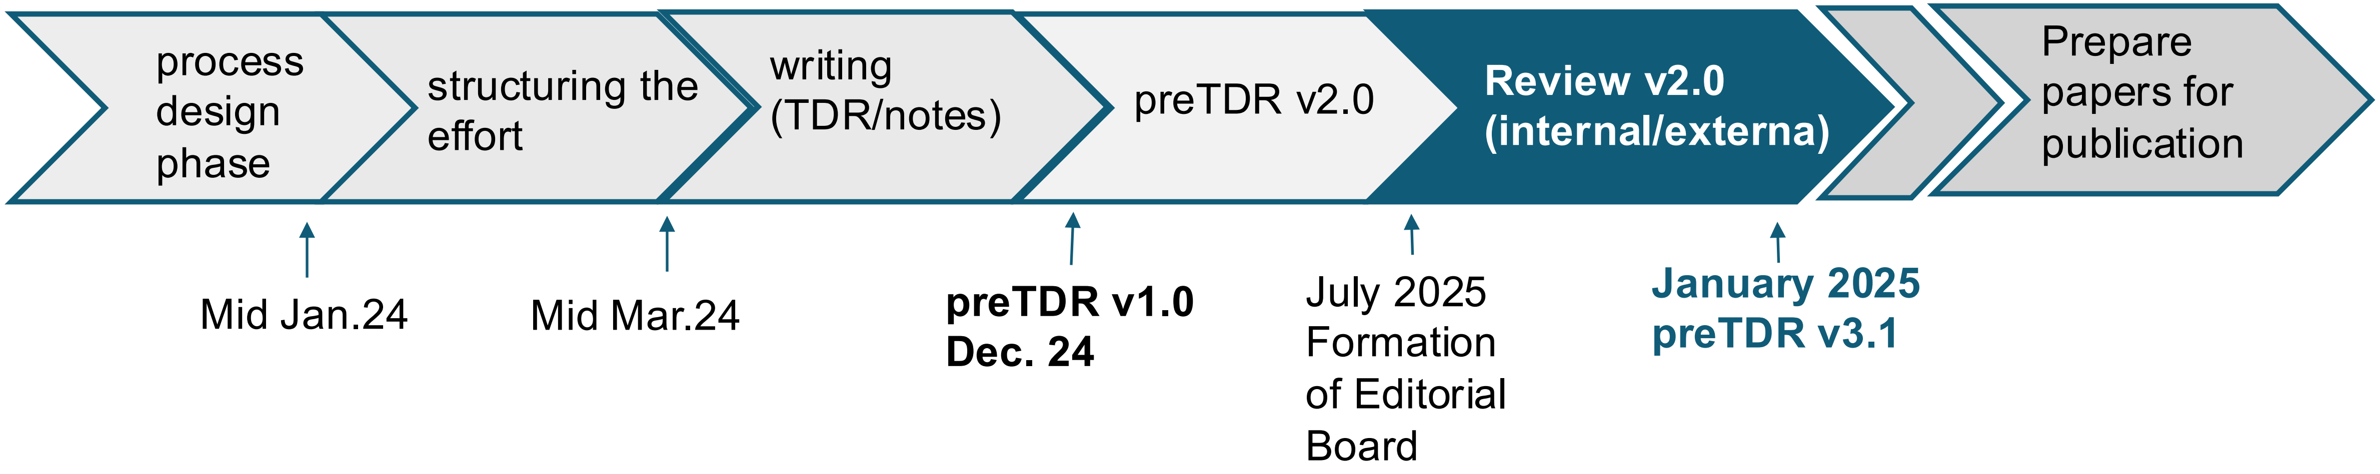

Recommendation (August 15th) - Use **October** (25.10) simulation campaign
List of features for October release finalized in September – all achieved.

Final deadline for finalizing (incl. review) of the contributions (October 7th)

Several software revisions (25.10.1, 25.10.2, 25.10.3, 25.10.4) were released, last of the second priority requests were delivered by late-**November**

Software development continues with the same cadence

Geometry validation

A major task of resolving discrepancies between engineering and simulation designs requires as much coordination as technical work!

DSC

Work with project engineers to converge on subsystem design and files issue and PR for simulation geometry to implement the change

Physics and Detector Simulation WG

Convert subsystem geometry to STEP file using npdet_to_step

Generate drawings using OnShape and provide them to the EIC project together with STEP files

Project

Provide project CAD files and comparison log back

Synchronizing DSC Design, ePIC Simulation Geometry and Project CAD (S. Rahman)

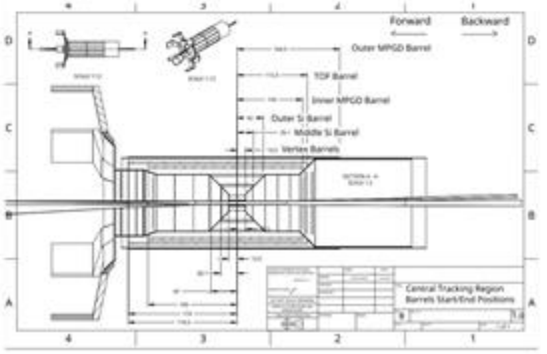

ePIC simulation geometry converted to CAD

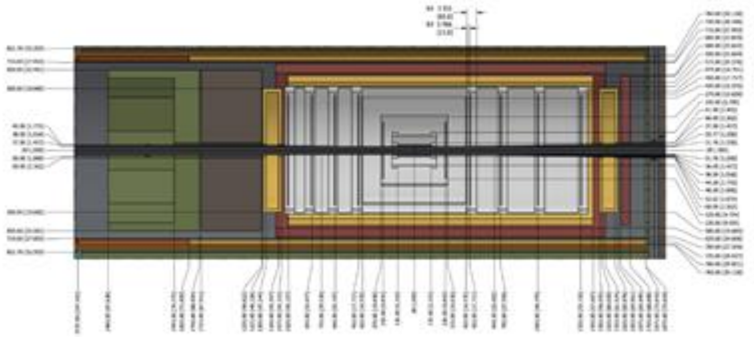

Project detector CAD design

As part of the ongoing design process, DSCs are responsible for addressing discrepancies between engineering and simulation designs. A first round initiated last March, most geometries were validated in October. Latest version of envelopes were released in September.

In 2026, we are looking to provide guides to help detector groups perform this task using tools available.

Background embedding and tracking performance

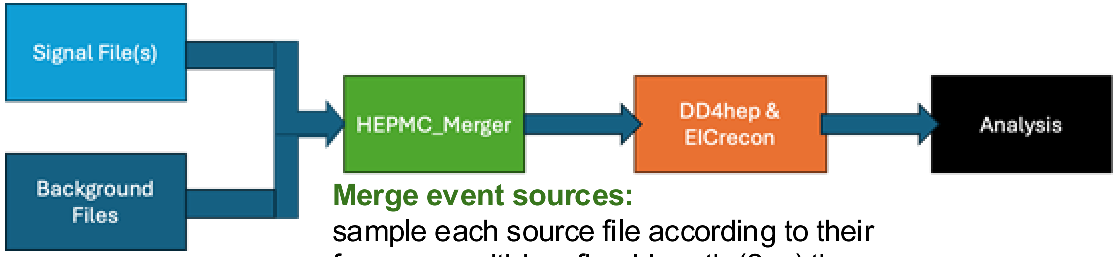

Merge event sources:
sample each source file according to their frequency within a fixed-length (2us) time window

- Merging of full set of backgrounds for various beam energies into physics events implemented in the campaign simulation chain – **enabling highly advanced background studies**
- Working with EIC project on integrating externally-simulated (at beamline volume) Synchrotron Radiation sample; required additional optimizations in the framework
- Especially important for studies on tracking detectors (on the right) and time-frame based reconstruction

eta distribution, 18x275, 1000 time slices

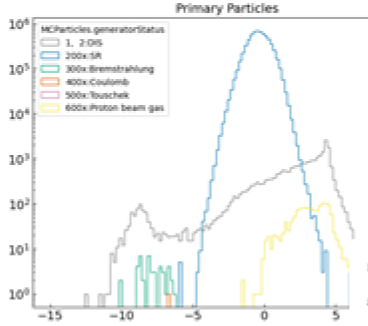

N_TRACK_HIT_MIN = 4 would reject most ghost tracks without hurting the efficiency significantly

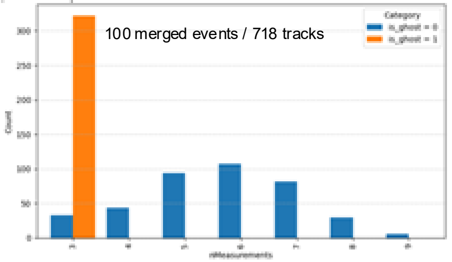

Work in progress by S. Li

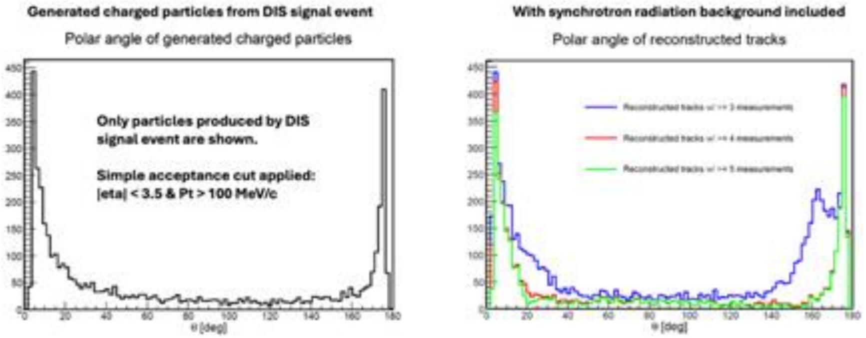

Work in progress by B. Schmookler

Simulation Campaigns

- We provide simulation productions tailored to the needs of the collaboration, as defined by the DSCs and PWGs.
- Simulation campaigns are conducted monthly, based on the software release for the corresponding month.
- These simulations serve as the standard for detector and physics studies for the preTDR and also the Early Science Program.

Monthly Core Hours on the Open Science Grid

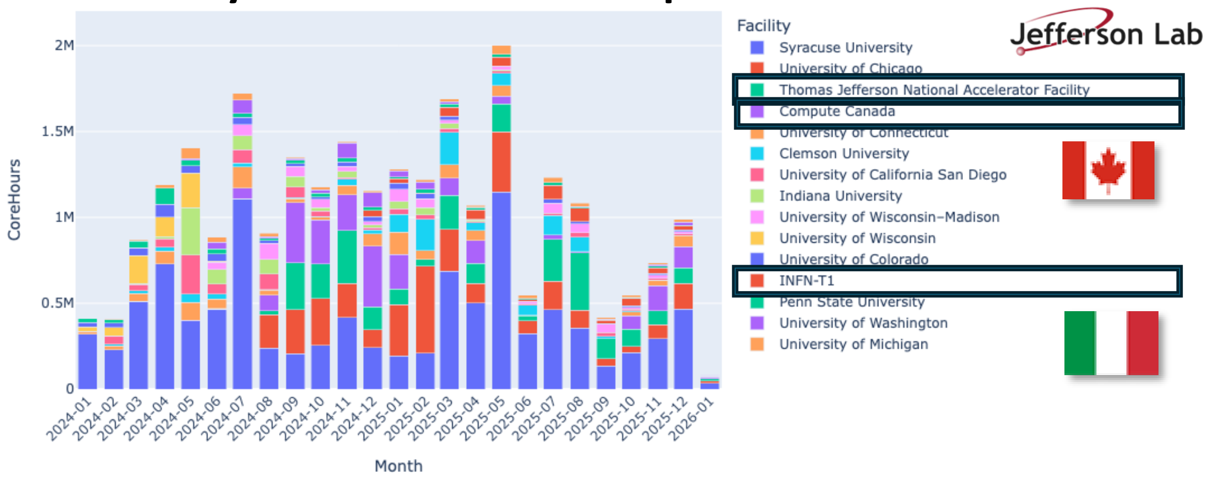

Monthly Storage Usage

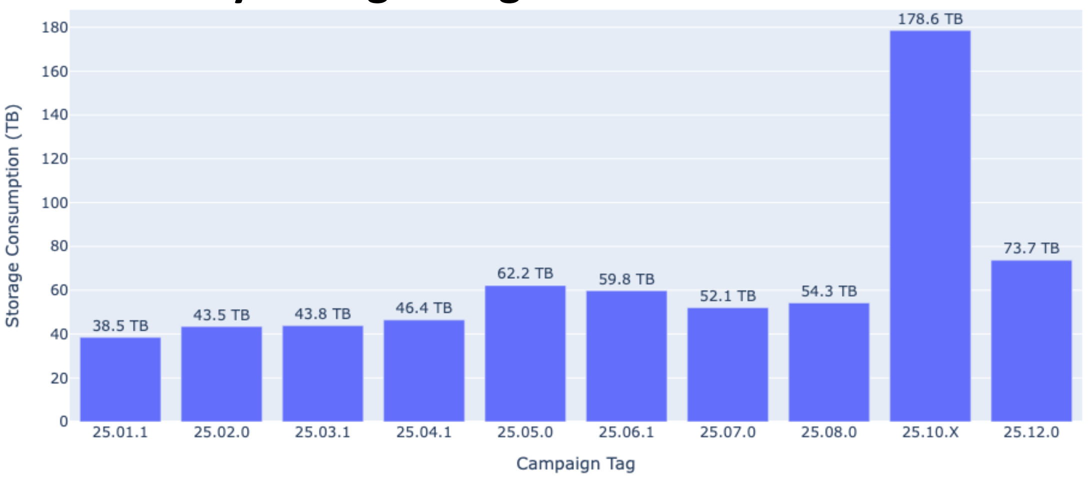

We can integrate new detector geometry and algorithms within a month, processing millions of events needed to assess scientific impact.

Background simulations lead to 2X increase in storage needs – critical time to extend our storage infrastructure in partnership with US and international sites.

Production Software and Infrastructure Development

Existing infrastructure undergoes continuous evolution

- Background event mixing at the HepMC3 file level implemented into the production workflows
  - Ongoing R&D to implement mixing at simulation file level
- Integration of BNL storage element JLab Rucio ongoing for campaign delivery
  - Increase available storage and production redundancy (e.g. for critical time of pre-TDR campaign)
  - Exercise the symmetric BNL-JLab computing model
- Work is ongoing to enable running campaigns using PanDA WMS
  - Integration of additional distributed resources beyond OSG (e.g. Perlmutter)
  - Advanced monitoring and logging facilities will become available
- Redundancy of Docker container building and delivery is being improved
  - Resiliency against CI system outages and CVMFS synchronization delays
- Distributed manpower – grown workforce to 3 production submitters

Developments in Vertex Finding

- The **primary vertex finder** using Acts' IterativeVertexFinder was demonstrated to have a reasonable performance
- **Secondary vertexing** using Helix method integrated, another based on Acts' AdaptiveMultiVertexFinder implemented

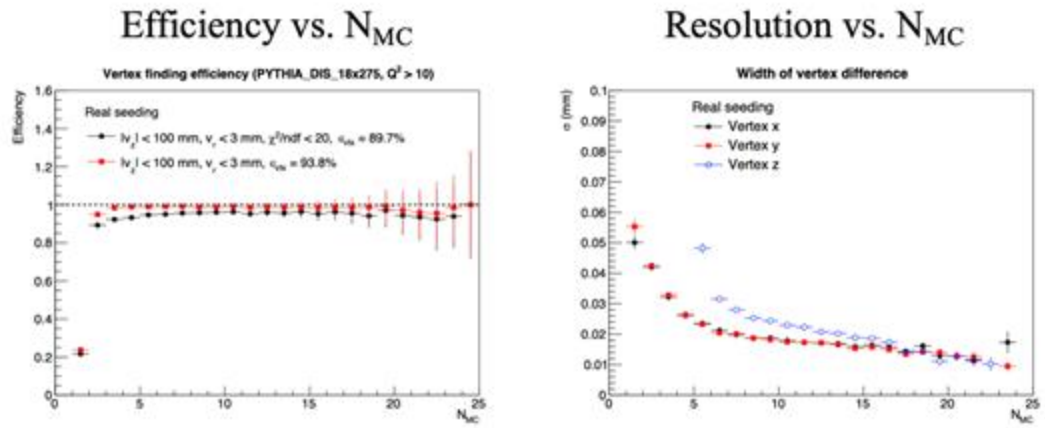

- High vertex finding efficiency at high multiplicity
- About 15 μm resolution at high multiplicity

Slide by Rongrong Ma

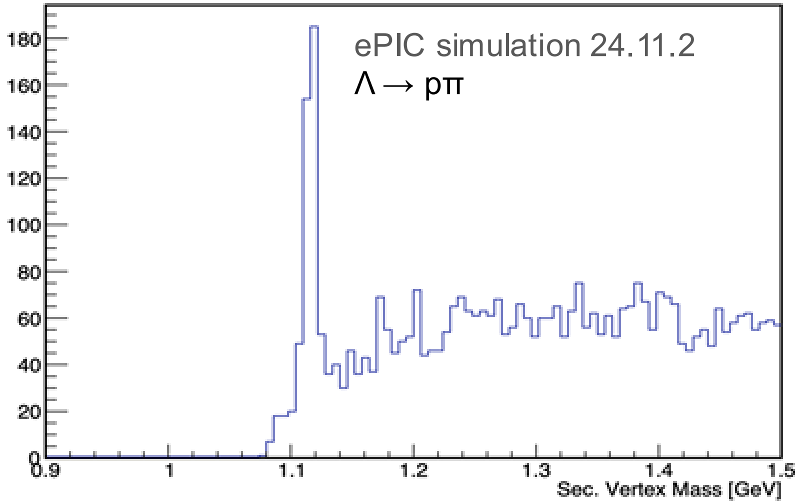

Work in progress by Bishoy H. Dongwi

Particle Flow: Development and Planning

- Initially deemed non-essential for TDR, but later **identified as a need by hadronic calorimetry DSCs as well as Jet/HF PWG**.
- Planning organized according to a staged approach with a goal for modularity in mind.
- Work on Particle Flows also advances implementations for electron ID and holistic reconstruction at ePIC.

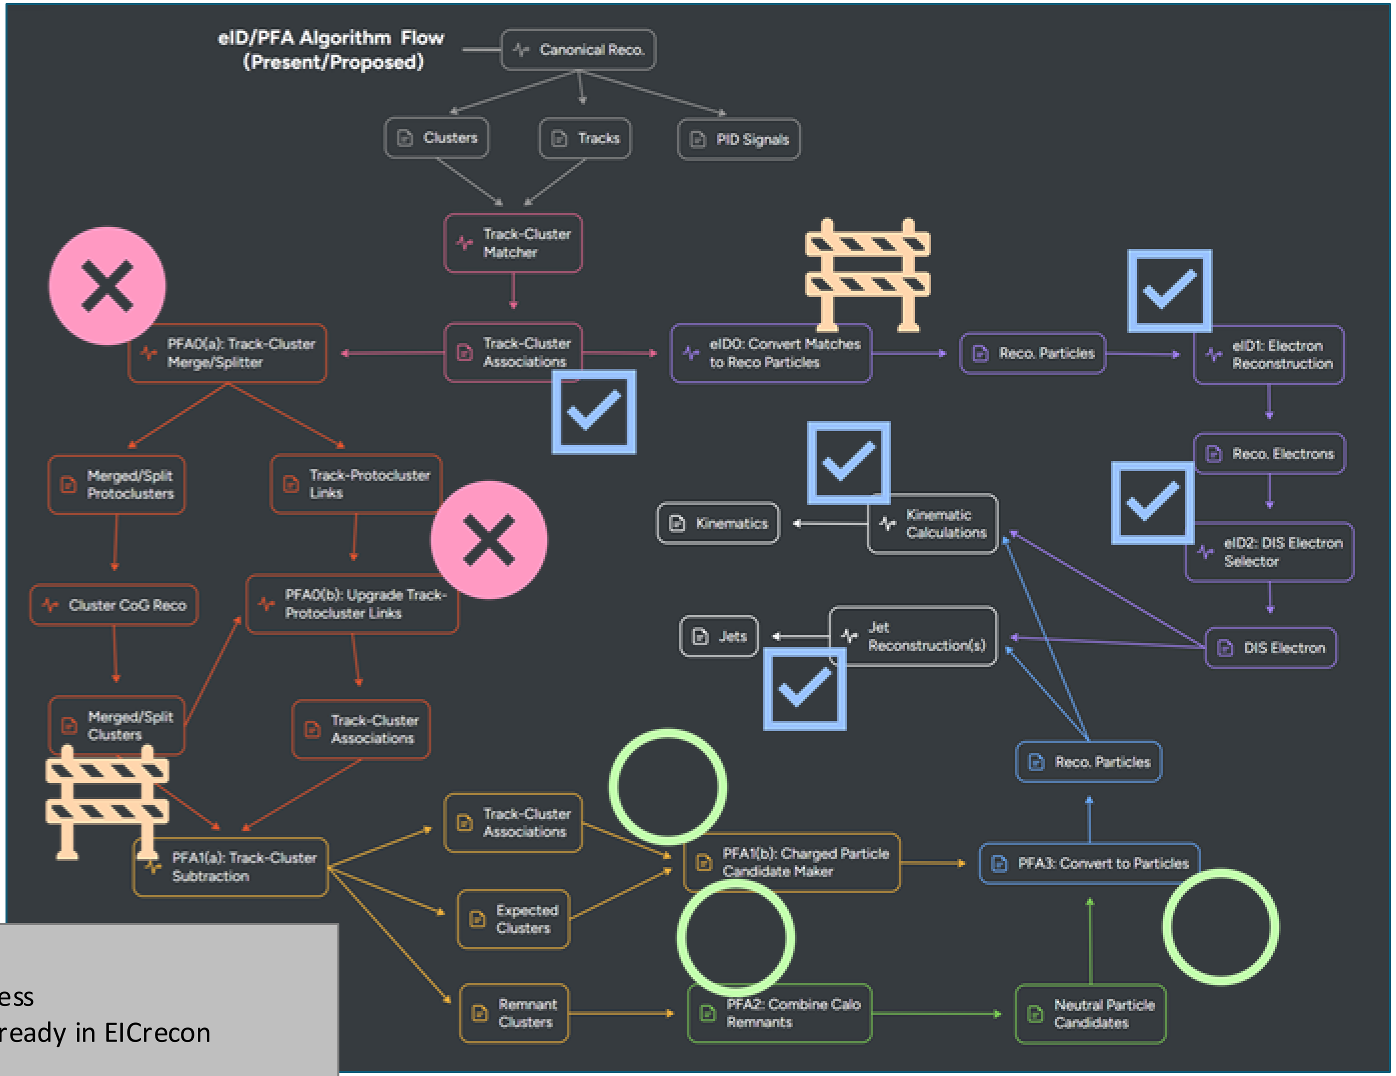

= To-do
= In progress
= Done/already in EICrecon
= Blocked

Chart by D. Anderson

Test Bench and Test Beam: Integration with DAQ and Data

Key priorities

- Getting Test Data into reliable common storage and Rucio tracking
- Planning additions to our PODIO Event Data Model needed to support processing in EICrecon
- Infrastructure for alternative test beam geometries (example implementation provided)

Presentations during last collaboration meeting from DAQ frameworks to enable Test Data collection:

- **RCDAQ** – Martin Purschke (BNL)
  - Ongoing work to implement RCDAQ-to-EICrecon interface
- **Streaming Coda** – D. Abbott, Vardan Gyurjyan, Dmitry Romanov (JLab)
- **nestDAQ** – Noboyuki Kobayashi (OU-RCNP)

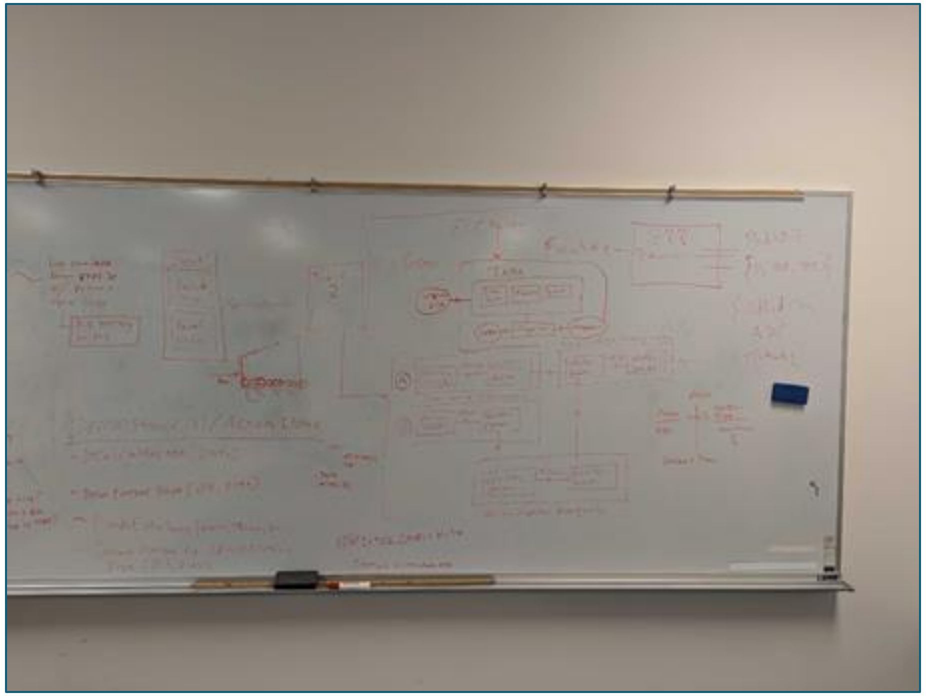

From brainstorming session for Test Beam interfaces.

Supports DSCs and advances simulations

Particle Identification at ePIC

Immediate needs for pTDR are addressed by providing external (standalone) simulations via LUTs. A more integrated approach is expected for the TDR studies.

pfRICH

Effort to integration IRT2 into mainline EICrecon:
- IRT2 namespace to prevent conflict with IRT1
- PRs available

Ongoing work by SBU group to implement an AI method (see talk by Charles)

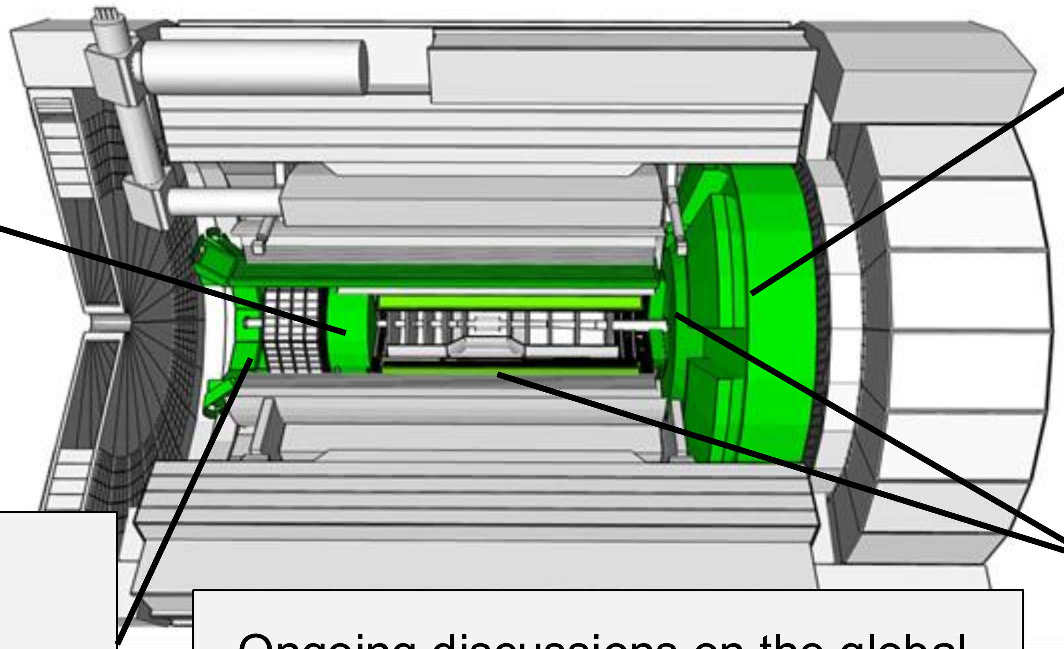

dRICH

- Strongly contributes to the IRT2 effort

Barrel & Encap TOF

- Driving force behind work on waveform-based digitization simulation
- PID reconstruction to be implemented

DIRC

- Promising AI/ML solutions available from several groups

Ongoing discussions on the global PID (e.g. SC-TC meeting)

PID efforts were strengthened by adding third convener to the Reconstruction WG – Chandradoy Chatterjee.

Reconstructing events from Time-frames

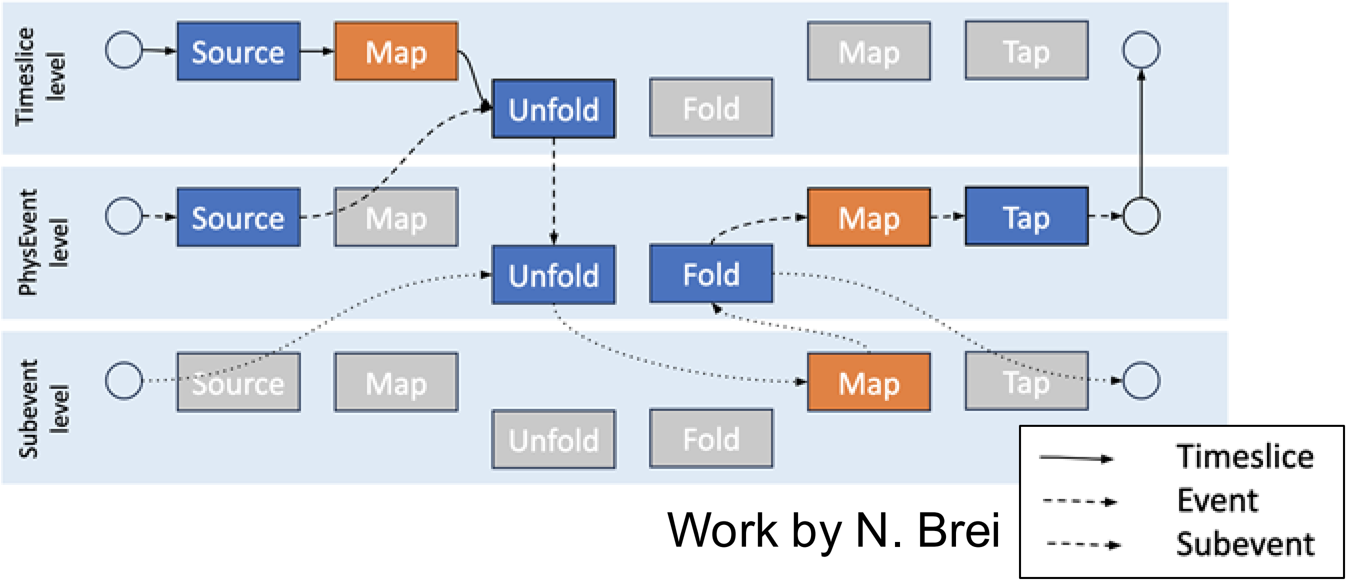

Work by N. Brei

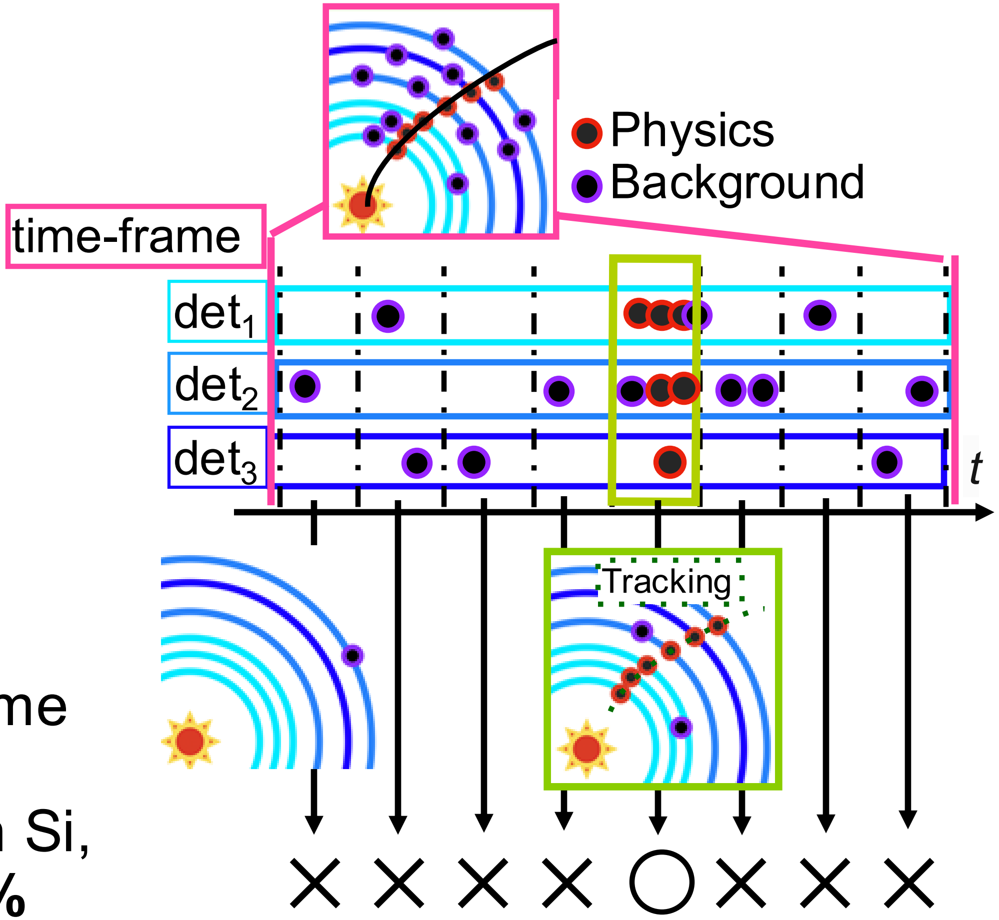

Work by T. Kumaoka

- Utilizes **advanced-topology event processing functionality** from JANA2
- A **prototype algorithm** based on slicing a time frame by ~20 ns window of a 2 μs time-frame
- A simple "trigger" condition on raw simulated hits in Si, TOF and MPGD **achieves >99% efficiency at <1% background**
- We are looking to port this functionality into EICrecon

AI/ML activities at ePIC community

ePIC does not have a dedicated AI/ML WG, unlike EICUG. Instead, WG apply AI/ML methods where appropriate. AI Town Halls support community building.

- Regular ePIC AI Town Hall meetings (1st , 2nd , 3rd)
- Strong ePIC representation at the last AI4EIC Workshop
- ePIC AI/ML Hackathon @ Frascati 2025
- High interest in integrating AI solutions
  - Reconstruction (several algorithms already in EICrecon using ONNX integration provided)
  - Simulation (existing R&D for bottlenecks needs to be integrated)
  - Initial MLOps (long term strategy under planning)
- AI challenge for signal/background discrimination in Timeframes (TBD 2026)

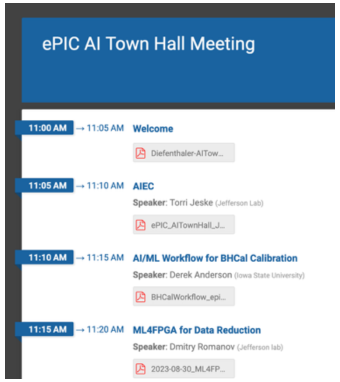

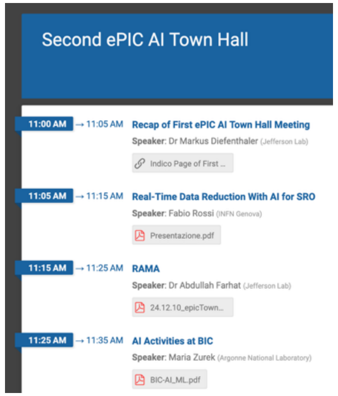

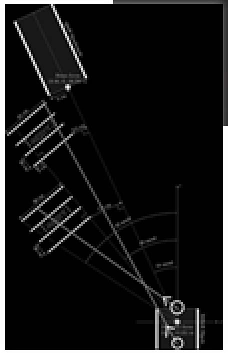

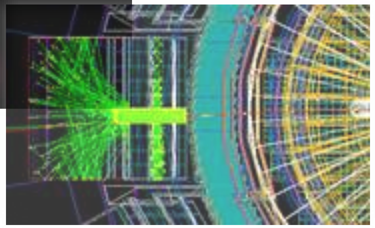

Two problems given:
Low-Q2 tagger calibration
and PID in DIRC

Accelerated Simulations

Heterogeneous computing provides an opportunity to significantly speed up simulations

- We are closely watching latest developments in GPU-accelerated simulation frameworks: AdePT, Celeritas, EIC Opticks
  - certain benchmark results are promising
  - pathways towards integration in the ePIC framework are considered
  - need to assess current and future computing resource availability and compatibility
- Inference for state-of-art Gen-AI simulation models is another avenue for GPU acceleration.

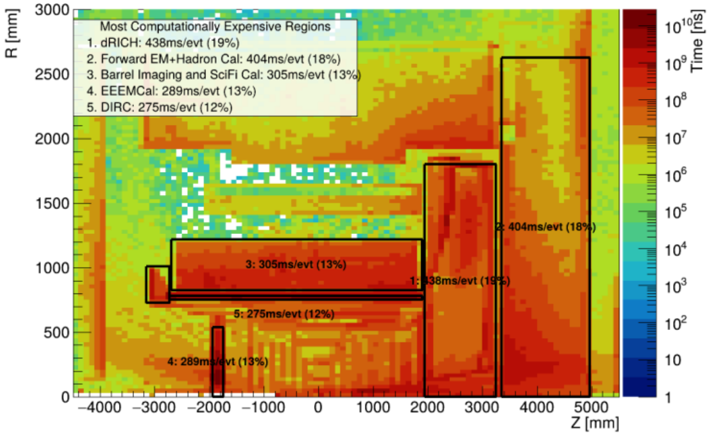

Simulation time budget at ePIC central detector
(W. Deconinck, S. Rahman)

Development Priorities in 2026

Input to discussion (December 2025):

Analysis Tools Review analysis workflows and requirements with the Physics WGs to enhance reconstruction algorithms and simulation output, and to identify priorities for the development of the analysis model and tools for ePIC.

AI at Scale Coordinate and support AI development within ePIC. Use the integration of AI into simulation, reconstruction, and analysis workflows as a primary success metric. Provide MLOps infrastructure and support.

Production Provide the simulation campaigns to finalize the ePIC detector design and validate its physics performance, enhance integration of Rucio within detector and physics analysis workflows, and automate production workflows.

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.

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.

We have collected feedback and additional suggestion by WGs. We will continue developing priorities in the early 2026.

Conclusion

- **In 2024, ePIC defined preTDR readiness for software and simulations —and successfully met those goals. In 2025, priorities were renewed at the Frascati Collaboration Meeting, with many topics already reaching an advanced stage and positioning us well for the next phase.**
- ePIC Software is successfully being used to deliver large-scale simulations to
  - Deliver crucial results needed for finalizing detector design (pTDR)
  - Validating the design’s fitness for the purposes of the EIC science program
- Ongoing software developments are kept aligned with the ePIC Computing Model
- Strong emphasis is placed on **coordination with Physics WGs** on simulation targets, reconstruction
- Surveys of DSCs helped us to have a clear picture for ongoing and planned work. We will repeat and strengthen this practice in 2026.

Backup

Software Release 25.10: Planned vs. Released Features

Planned vs. Released Features

- **Geometry and Materials**: ✅ Updated silicon outer barrel geometry, ✅ revisions to MPGD geometry, ✅ and material map update for FTOF.
- **Reconstruction**: 🚧 ACTS patches to undo regressions with new material map generation, ✅ reworked low-Q^2 tagger momentum reconstruction, ✅ and clustering in TOF.
- **Ongoing Discussions**: ✅ EEEMCal performance to be addressed, ❌ forward ECal geometry, ❌ random noise injection, ❌ and secondary vertexing under discussion.

We met all goals set for the 25.10 software release.

What goes into productions?

Datasets are submitted via a form by Physics Working Groups with generator, its steering files and respective validated HepMC3 files

Dataset priority is defined in cooperation with Analysis Coordination

DSC or PWG	Simulation Path	Generator Config	Background	Events	New Request	Pre-TDR Use	Early Science Use
Exclusive/Diffractive/Tagging	Input Files: /w/eic-scshelf2104/users/sjdkay/Jul2025_Campaign_Input/Afterburner_Output/pion	DEMPgen1.2.4	No		Yes	Maybe	Yes
Exclusive/Diffractive/Tagging	/w/eic-scshelf2104/users/ gbxalex/SimCampaign_Input	lAger3.6.1	No		Yes	Yes	Yes
Exclusive/Diffractive/Tagging	/w/eic-scshelf2104/users/ gbxalex/SimCampaign_Input	lAger3.6.1	No		Yes	No	Yes
Exclusive/Diffractive/Tagging	/gpfs02/eic/wlin/docu/eHe3_10x166_1M	BeAGLE	No		Yes	No	Yes
Exclusive/Diffractive/Tagging	/gpfs02/eic/wlin/docu/eHe3_18x110_1M	BeAGLE	No		Yes	Yes	No
Exclusive/Diffractive/Tagging	/gpfs02/eic/wlin/docu/eHe3_10x110_1M	BeAGLE	No		Yes	Yes	No
Exclusive/Diffractive/Tagging	/gpfs02/eic/wlin/docu/eHe3_5x41_1M	BeAGLE	No		Yes	Yes	No
Exclusive/Diffractive/Tagging	/gpfs02/eic/jkim/DVpi0P/sim_data/	EpIC	Yes		Yes	No	Yes
Exclusive/Diffractive/Tagging	/gpfs02/eic/jkim/DVpi0P/sim_data/	EpIC	No		Yes	Yes	Yes
Exclusive/Diffractive/Tagging	/gpfs02/eic/jkim/DVpi0P/sim_data/	EpIC	Yes		Yes	Yes	No
Exclusive/Diffractive/Tagging	/gpfs02/eic/gpenman/DDVCS_18x275	EpIC	Yes		Yes	Yes	No
Exclusive/Diffractive/Tagging	/gpfs02/eic/gpenman/DDVCS_18x275	EpIC	No		Yes	Yes	No

… 50+ simulation requests processed

Simulation geometry validation status

1	ePIC March 2025 Envelope CAD	https://indico.bnl.gov/event/27186/	
2	ePIC 25.08.0 Simulation Geometry CAD		
3			
4	**System**	**Validated Against Envelope from March, 2025?**	**Validated Against Envelope from September, 2025?**
5	Beampipe	Yes	
6	SVT Inner Barrel		
7	SVT Outer Barrel		
8	SVT Disks Forward	?	?
9	SVT Disks Backward	?	?
10	SVT Support and Services		
11	MPGD Inner Barrel	Differences are Understood and Addressed	
12	MPGD Outer Barrel	Differences are Understood and Addressed	
13	MPGD Disks Forward	Differences are Understood and Addressed	
14	MPGD Disks Backward	Yes	
15			
16	ECal Barrel (BIC)	Differences are Understood and Addressed	
17	ECal Forward (FEMCAL)	?	?
18	ECal Backward (EEEMC)	Yes	
19			
20	HCal Barrel	Not in Agreement	
21	HCal Forward	Differences are Understood and Addressed	
22	HCal Backward	Not in Agreement	
23			
24	AC-LGAD Barrel	Not in agreement	
25	AC-LGAD Disk Forward	Yes	Yes

Software Release 25.10: Patches

Release Version	Release Date	Description
25.10.0	10/17	Initial Release
25.10.1	10/30	EICrecon was updated at Alex Jentsch's request to enable polynomial interpolation in far-forward reconstruction, with the agreement of the Exclusive PWG and PACs.
25.10.2	11/03	Correction to 25.10.1.
25.10.3	11/06	Fixed a memory issue related to backgrounds.
25.10.4	11/11	Correction to 25.10.3.

Enabling pTDR and Early Science Physics Studies

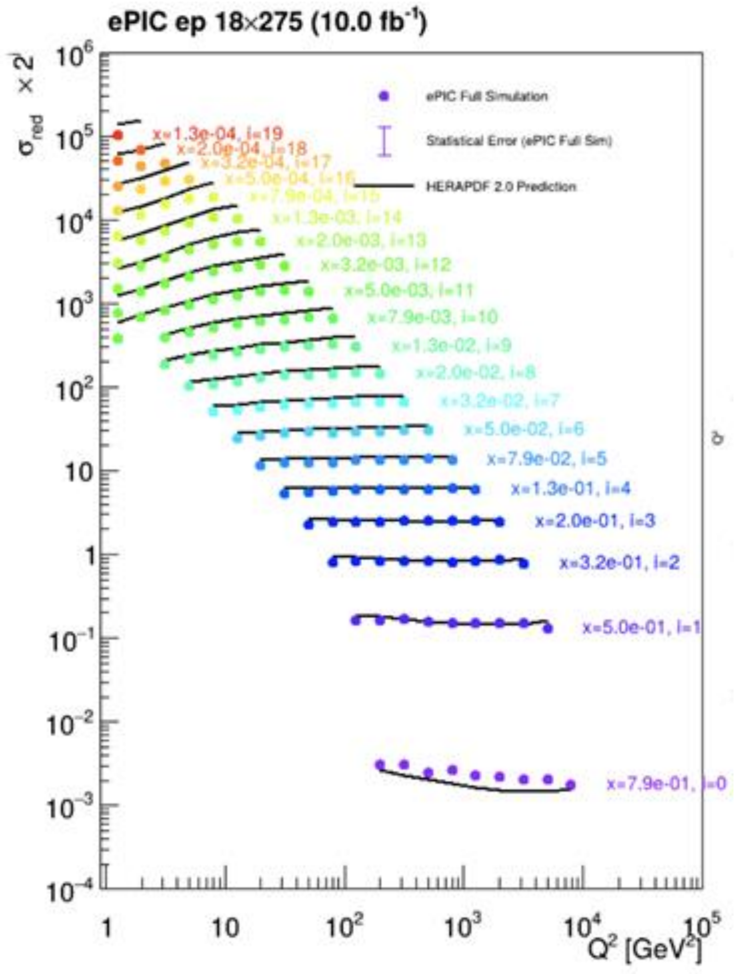

Plot by S. Maple

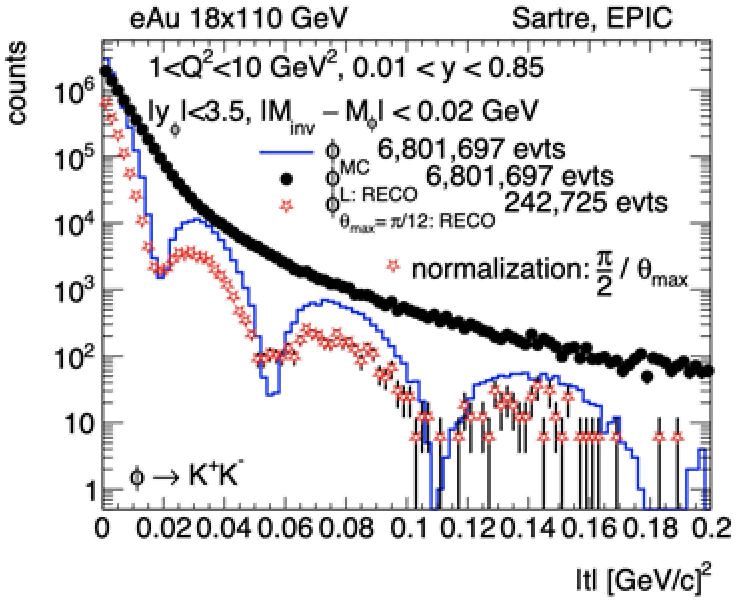

Plot by M. Kesler

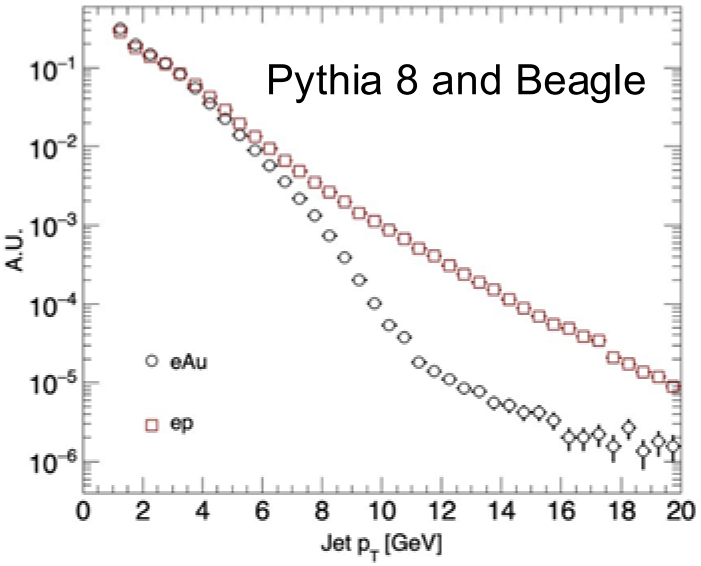

Plot by B. Page and D. Lemos

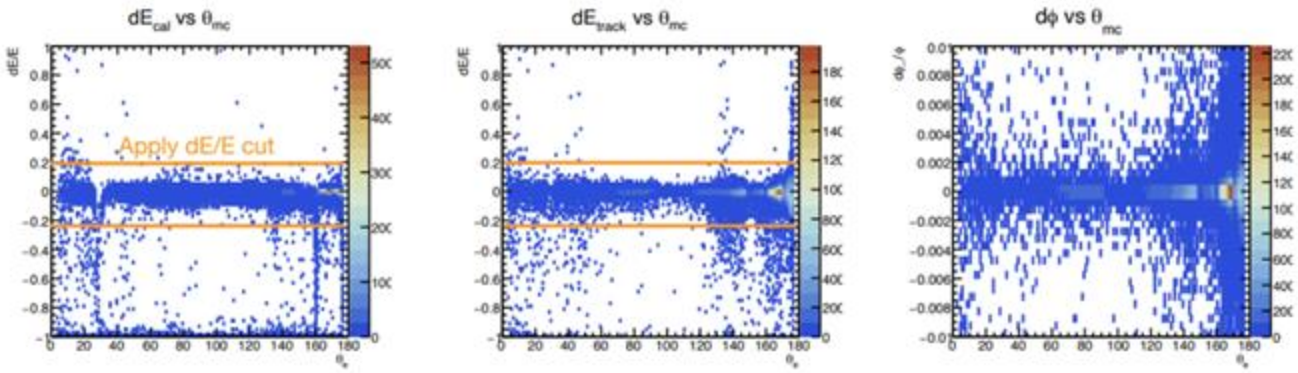

Plot by W. Lin

Enabling pTDR and Early Science Physics Studies

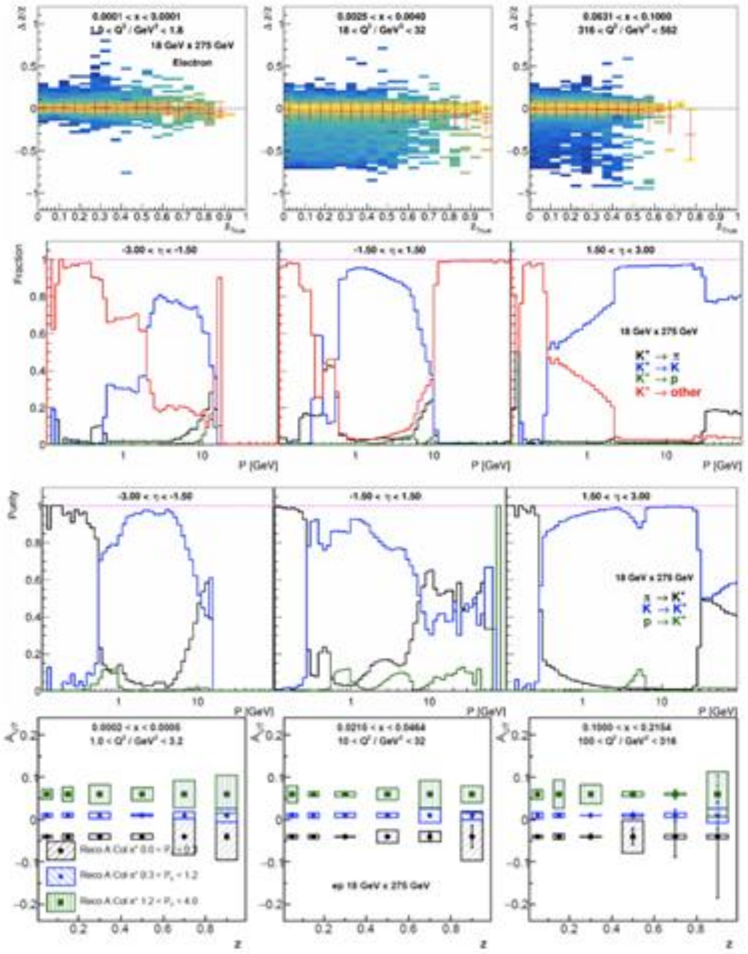

Plot by R. Siedl

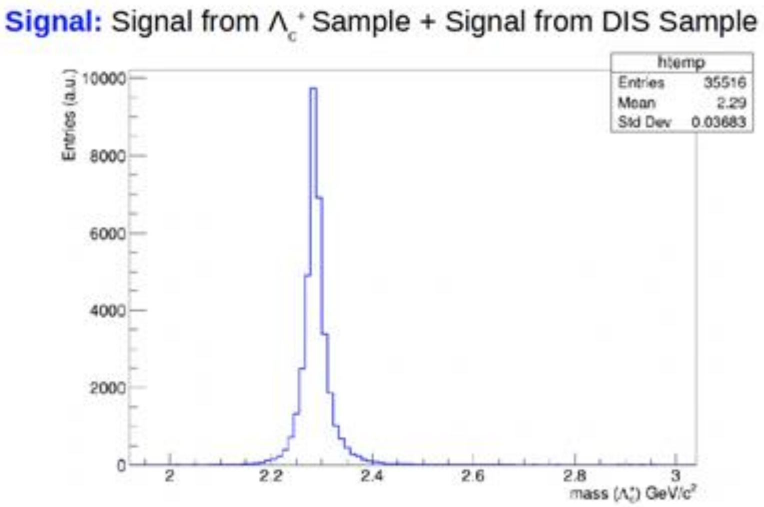

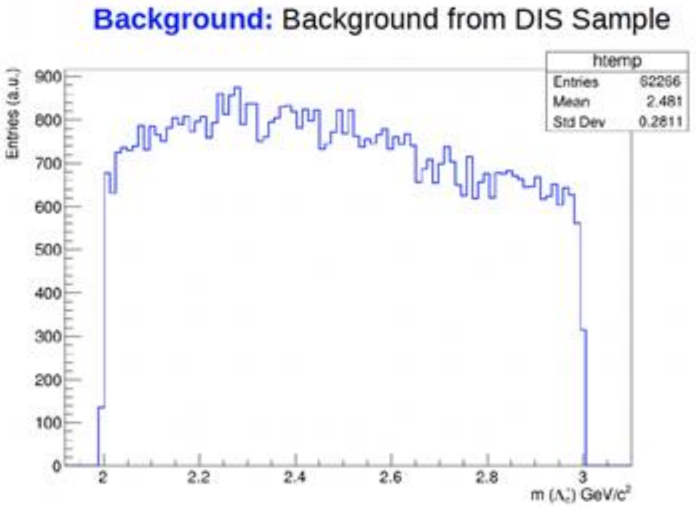

Plots by S. Kumar

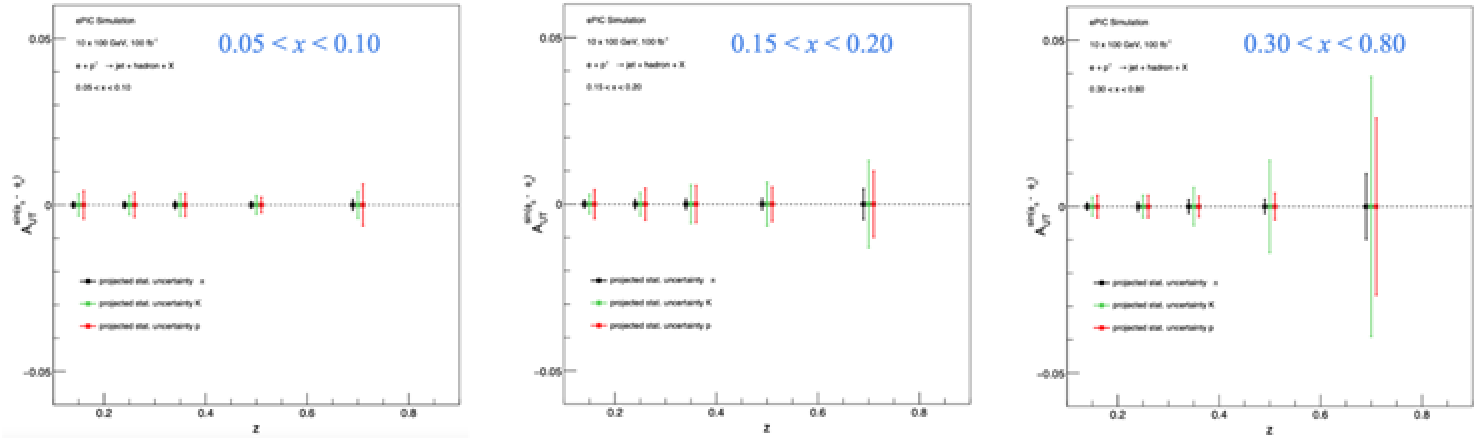

Plots by K. Adkins

MLOps at ePIC

ePIC will deploy more and more AI models, that will all need to be kept up to date with the latest simulations and calibrations.

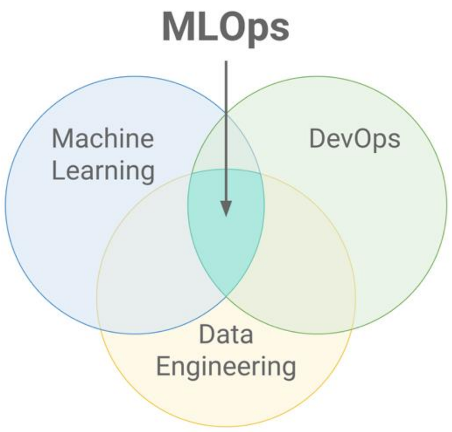

[Attribution CC BY-SA 4.0]

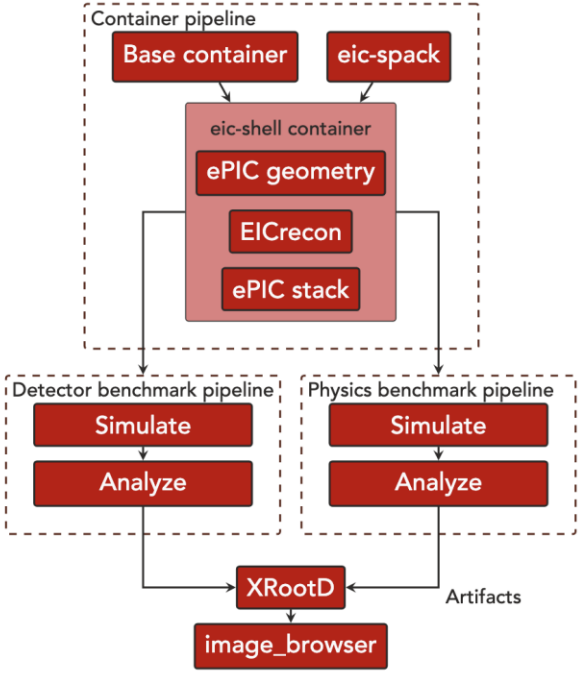

CI pipelines

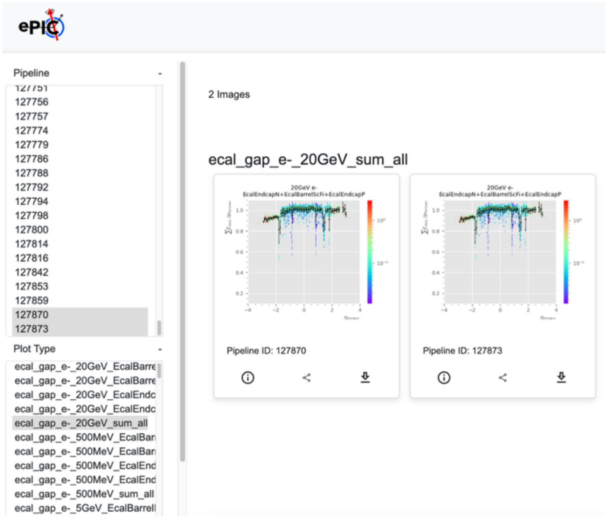

Image artifacts delivered to image_browser

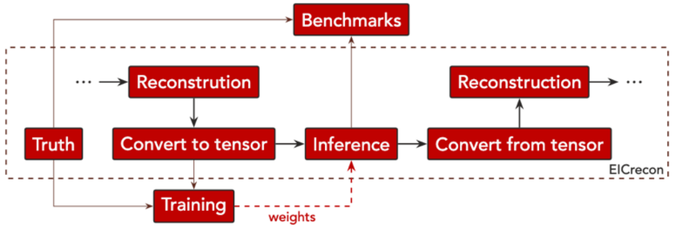

Standartized ONNX factory for EICrecon with model training and validation on CI