Scaling ePIC Simulation Production:
Distributed Workflow and Data Management

Sakib Rahman

on behalf of the ePIC Collaboration

Brookhaven National Laboratory

Computing in High Energy and Nuclear Physics (CHEP)
Chulalongkorn University, Bangkok, Thailand
26 May 2026

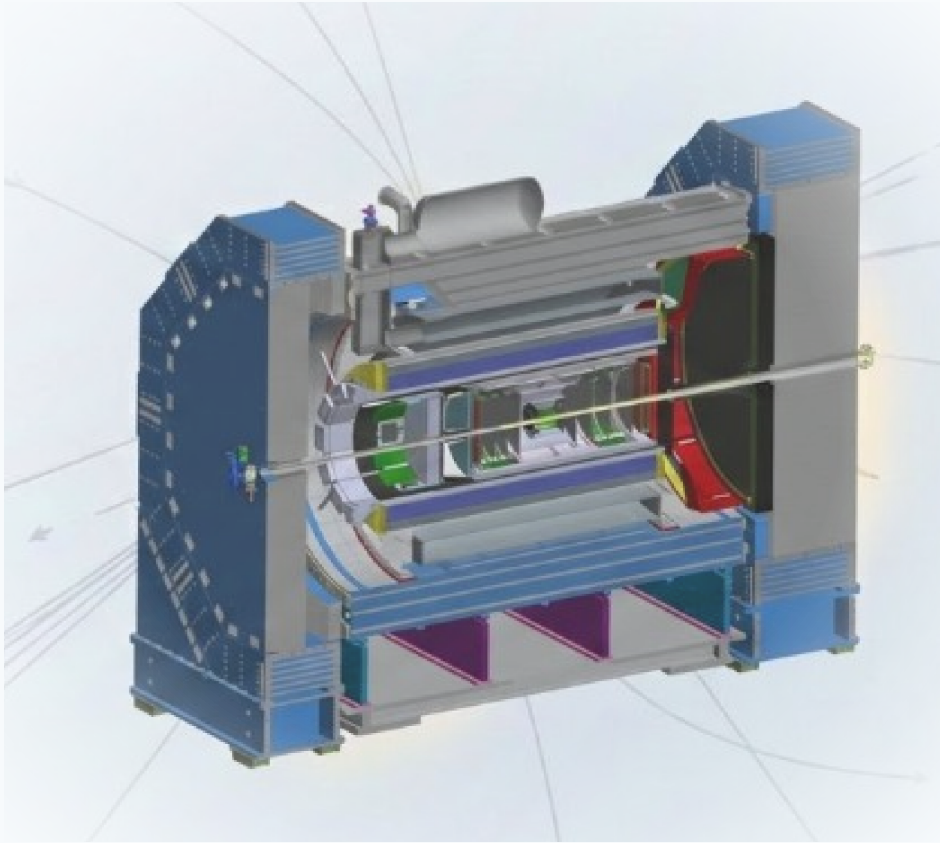

Slide template generated with Claude Sonnet 4.6

The ePIC Experiment at the Electron-Ion Collider

1157

Collaborators

183

Institutions

26

Countries

ePIC will be the first detector at the future Electron-Ion Collider. Data taking planned for the **early 2030s**.

Being realized through a partnership between host labs: **Brookhaven National Laboratory (BNL)** and **Jefferson Lab (JLab)**

Will enable precision studies of nucleons and nuclei at the scale of sea quarks and gluons

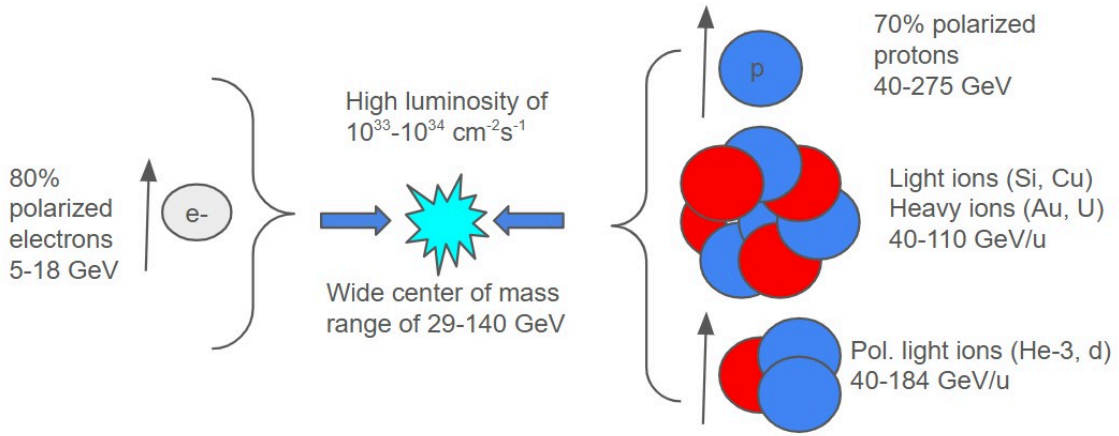

EIC Beam Specifications

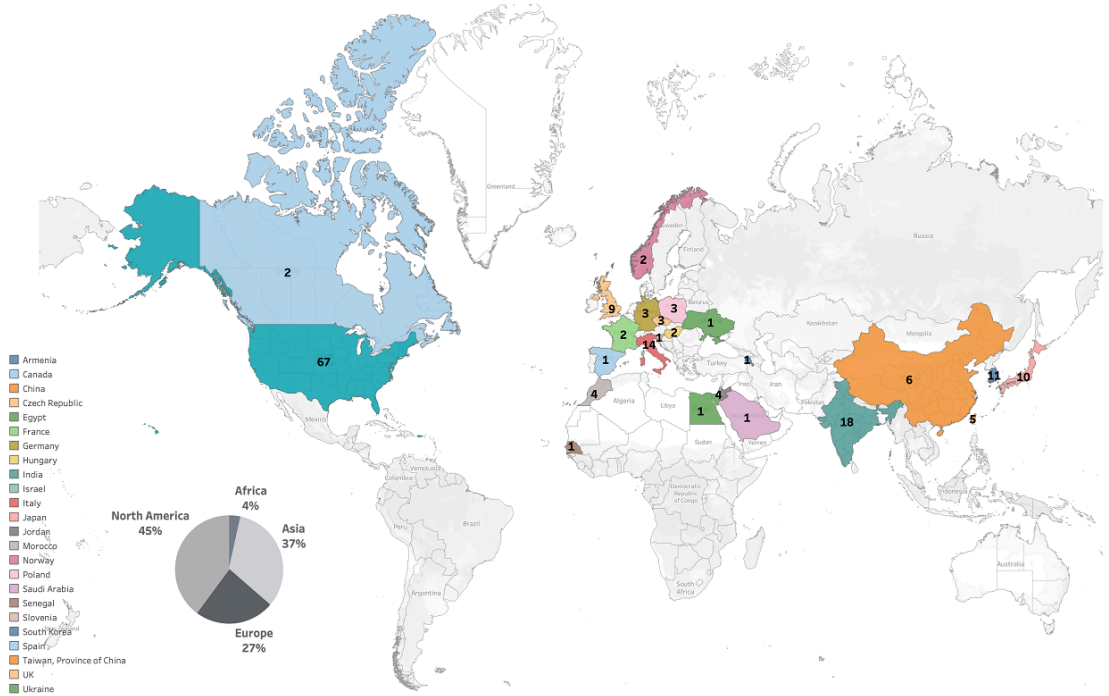

ePIC: A Global Collaboration

The ePIC Production Working Group

- Sits within the **Software & Computing** branch of the ePIC collaboration
- Runs large-scale monthly simulation campaigns on HTC and HPC resources — opportunistic sites via the Open Science Grid (OSG) and allocated domestic & international sites
- Enables detector and physics studies for the Technical Design Report

SPOKESPERSON'S OFFICE		
TECHNICAL COORDINATION	**SOFTWARE & COMPUTING**	**ANALYSIS COORDINATORS**
Tracking	Physics & Detector Simulations	BSM & Precision EW
Electronics, Readout & DAQ	Reconstruction	Exclusive, Diffraction & Tagging
AC-LGAD	Streaming Computing	Jets & Heavy Flavor
Calorimetry	User Learning	Inclusive Physics
PID	⭐ **PRODUCTION**	Semi-Inclusive Physics
	Validation	

Campaigns

Running continuously since **May 2023**, growing to **>1M core hours** and **~150 TB** per month

Infrastructure

Develop and maintain workflow and data management systems that leverage heterogeneous resources efficiently (See my other talk on acceleration pathways)

Operations

Deliver production-scale simulated data to Physics Working Groups (PWG) and Detector Subsystem Collaborations (DSC)

Simulation Production: Compute

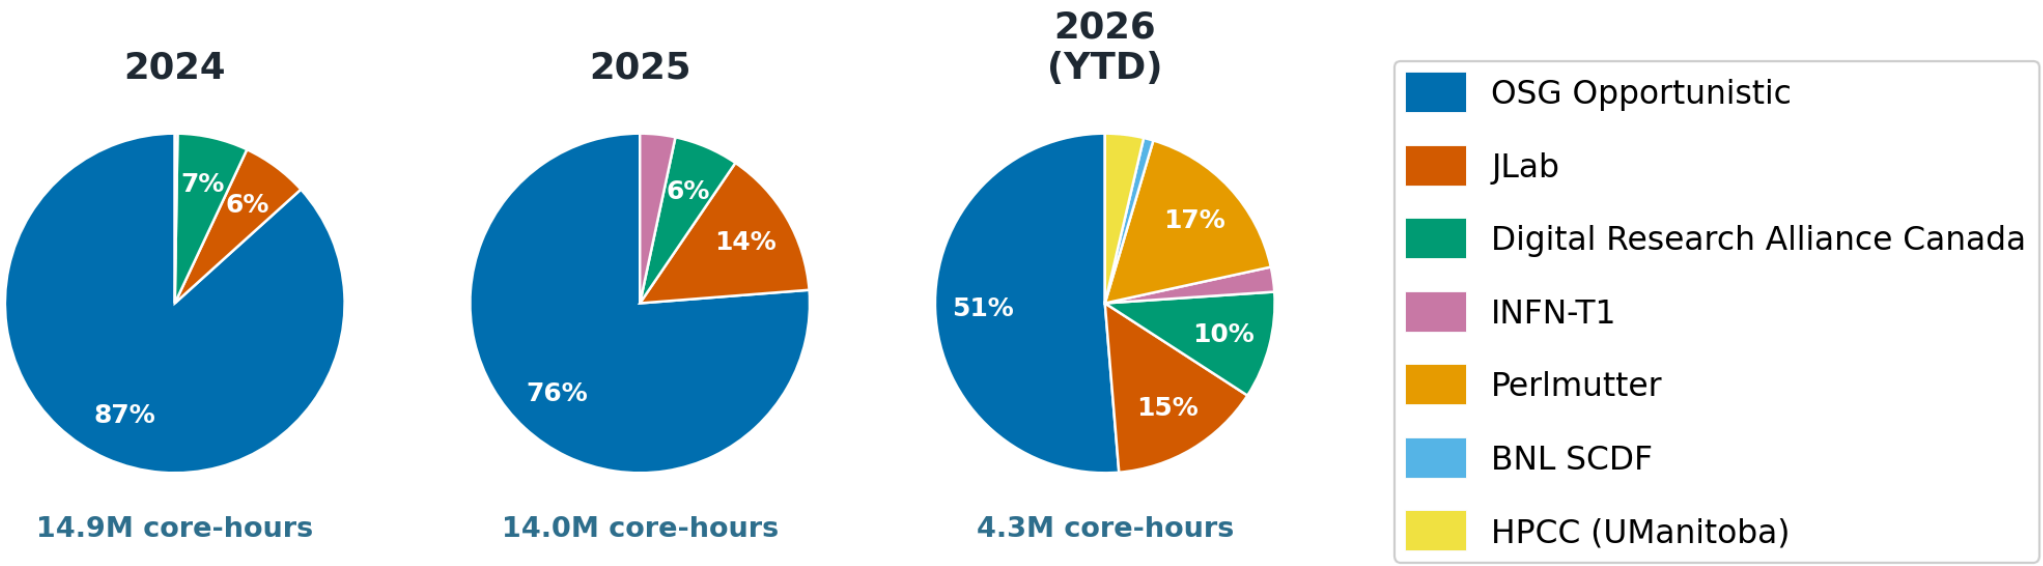

ePIC Computing Usage by Site (2024–2026)

- Continue to rely heavily on opportunistic resources from OSG
- Added more allocated resources in 2026, anticipating growing compute demand for new physics studies and background mixed simulations

Simulation Production: Storage

- **750 TB** disk added at BNL (March 2026) to meet growing collaboration demand
- **500 TB** disk + scalable tape at JLab
- Additional resources being discussed with international partners

- **Storage backends:** BNL uses ZFS, JLab uses Lustre
- Access via **XRootD**, dual instance to simultaneously allow authenticated write and unauthenticated public read access

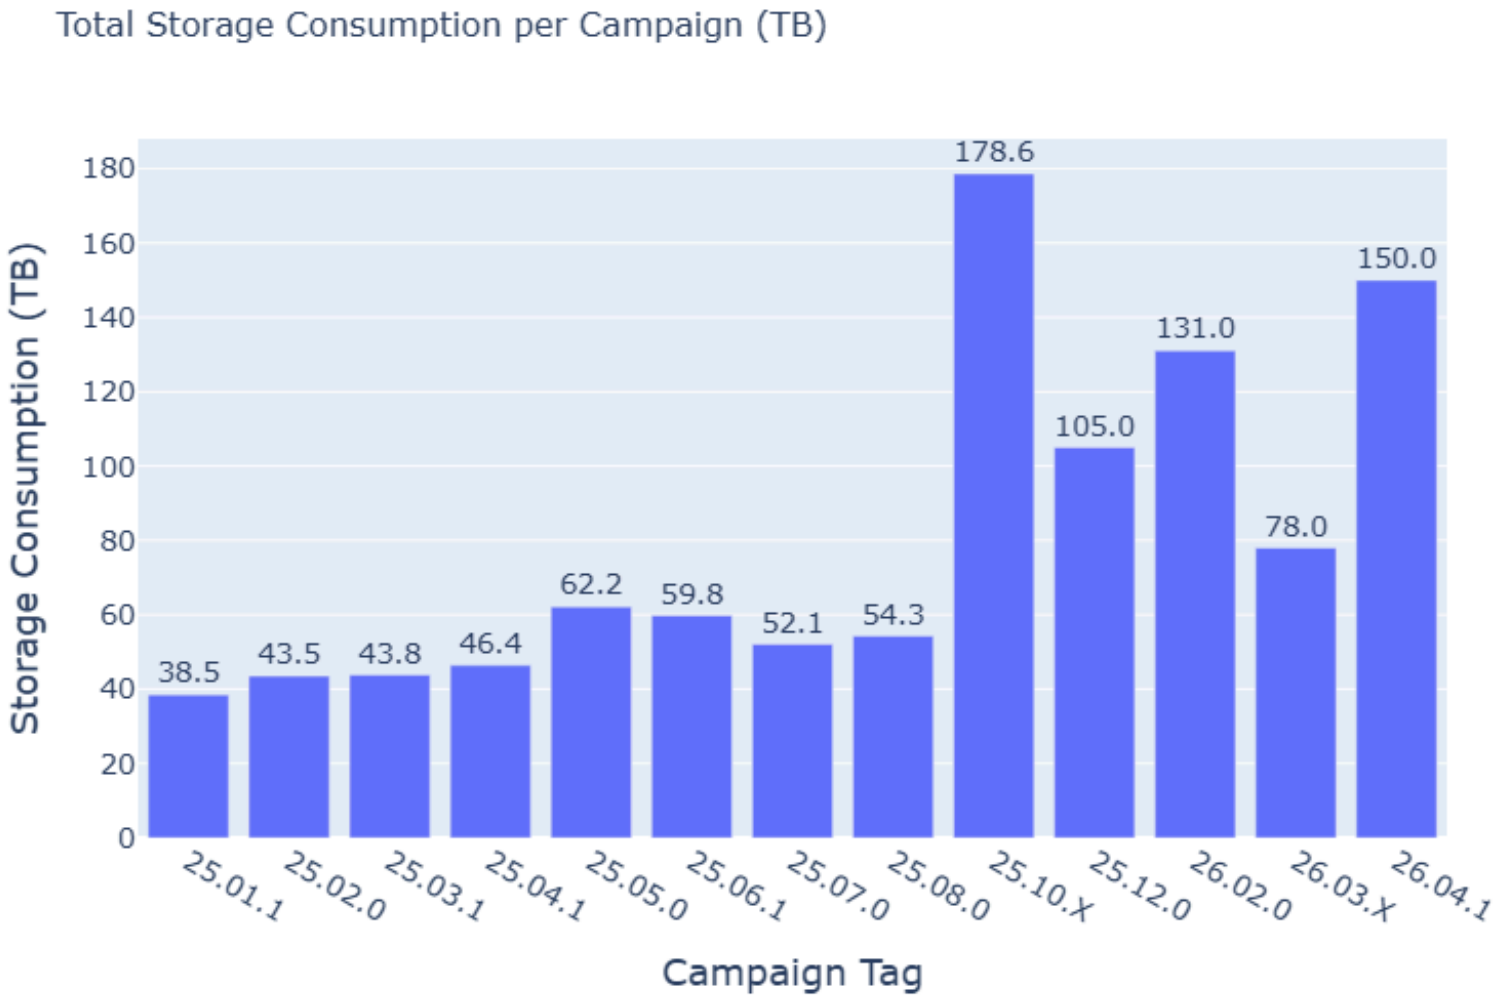

ePIC Storage Usage (2025–2026)

Simulation Production: Infrastructure Evolution

Workflow & Data Management Stack

PanDA WMS / iDDS

Automated production workflow and adaptive job brokerage

Rucio

Distributed data management and replication across storage elements

AI-Powered Monitoring and Self-Documentation

MCP server stack powering an AI agent for monitoring across systems

Physics Configuration System (PCS) ⚠ beta

Simulation task provenance and configuration record

Integrating International Resources

OSG Umbrella

INFN-T1, DRAC, and UManitoba HPCC onboarded as allocated resources under the OSG umbrella

Next Steps

Direct PanDA submission and Rucio replication between storage elements are being validated at UManitoba HPCC to develop a model for future sites.

ePIC Computing Model

Infrastructure planning guided by the collaboration's computing model (See Holly Szumila-Vance's talk)

Each component elaborated in follow-up slides →

PanDA WMS + iDDS: Improved Workflow Management

Started with lightweight shell scripts and job submission templates (CHEP 2024) — but have now reached a scale of complexity that necessitates a dedicated **Workflow Management System (WMS)**.

PanDA/iDDS meet our urgent needs:

- **PanDA WMS** — task brokerage, scalability, adaptive scheduling
- **iDDS** — data transformations and task dependency chains

Production orchestration will be consolidated around **BNL PanDA/iDDS** and **JLab Rucio** by **2026**. WMS options for streaming orchestration will be re-evaluated in **FY2028**.

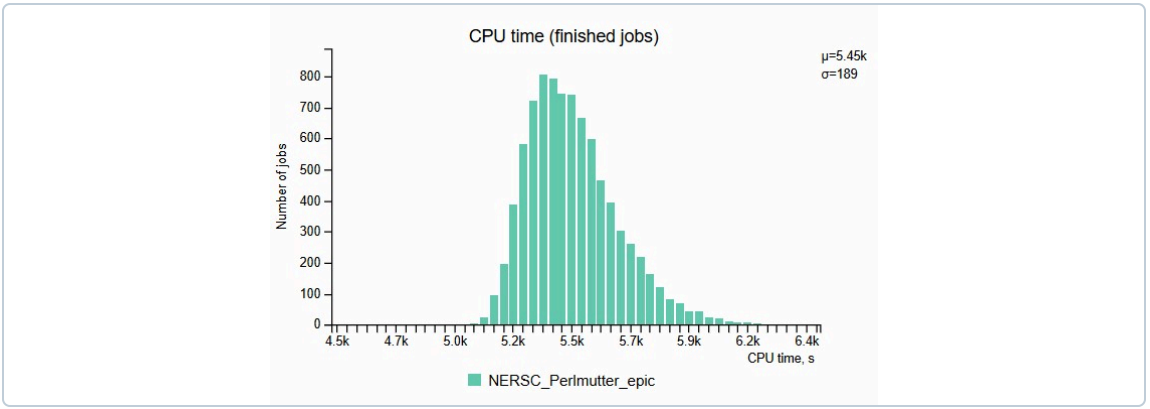

Example of PanDA Monitoring

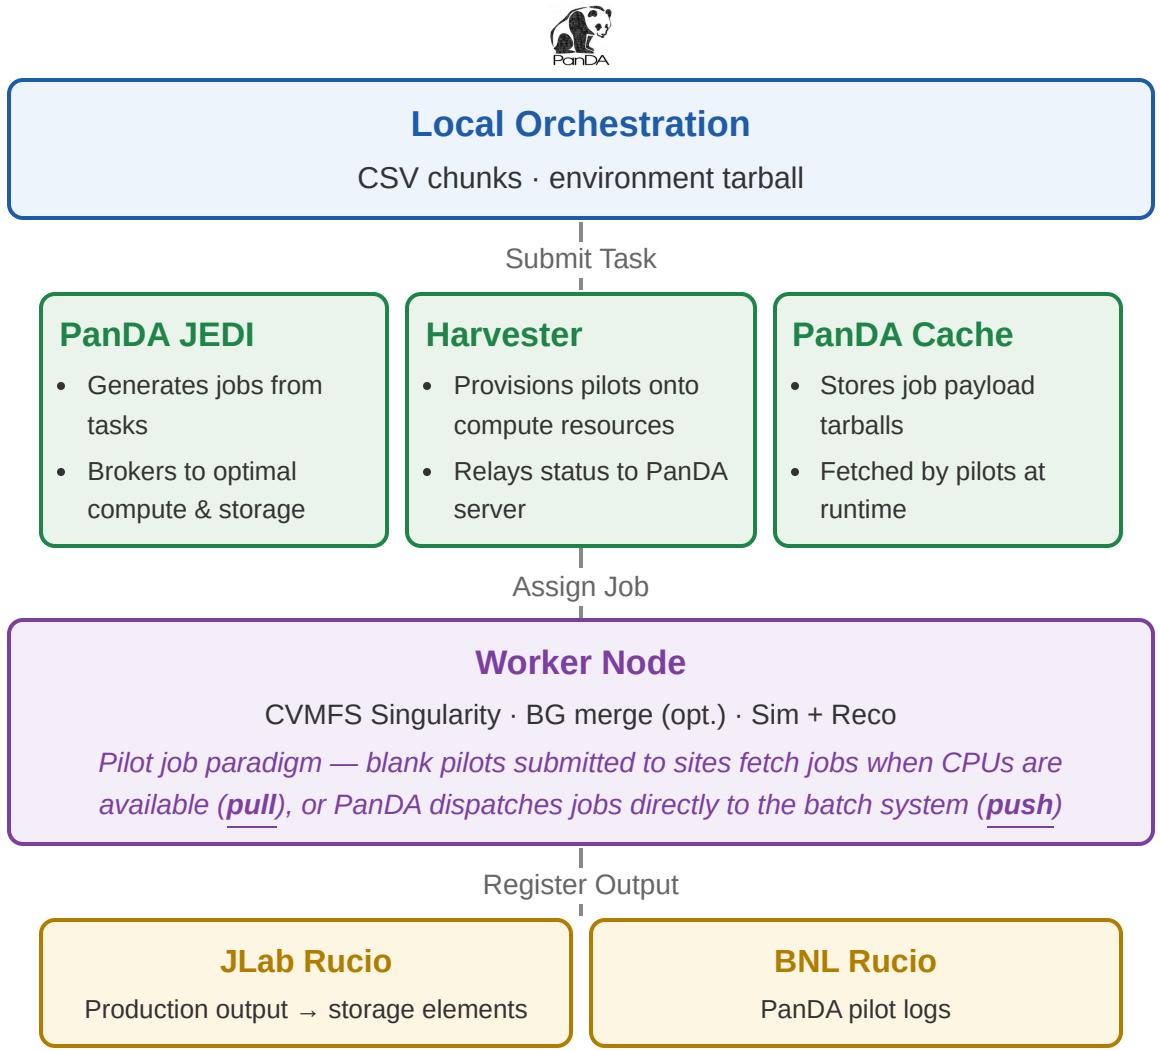

Current State of PanDA Integration for ePIC

Rucio for Data Management

Current State

- Server at JLab in production since Jan 2025
- File & dataset-level Data Identifiers
- FTS transfers via subscription rules
- X.509 auth for uploads; token-based auth with **CILogon** being developed

Naming Convention

`epic:/dataLevel/path/to/dataset/file`

`protocol://host:port/prefix/dataLevel/...`

Strategy: full filesystem path reflected in Data Identifier for human readability

Limitations:

- Rucio character limits constrain deeply nested paths
- PanDA interprets `/` as container boundary separator

Rucio Metadata Rollout

- Data made **Findable** and **Accessible** under **FAIR data principles**
  - Metadata tags (energy, event generator, physics working group, detector geometry, etc.) for precise discovery and filtering
  - Rucio CLI available inside the eic-shell container
- Rucio tools introduced to collaboration via user learning & training efforts (See Alexandr Prozorov's talk)
- Further reading: Rucio for ePIC — Rucio Workshop 2025
- Tutorial: EIC File Access

AI-Powered Monitoring and Self-Documentation

MCP Server Stack

Production Systems

`SWF-monitor` — PanDA jobs/tasks/queues, PCS tags, error summaries

`rucio-eic-mcp-server` — Data Identifiers, replicas, replication rules

Storage & Data

`xrootd-mcp-server` — browse & read files on EIC storage

`uproot-mcp-server` — inspect ROOT file structure, histogram branches

Code & Documentation

`lxr-mcp-server` — cross-reference EIC source code

`github-mcp-server` — issues, PRs, discussions

`zenodo-mcp-server` — search & download records

`corun-mcp-server` — generate ePIC documentation pages

Examples: AI Chatbot on Mattermost

> *"Summarize failed jobs on Perlmutter in the month of April and specify the most dominant error"*

PanDA bot BOT 10:46 AM

Failed jobs on Perlmutter (last 30 days): 78,194 total

Error Source	Error Code	Diagnostic	Count	% of Total
Pilot	1356	Pre-process command failed	69,632	39.2%
TaskBuffer	100	kill by Sakib Rahman	30,821	17.3%
Executor	5303	Failed to download tarball (various)	~59,560+	~33.5%
DDM	200	Data identifier not found	6,914	3.9%
Pilot	1305	Failed to execute payload (connection/validation)	6,454	3.6%

> *"List what's available under the /volatile/eic/EPIC/EVGEN directory on JLab XRootD"*

Name	Type	Size	Last Modified
BACKGROUNDS	directory	25.6 KB	2025-03-11
CI	directory	25.6 KB	2025-03-11
DDIS	directory	25.6 KB	2025-12-22
DIS	directory	25.6 KB	**2026-05-14**
Djangoh	directory	25.6 KB	2023-02-27

Physics Configuration System ⚠ beta

PCS follows a **user-centered design** — mapping collaboration requests to job settings, automating request tracking, and bridging the gap between physics groups and production infrastructure. Authoritative provenance and configuration record for simulation tasks, with direct PanDA job injection as the long-term goal.

Stores Per Task

Physics & EvGen tags	Simulation & reco tags
Input/output datasets	Workflow mode & stages
Validation status	Public catalogue entry
Downstream task IDs	Request state

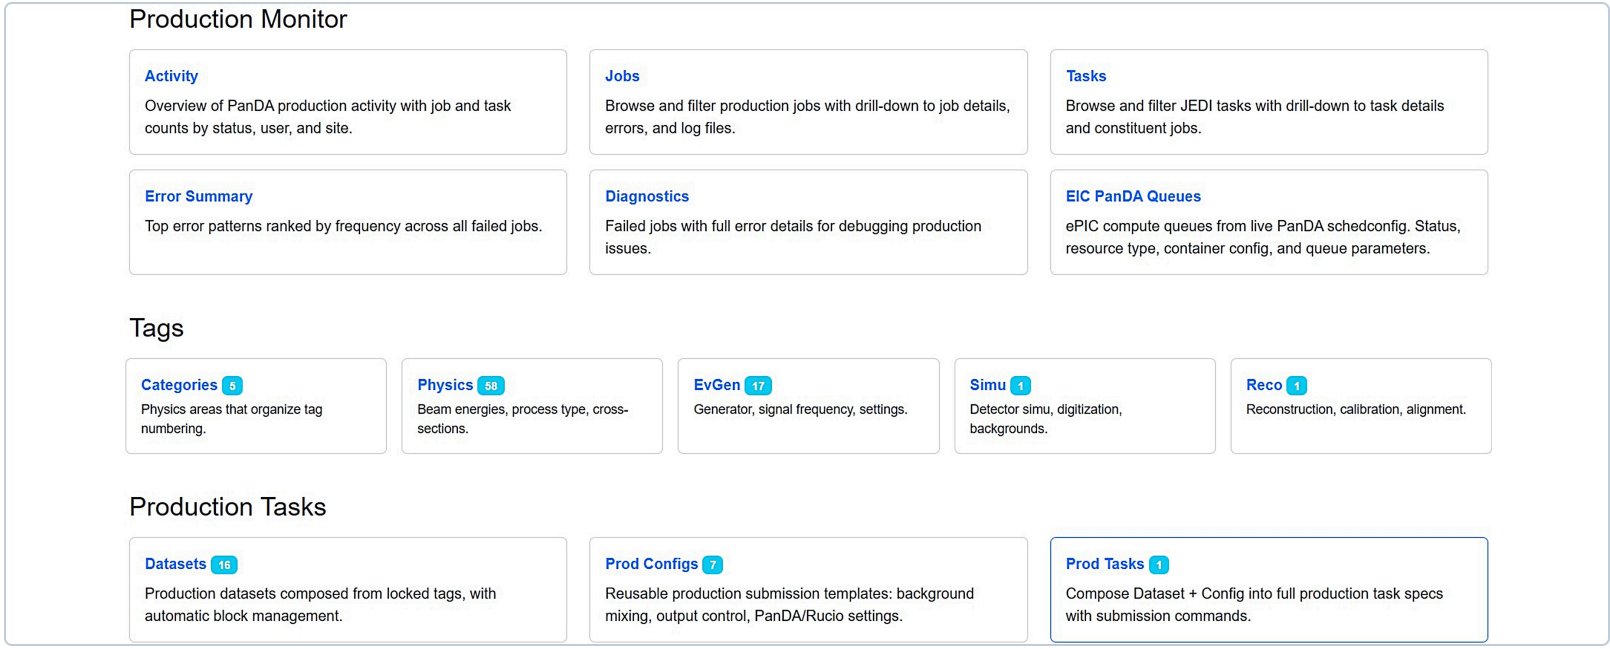

Preliminary IAM Model — ePIC Phonebook defines groups/roles → **COmanage** translates to registry groups → **CILogon** issues tokens (JWTs) → PCS enforces local community permissions.

PCS in Action: Example PanDA Task

> Dataset naming convention inspired by the **ATLAS Metadata Interface (AMI)** — Odier et al., EPJ Web Conf. 214 (2019) 05046

Dataset Naming Convention

```
group.EIC.26.02.0.epic_craterlake.p4017.e7.s1.r1
```

26.02.0	Campaign version
epic_craterlake	Detector geometry
p4017	Physics tag
e7	EvGen tag
s1	Simulation tag
r1	Reconstruction tag

PanDA Jargon

Task	Unit of work submitted to PanDA; maps to one PCS dataset
Job	Individual pilot execution within a task
Queue	Site-level resource endpoint for job brokerage
JUG_XL	ePIC software container image tag

group.EIC.26.02.0.epic_craterlake.p4017.e7.s1.r1

Dataset

Name	group.EIC.26.02.0.epic_craterlake.p4017.e7.s1.r1
Detector	26.02.0 / epic_craterlake

Type	Tag	Description	Parameters
Physics	**p4017**	Photoproduction J/psi 18x275 ep edecay	process=PHOTOPRODUCTION_JPSI, decay_mode=edecay, beam_species=ep, beam_energy_hadron=275, beam_energy_electron=18
EvGen	**e7**	eSTARlight 1.3.0-1.0, no backgrounds	generator=eSTARlight, generator_version=1.3.0-1.0
Simu	**s1**	npsim 26.02.0, standard filters	sim_version=26.02.0, detector_sim=npsim
Reco	**r1**	eicrecon 26.02.0, standard	reco_config=standard, reco_version=26.02.0

Config: 26.03.0 Standard Production

JUG_XL	26.03.0-stable
Container	/cvmfs/singularity.opensciencegrid.org/eicweb/eic_xl:26.03.0-stable
BG Mixing	No
Copy	Reco, Log
Rucio	Yes (EIC-XRD)
Hours/Job	2.00
Events/Task	200000

Simulation Payload Optimizations: Background Merging

Every reconstructed event is an overlay of signal and multiple background sources. The slowest detector subsystem has a **2 μs readout window**:

Background inputs at 10×275 GeV

Process	Frequency	Status
Synchrotron radiation (synrad)	**13.277 GHz**	2000
Electron beam gas	3.177 MHz	3000
Electron Coulomb	29.56 kHz	4000
Electron Touschek	233.5 kHz	5000
Proton beam gas	32.6 kHz	6000
Signal (e.g. DIS)	**500 kHz**	-

Re-simulating background signals per physics sample is prohibitively expensive.

Solution: simulate background pools once and merge statistically at the 2 μs frame level — one synrad campaign reused across **all physics analyses**. Processes tagged via `MCParticles.generatorStatus` for post-hoc separation.

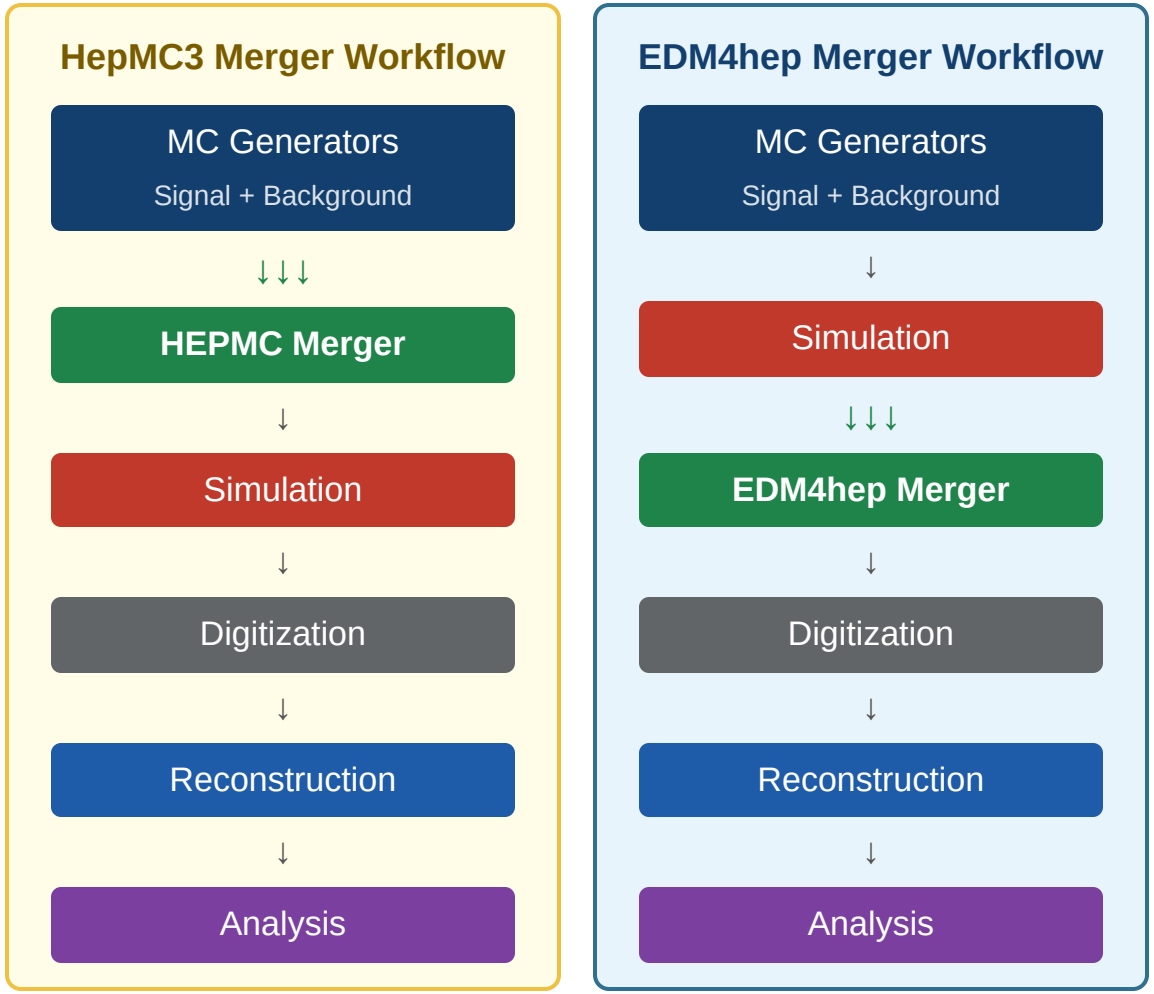

Merge before simulation (HepMC3) vs. after simulation (EDM4hep)

Background mixed events are highly relevant to streaming reconstruction algorithm development (See Takuya Kumaoka's talk)

Summary: Scaling ePIC Simulation Production

Experiment and Collaboration

- Primary detector at the EIC, realized by the ePIC Collaboration in partnership with the EIC Project, BNL and JLab
- Worldwide footprint: more than 1100 collaborators across 181 institutions in 25 countries

Infrastructure and Resource Evolution

- Monthly simulation campaigns run on HTC and HPC resources in preparation for the TDR
- Opportunistic sites via OSG and allocated domestic & international sites
- Storage is scaling with growing collaboration demand

Workload and Data Management

- PanDA WMS / iDDS enables automated production workflows and job brokerage
- Rucio enables distributed data management, replication, and user data access
- PCS maps user requests from PWGs and DSCs to production tasks

Payload Optimizations

- Optimization of expensive background mixed simulations

> ePIC builds on Grid computing experience and lessons learned from the NP & HEP community and contributes back to the community with our development efforts.

Related ePIC Talks at CHEP 2026

The ePIC Streaming Computing Model
Holly Szumila-Vance

Exploring Potential Pathways to Accelerate ePIC Detector Simulation
Sakib Rahman

Development of Streaming Data Reconstruction for ePIC Experiment at EIC
Takuya Kumaoka

ePIC User Learning Training and Documentation Strategies
Alexandr Prozorov