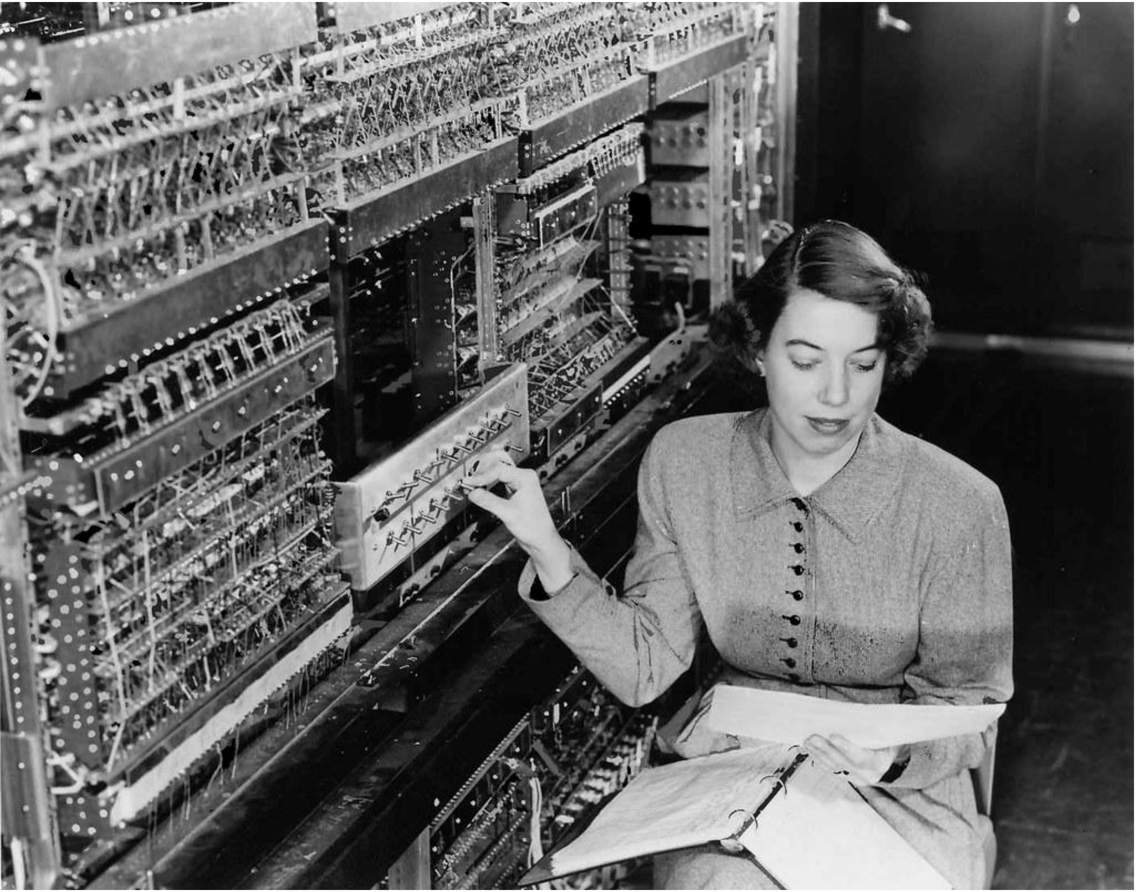

ATHENA

Computing Coordination Meeting

Monday 2021-10-18

The Software and Computing WG Conveners:
Andrea Bressan (University of Trieste and INFN) ,
Dmitry Romanov (Jefferson lab) ,
Sylvester Joosten (Argonne National Laboratory) ,
Whitney Armstrong (Argonne National Laboratory) ,
Wouter Deconinck (The University of Manitoba)

Philosophy*:* Let's prepare for our future at the EIC!

ATHENA

- **Build forward-looking team of developers to ensure the long-term success of the EIC scientific program in software & computing.**
- Focus on modern scientific computing practices
    - Strong emphasis on modular orthogonal tools.
    - Integration with HTC/HPC, CI workflows, and enable use of data-science toolkits.
- Avoid "not-invented-here" syndrome, and instead leverage cutting-edge CERN-supported software components where possible.
    - Build on top of mature, well-supported, and actively developed software stack.
    - Externalize support burden where possible.
- Actively work with the EICUG SWG to help develop and integrate community tools for all collaboration.

Software Stack In A Nutshell

DD4hep: TGeo and Geant4 geometry definition, detector plugin library, wrappers to run Geant4.

Juggler: Digitization and reconstruction algorithms (based on Gaudi with Podio-based data model and genfit/ACTS for tracking).

Gaudi: Generic open project for building auto-scaling event processing frameworks. Enables task-based concurrent execution in a heterogeneous computing environment. Encourages efficient coding practices.

ACTS: Experiment-independent tracking toolkit (geometry constructed from DD4hep via plugin).

Podio: Robust data model definition to cross the boundaries between tools, independent of file format.

MCEG

DD4Hep,
Geometry,
Geant4 (DDSim)

Juggler,
Gaudi - processing
ACTS - tracking,
Tensorflow - ML
Custom reco algorithms

Analysis/Benchmarks

ATHENA

ATHENA Software Ecosystem: Emphasis On Modularity

ATHENA

DD4hep

DD4hep (CLiC, SLD): define active and passive geometry, readout segmentation

Simulation in npsim (DDG4) provides hits for Gaudi

Gaudi (originally LHCb): connect geometry and data to algorithms

S3 XRootD

Flat output files (podio) enable flexible workflows (ROOT, python,…) without additional dependencies, available worldwide (S3, xrootd)

G4

G4 backend (can be swapped out, e.g., celeritas, opticks,…)

DD4hep geometry service in Juggler

ACTS has built-in support for DD4hep geometries

PyTorch TensorFlow

Juggler: I/O, digitization, reconstruction, clustering and tracking algorithms

Call ACTS for tracking & vertexing, providing hits and detector geometries

acts

ACTS (ATLAS): provide general track reconstruction algorithms

The DD4hep community

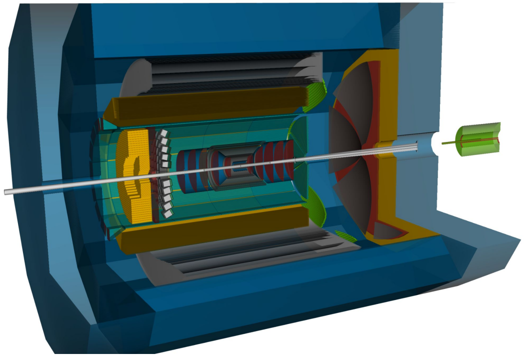

Super Charm Tau Factories

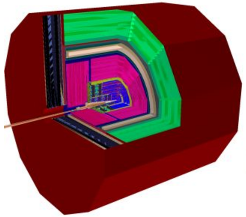

Production

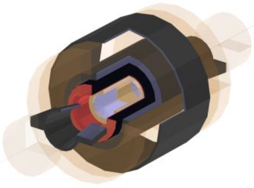

Production

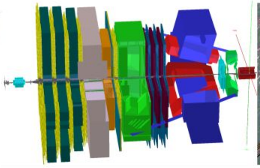

Production
Run 3

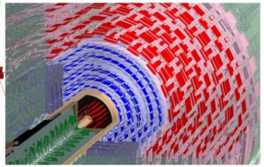

Under investigation
Run 3

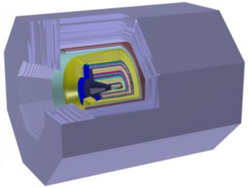

Production

“framework for providing a complete solution for full detector description (geometry, materials, visualization, readout, alignment, calibration, etc.)”

The ACTS community

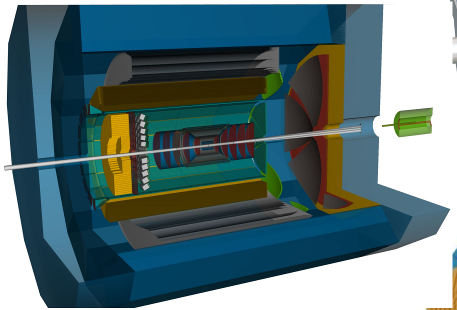

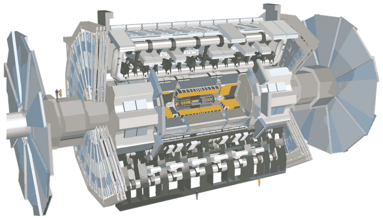

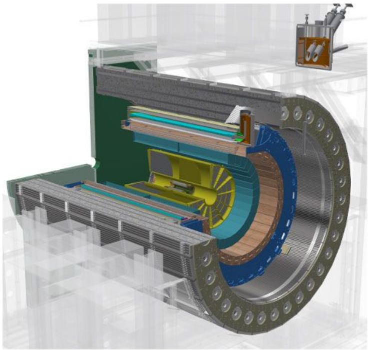

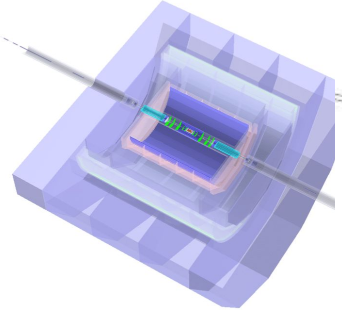

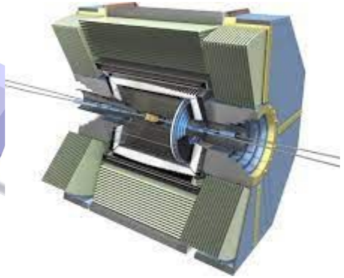

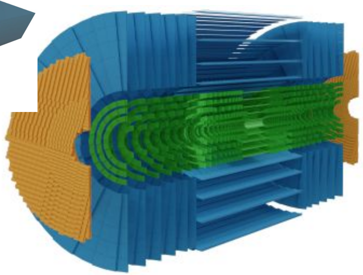

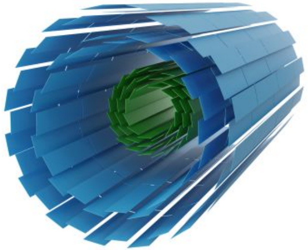

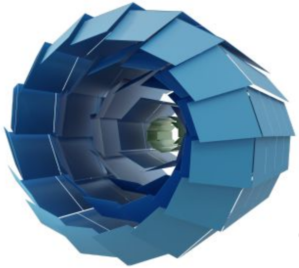

Automated Workflows at eicweb

GitLab server (eicweb.phy.anl.gov)

- mirrored on github.com/eic
- continuous integration
- dedicated build cluster

Runs automatically on each user commit, executing workflows running multiple tests, benchmarks and analysis

Automated containers

Both Docker and Singularity images are created nightly or on demand (commit) providing:

- reproducibility,
- production level images
- latest updates for those working locally

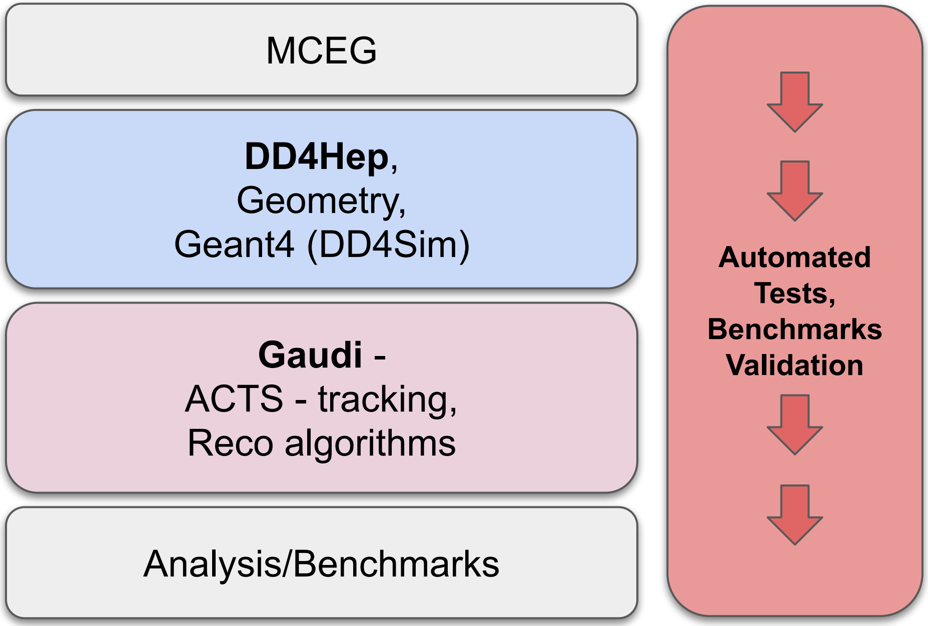

Why use a custom GitLab server?

- CI system loses benefits with shared, queued jobs
- Our collaboration controls users, yet access to HPC

Automated Workflows at eicweb

High Level Requirements

- Failures should show up in minutes, but benchmarking may take longer.
- Merge request cannot be held up by checks for more than ~two hours.
- Modular framework with large dependencies (root, geant4, ACTS) results in large containers, which must be distributed efficiently to CI nodes.
- Access for new users at institutions without MOU/NPUA (e.g. EIC India) without a lengthy approval process.
- Compliance with data management, export controls (private repositories).

Infrastructure at ANL

- Dedicated servers with 384 job slots to run ~10-minute to 3-hour long CI pipelines of ~50 to ~100 individual steps for commits and merge requests. (Github free tier: 2k mins / month)
- Docker build server of 128 cores and high bandwidth to distribute to nodes with modified gitlab-runners running unprivileged singularity containers.
- Integrated development environment with kubernetes for container shells.
- Full control over users accounts, no lab account required.

This infrastructure was not available on short notice at either BNL or JLab.

Automated Workflows: Local Environment Access

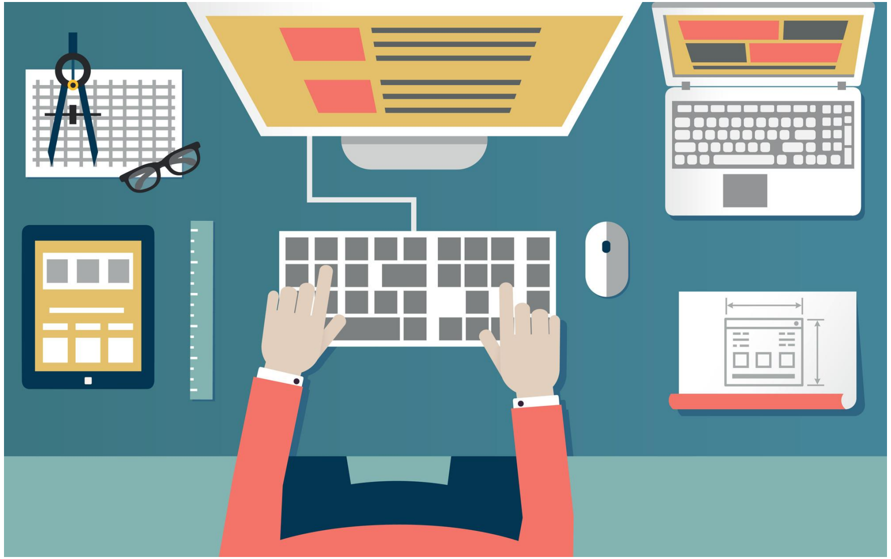

`curl -L get.athena-eic.org | bash`

- Uses images on `/cvmfs` when available, downloads singularity sifs otherwise.
- Basis of scalable computing on OSG: same containers are used everywhere.

Benchmarks, Documentation, Containerization

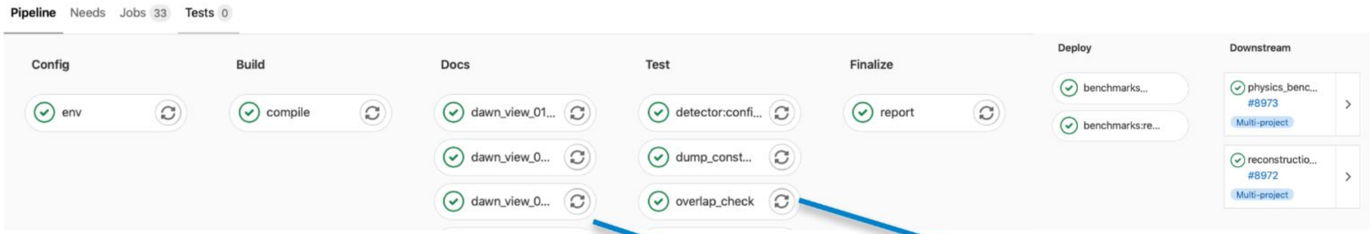

Geometry overlap checks running as part of every merge request

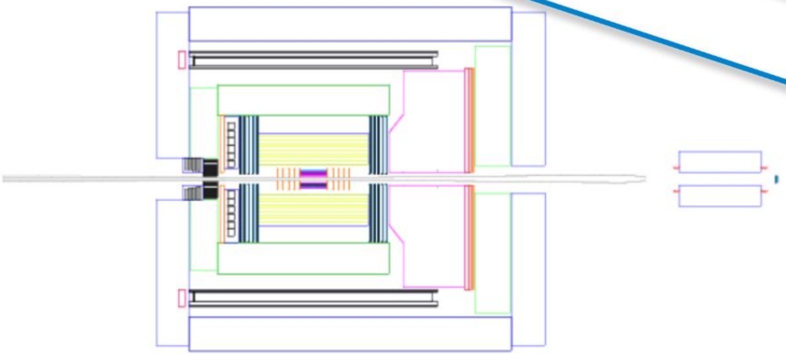

Automatic visualizations for detector geometries, saved as job artifacts (browsable!)

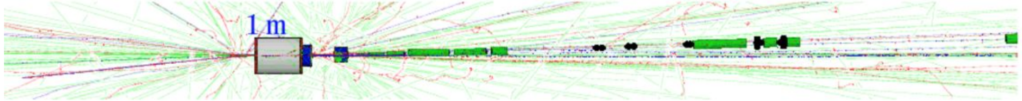

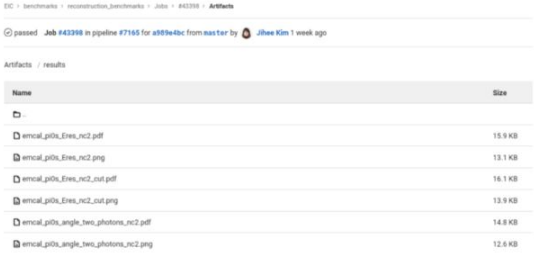

Automatic Visualization

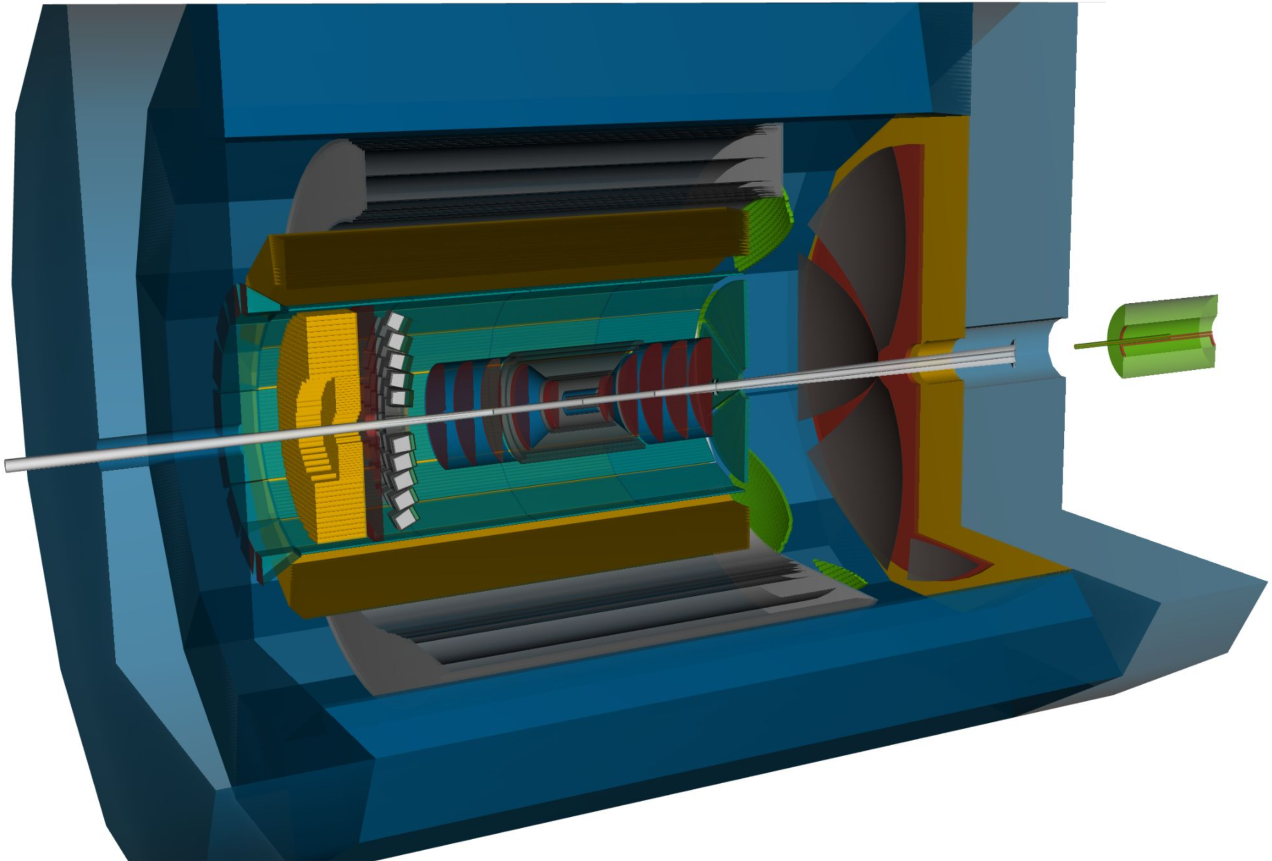

Automatic Visualization

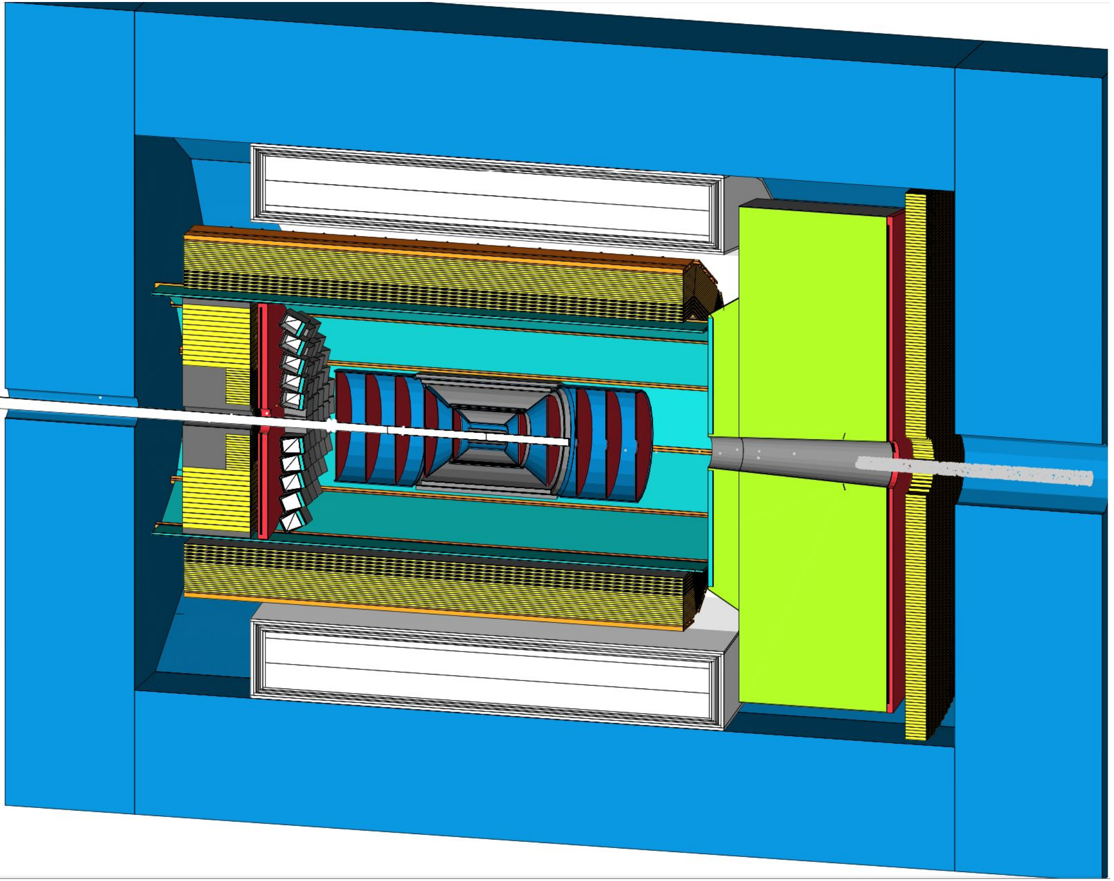

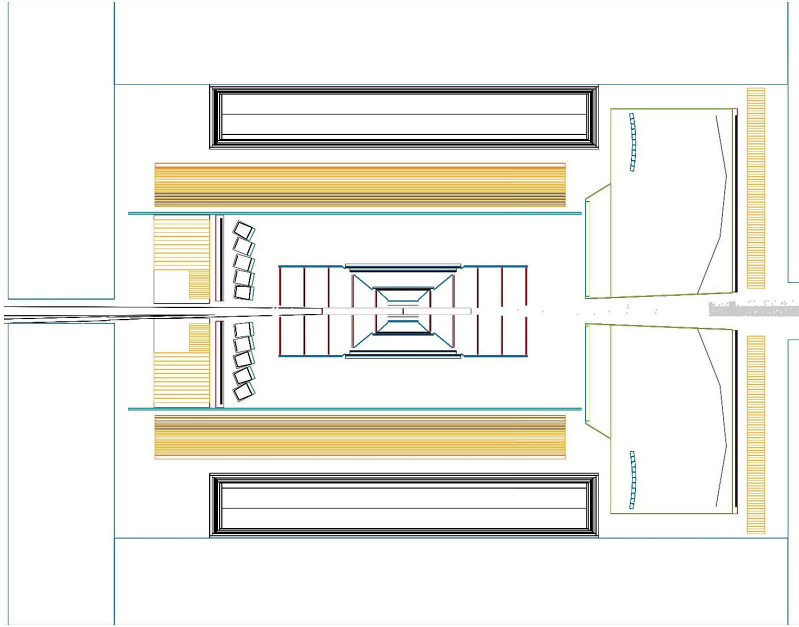

Development Boards on eicweb

- https://eicweb.phy.anl.gov/groups/EIC/-/boards
- Working to polish/integrate task list to make it easier for people to find/check out a task

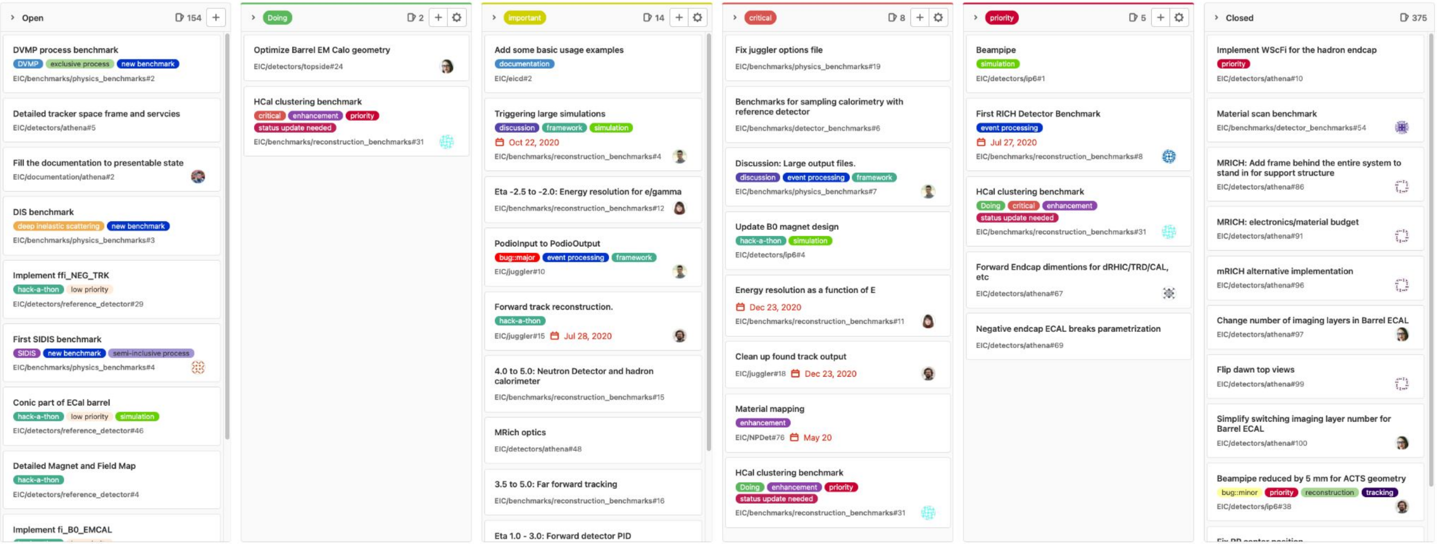

Detailed geometry implementation

EXAMPLE: Tracking Systems

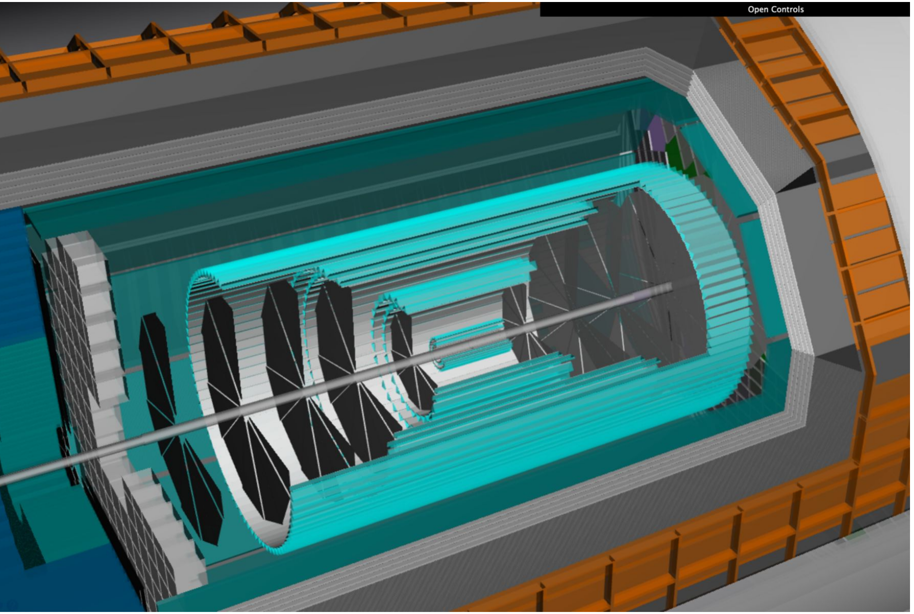

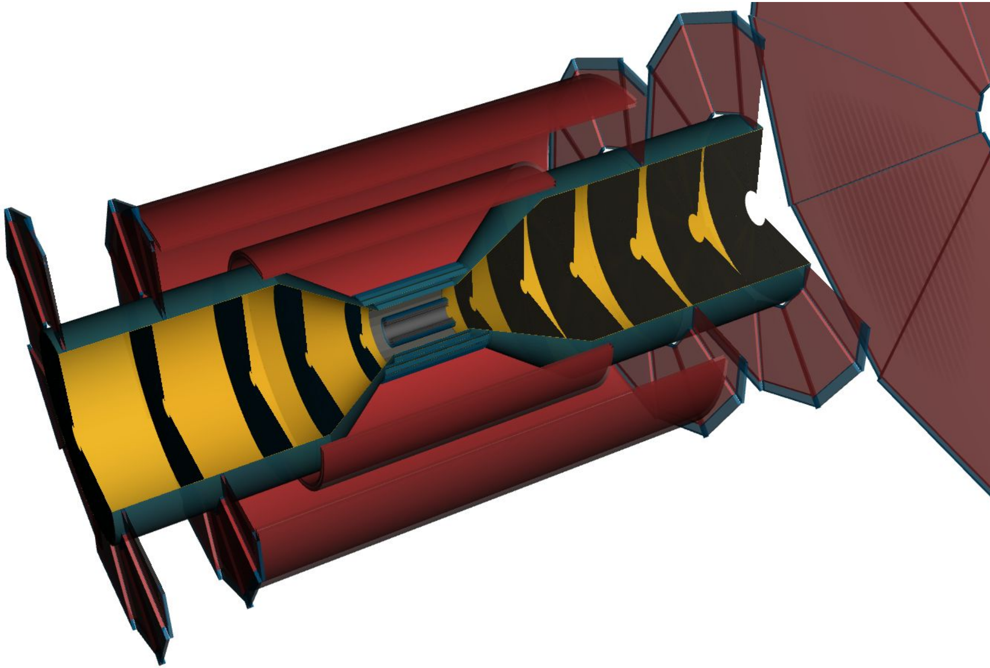

Computing Philosophy

Encourage Upstream Contributions

- Requirements of well-formed HepMC as input has resulted in real improvements to multiple MCEGs used by EIC community.
- Various upstream contributions to DD4hep, ACTS, Spack, uproot,...

Encourage Social Coding

- CI platform provides the incentive for developers to commit code frequently: achieving data management and analysis preservation goals.
- Merge request reviews ensure higher quality code and build developer skills.

Enable Access Without Restrictions

- ATHENA collaboration members include INFN, EIC India, LBL, ANL,...
- Data 'publicly' available through BNL S3 and publicly available through JLab xrootd.
- Flat data structures (i.e. could be a csv), stored as ubiquitous ROOT trees without need for data structure libraries.
- Support for uproot using numpy library (awkward not needed).

Under Evaluation

- Rucio for data management
- Reana for analysis workflows

Anatomy Of ATHENA Jobs

1. HepMC3 files generated by physics working groups (or internal single particle generator):
   a. Pythia8: full beam crossing angle and divergence effects included in hepmc3
   b. Other generators: hepmc3-to-hepmc3 afterburner to boost head-on collisions (independent of ATHENA: contribution to the EIC community)
2. CI analysis of HepMC3 files (produces csv artifacts):
   a. Running test, smoke tests, sanity checks, time-per-event determination
3. Job submission (identical syntax for slurm and condor systems)
   a. Automatic retrieval of csv artifacts, automatic job strategy determination
   b. No user code is needed: all submission support is available on `/cvmfs`
4. Job progression on node (entirely inside on container on `/cvmfs`)
   a. S3 download of HepMC3 file
   b. Run full simulation (downloads additional artifacts as needed)
   c. S3 upload of full simulation podio output
   d. Several reconstruction strategies (downloads additional artifacts as needed)
   e. S3 upload of reconstruction podio outputs
5. Mirroring from S3 to xrootd

Operational Benefits of OSG Jobs

Running at JLab (capacity 25k job slots, 14% for EIC)

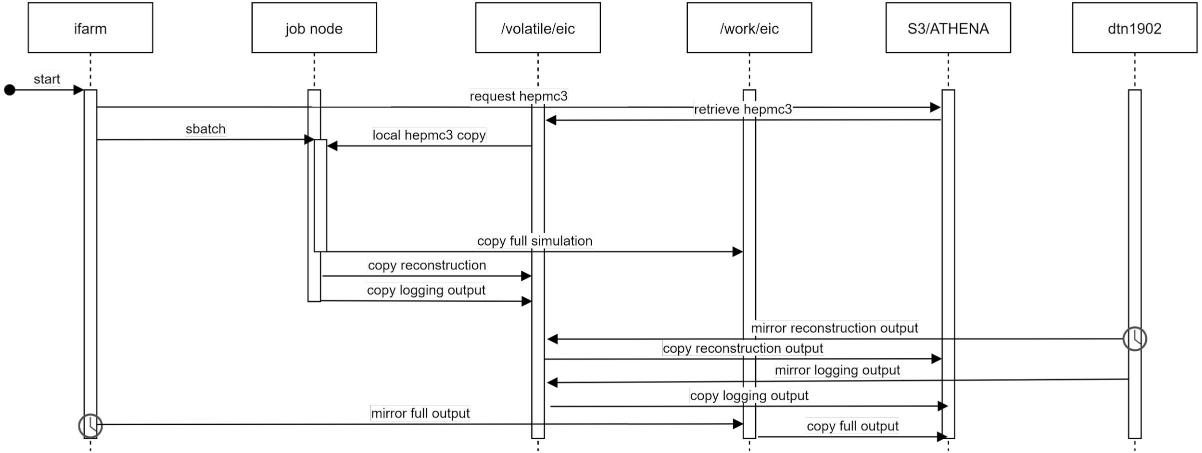

Now 500 TB each on /work/eic{2,3}; larger EIC xrootd service

Running on OSG (capacity 50k)

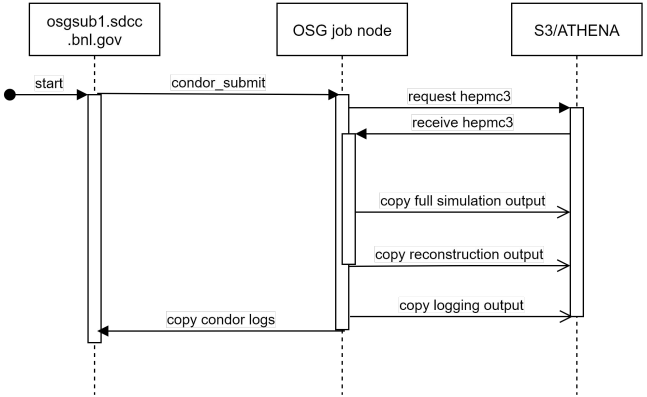

Also mirroring S3 to xrootd at JLab

Questions?