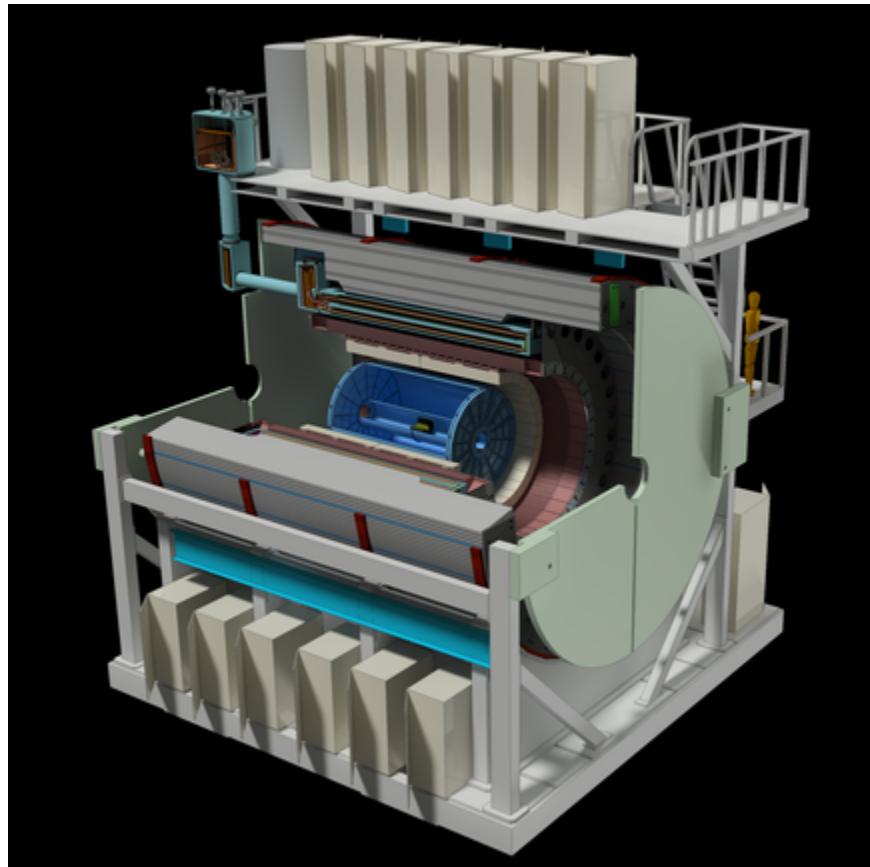
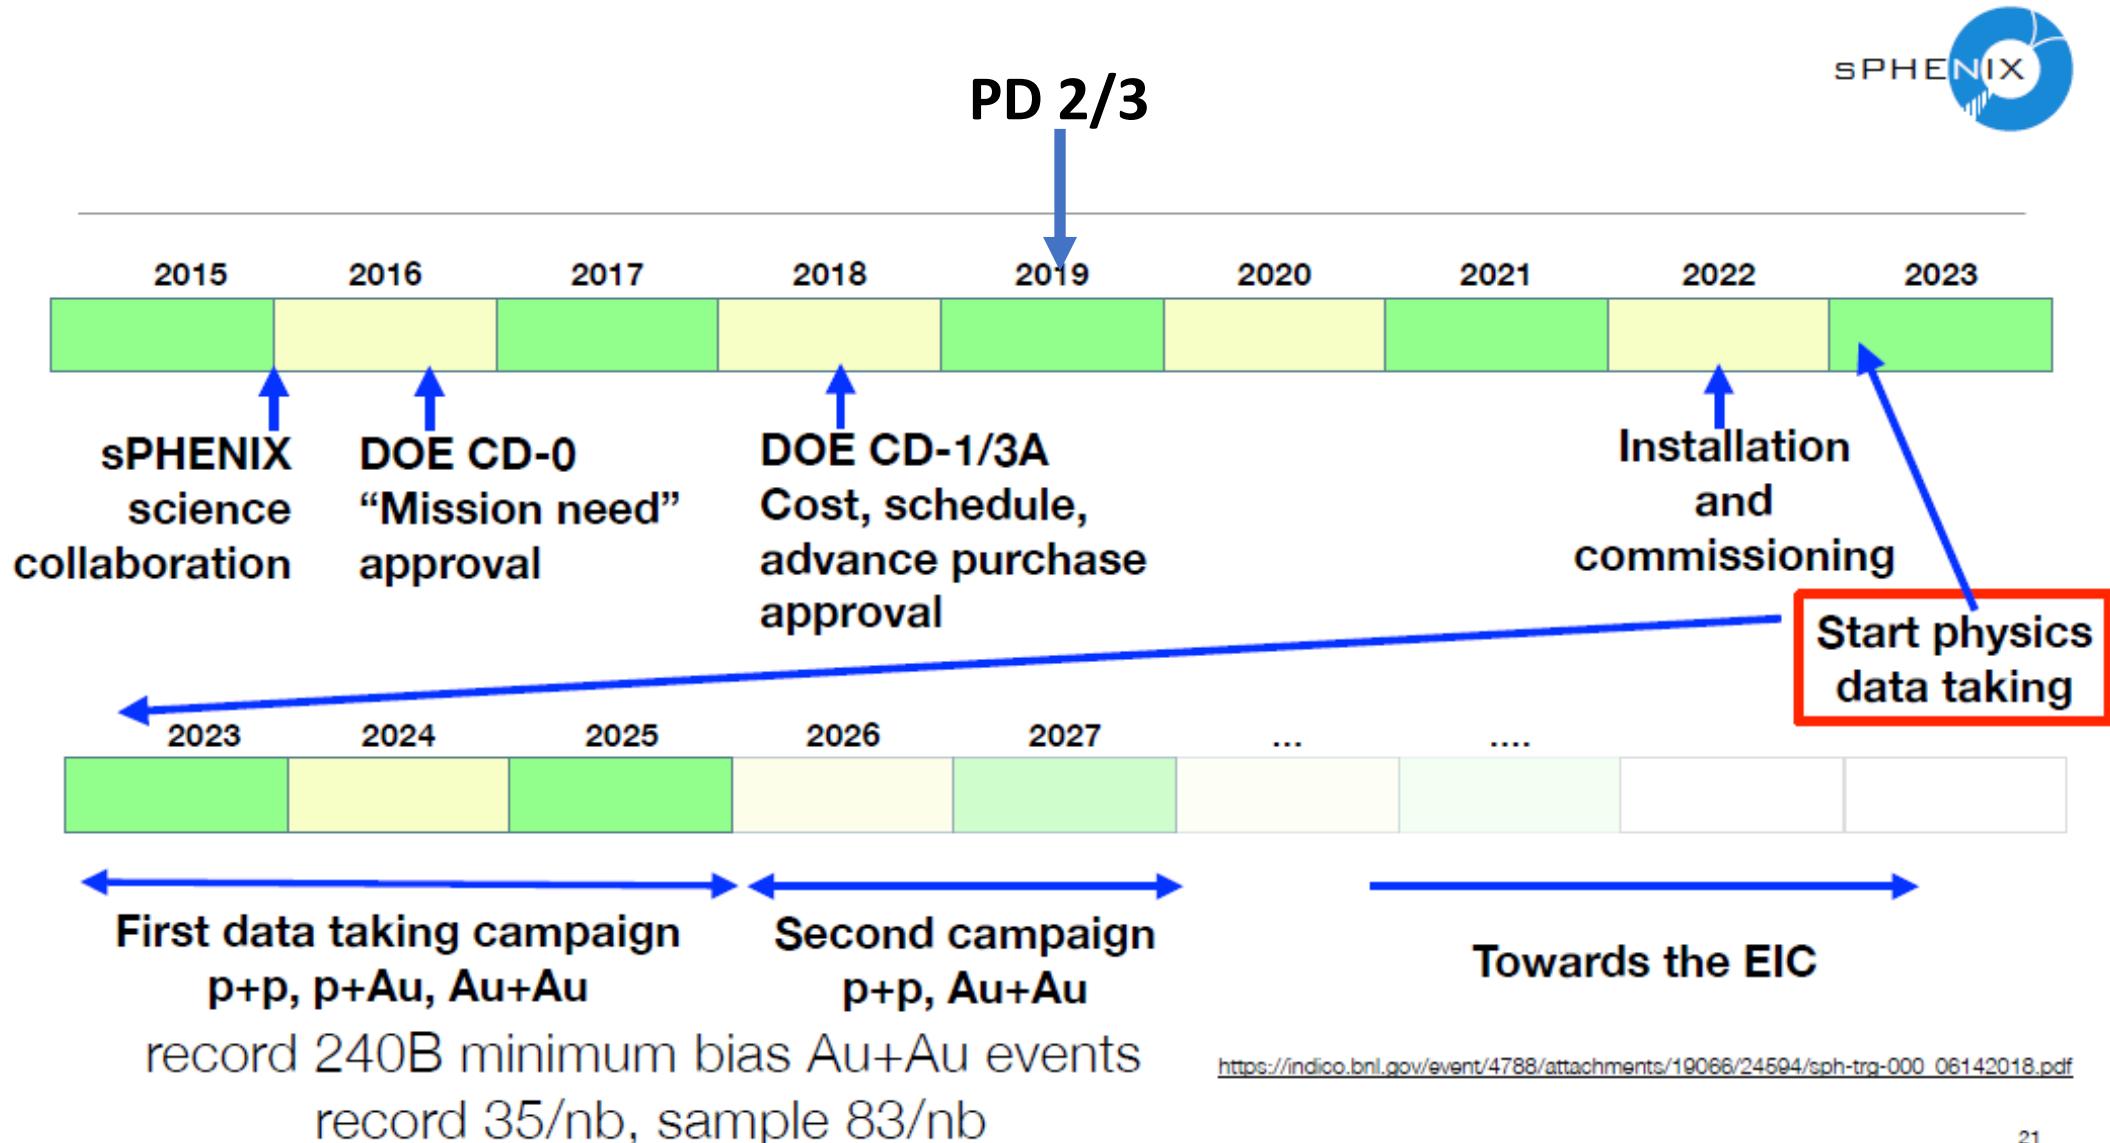


sPHENIX computing



sPHENIX timeline



1st sPHENIX workfest, 2011 in Boulder



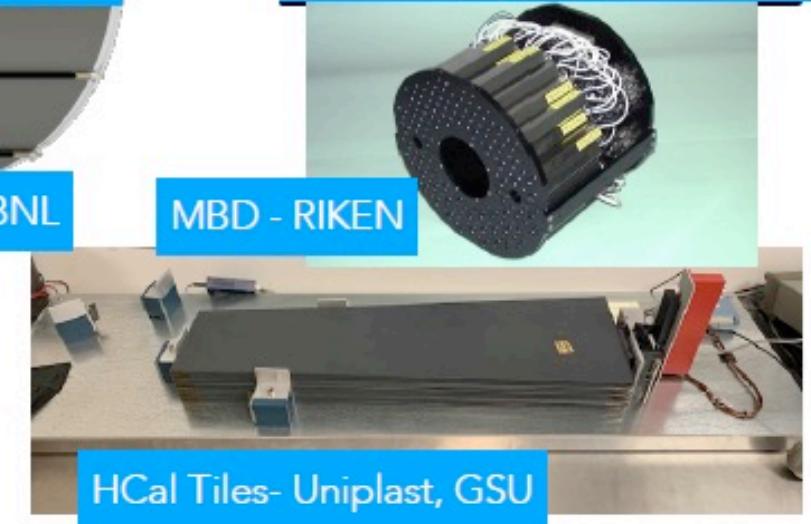
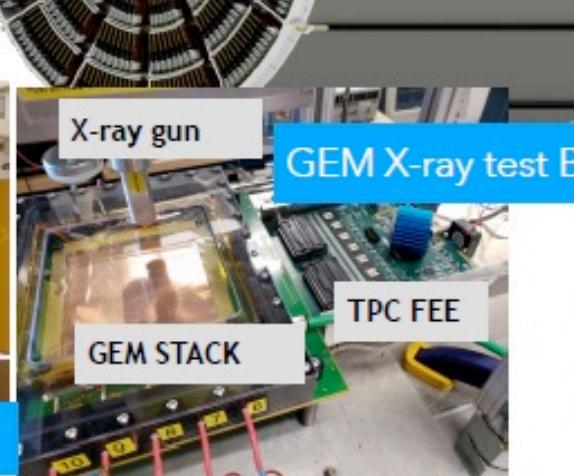
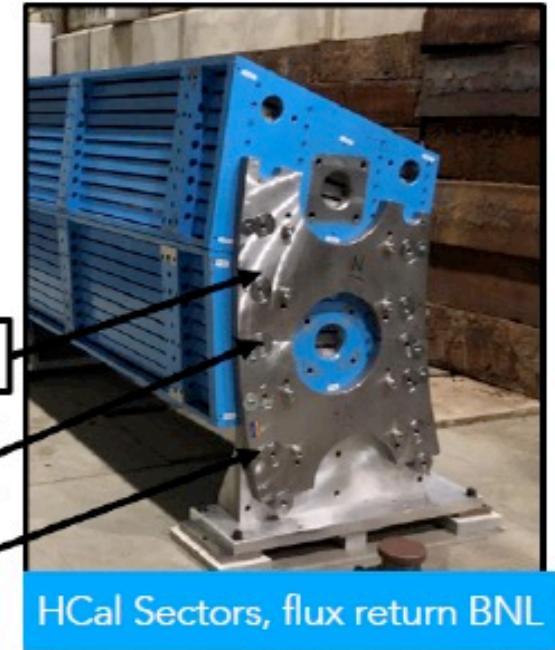
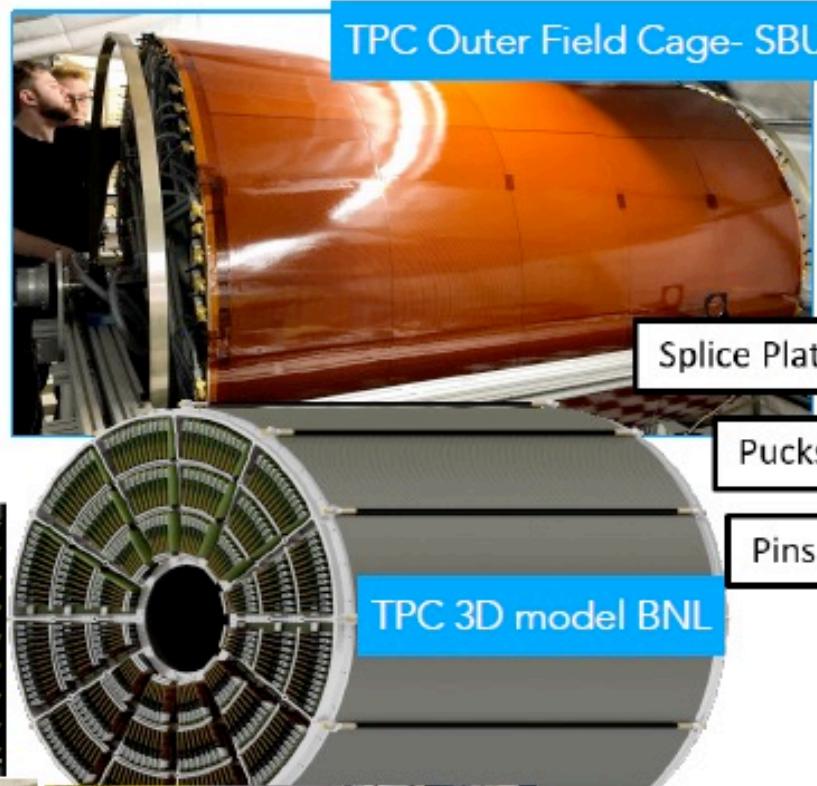
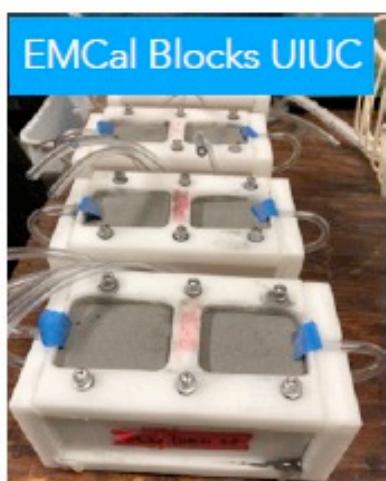
Computing corner:
Adding G4 to PHENIX software

A week later

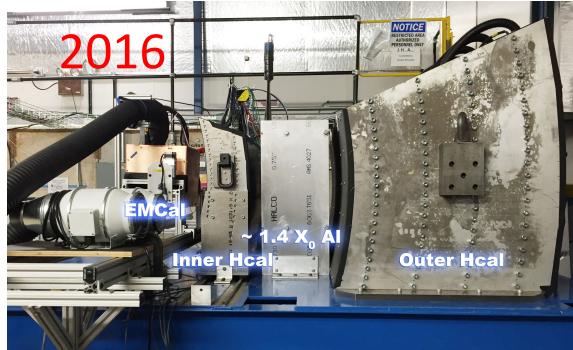
The first sPHENIX event display
(of a pythia event)



Fast forward to 2019, it is really happening



Test Beam, Test Beam and more Test Beam



Not only simulations,
We also get exposed to real data written in our envisioned daq format

Executive Summary Framework

Fun4All was developed for data reconstruction by people who used it to analyze data

Data was coming in already

Had to combine a zoo of subsystems each with their own ideas how to do things

No computing group, manpower came from the collaboration

It has been in production since 2003 (constantly updated for new needs)

Used for raw data reconstruction, analysis, simulations (G3 runs separately), embedding

Processed PB of data with many B of events

sPHENIX branched off in April 2015

Major revamping – made large effort to fix all those lessons learned things

Code checkers: cppcheck, scan, clang, coverity, valgrind, insure, 100% reproducible running

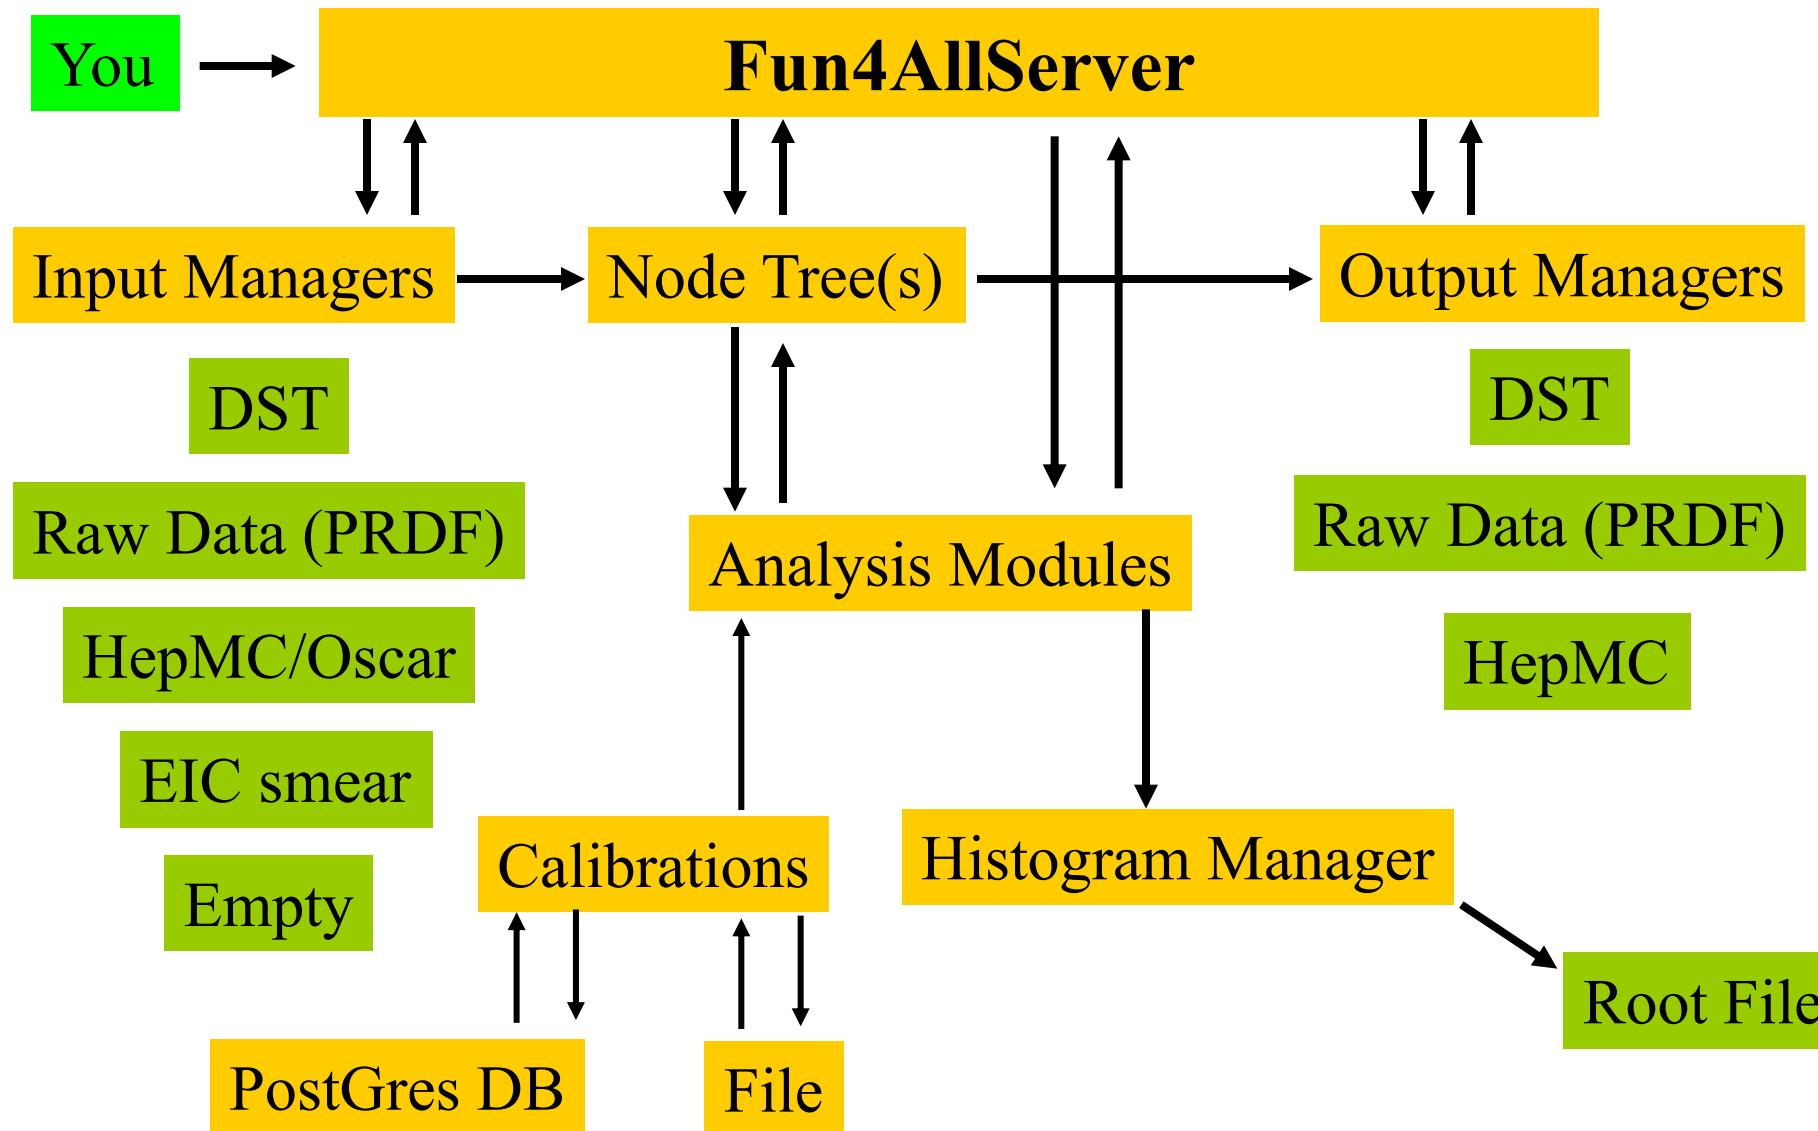
Good – PHENIX subsystem zoo reduced to basically 2 types:

Calorimeters (means need for clustering, use PHENIX clusterizer)

Inner Tracking silicon+tpc (means need for tracking, based on genfit)

Fermilab E1039 (Seaquest) branched off in 2017

Structure of our framework Fun4All



That's all there is to it, no backdoor communications – steered by ROOT macros

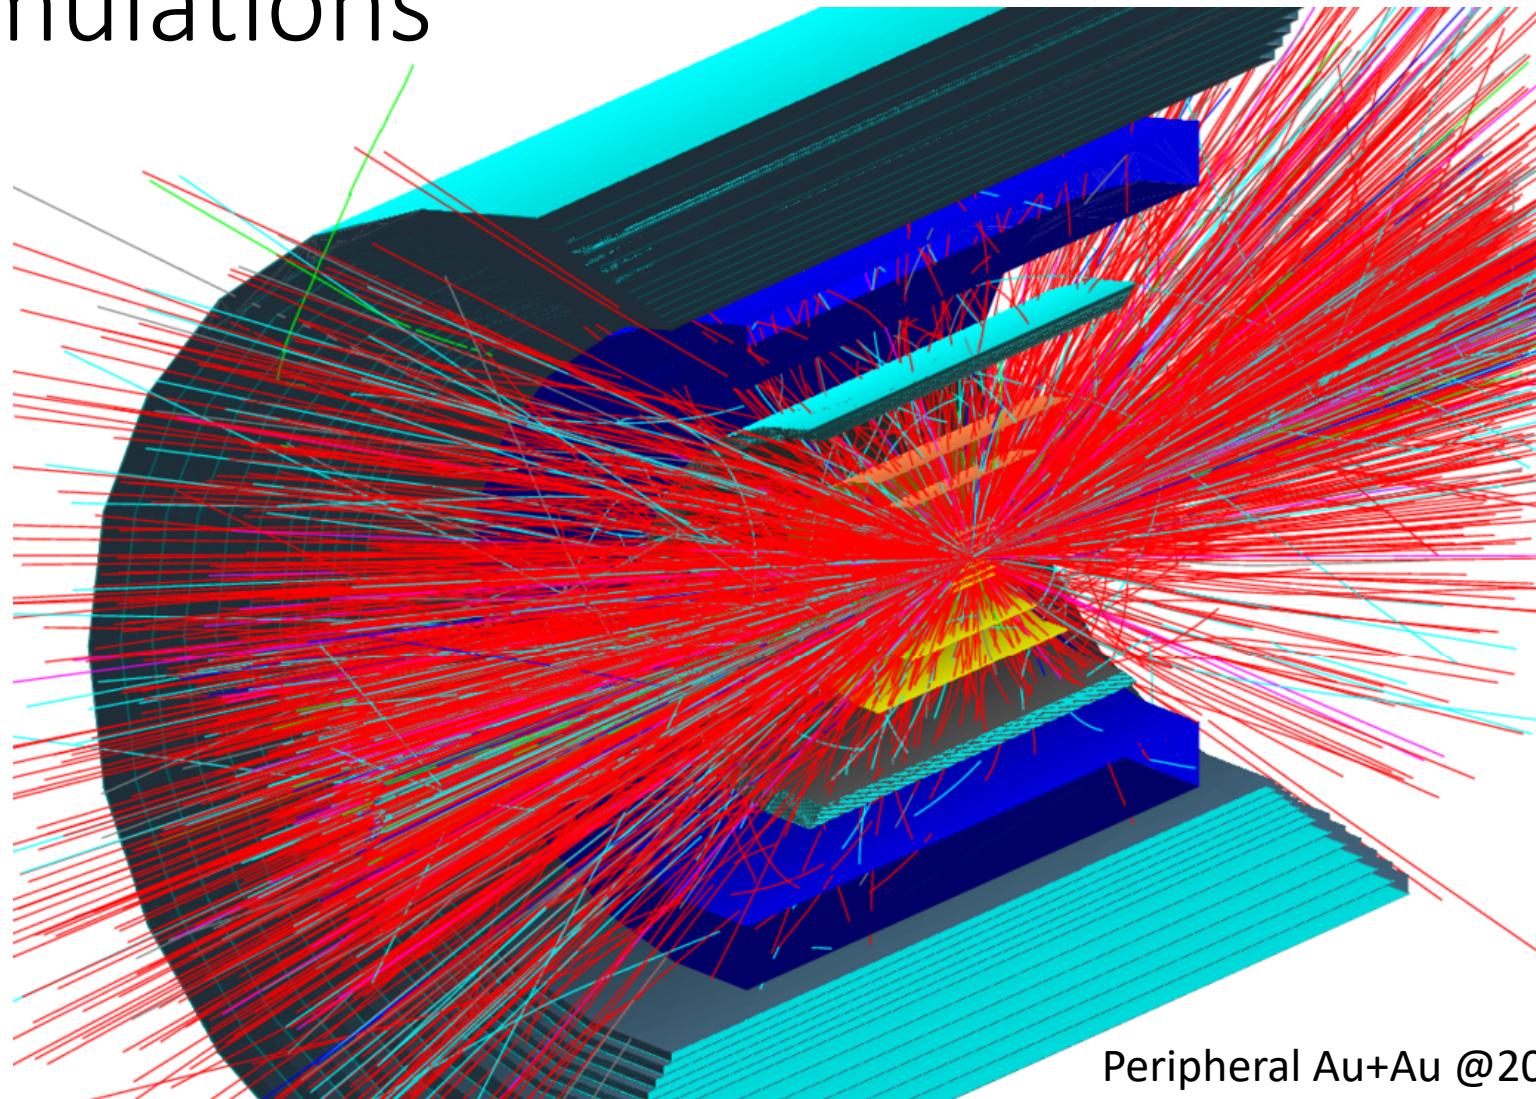
Keep it simple – Analysis Module Baseclass

You need to inherit from the SubsysReco Baseclass (offline/framework/fun4all/SubsysReco.h) which gives the methods which are called by Fun4All. If you don't implement all of them it's perfectly fine (the beauty of base classes)

- `Init(PHCompositeNode *topNode)`: called once when you register the module with the Fun4AllServer
- `InitRun(PHCompositeNode *topNode)`: called before the first event is analyzed and whenever data from a new run is encountered
- `process_event (PHCompositeNode *topNode)`: called for every event
- `ResetEvent(PHCompositeNode *topNode)`: called after each event is processed so you can clean up leftovers of this event in your code
- `EndRun(const int runnumber)`: called before the `InitRun` is called (caveat the Node tree already contains the data from the first event of the new run)
- `End(PHCompositeNode *topNode)`: Last call before we quit

If you create another node tree you can tell Fun4All to call your module with the respective `topNode` when you register your module

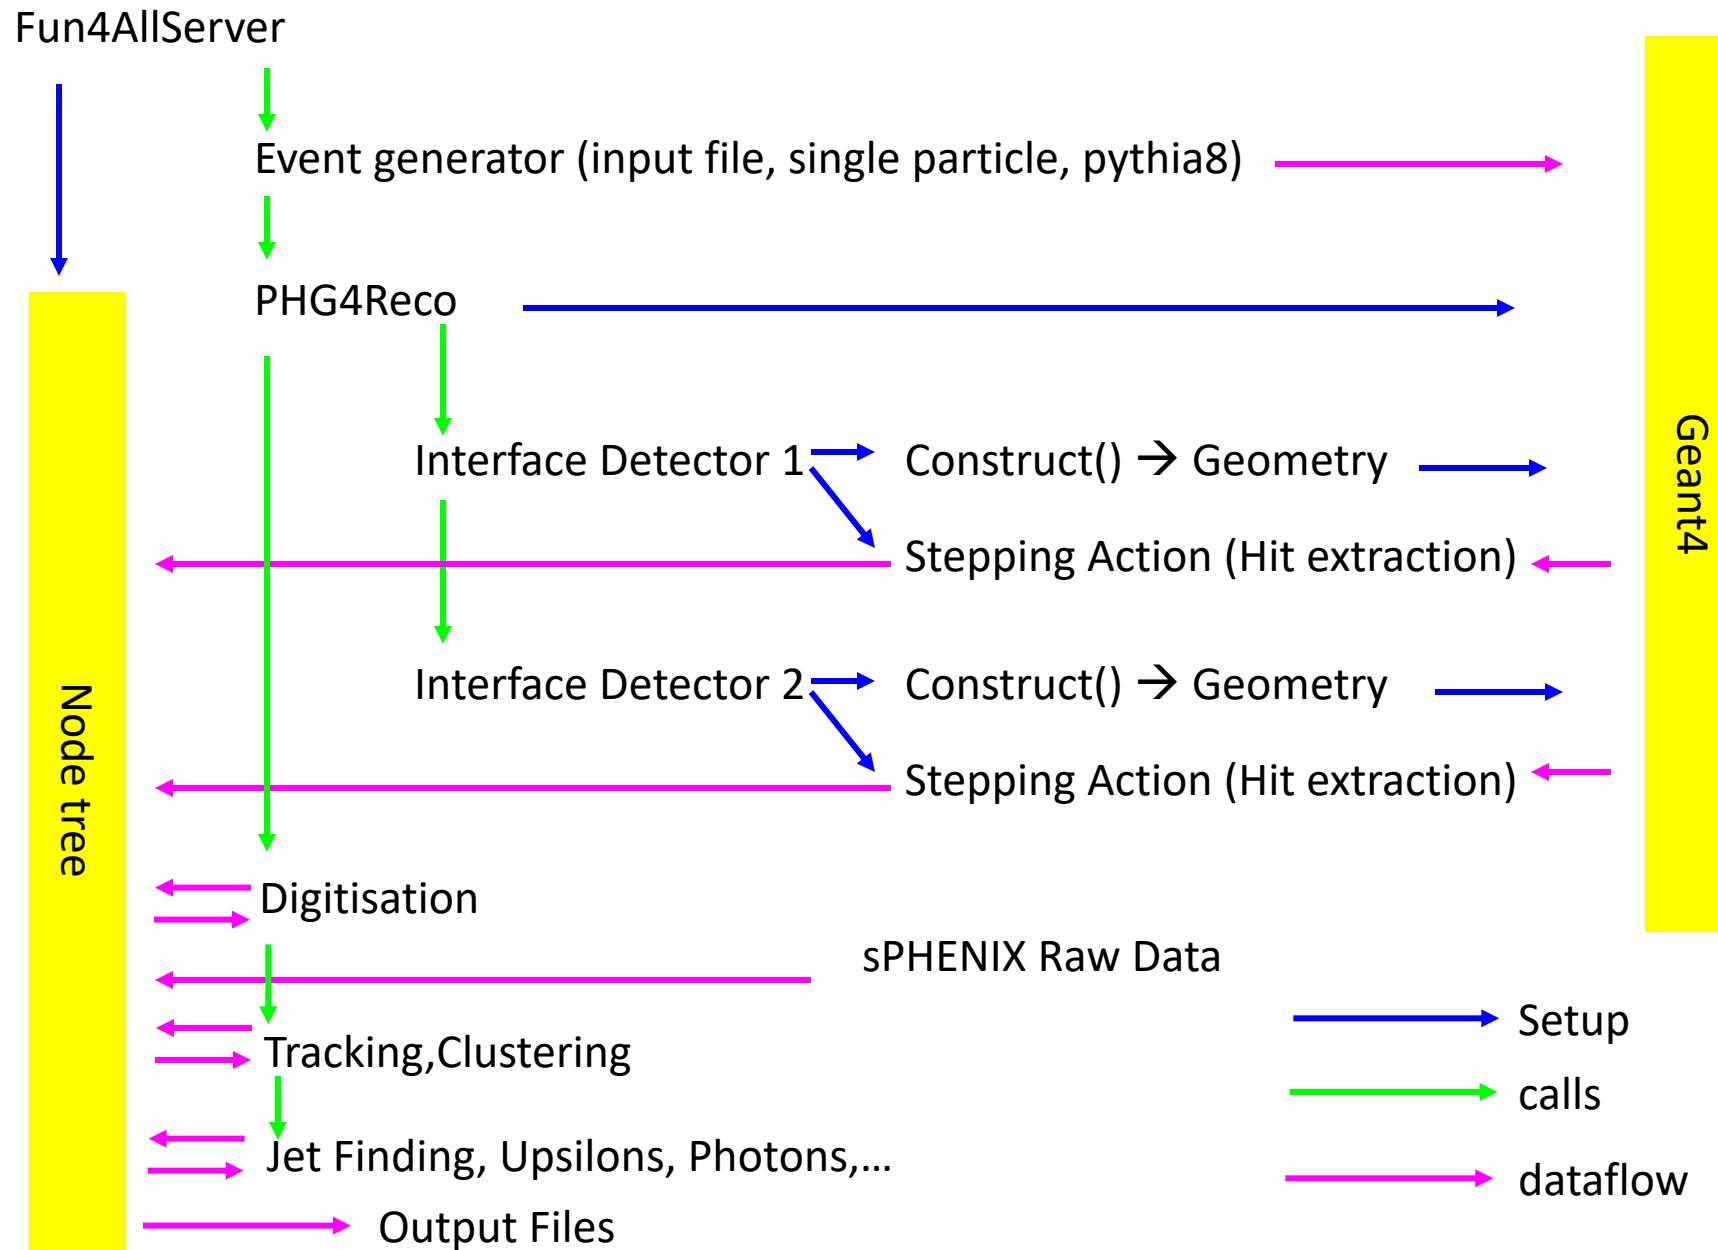
Simulations



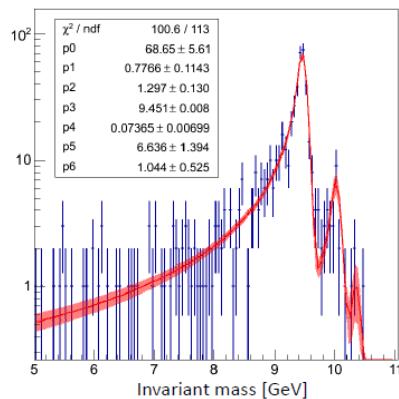
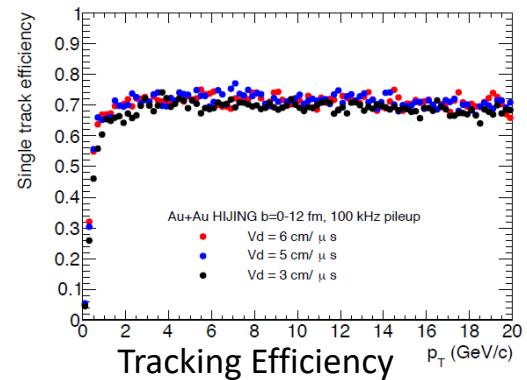
Peripheral Au+Au @ 200 GeV Hijing event

Event displays are the easy part, how do we actually analyze this mess?

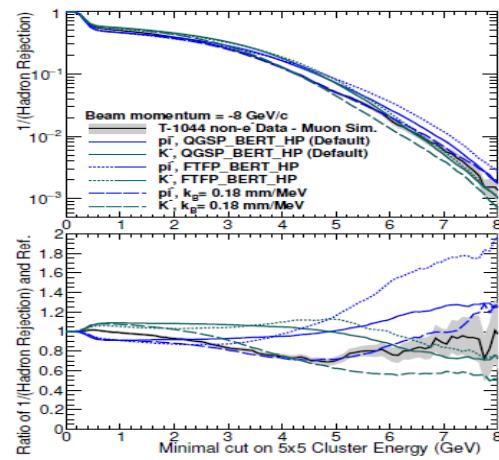
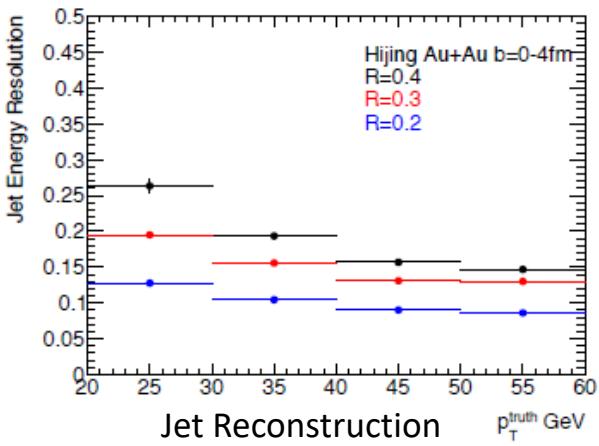
G4 program flow within Fun4All



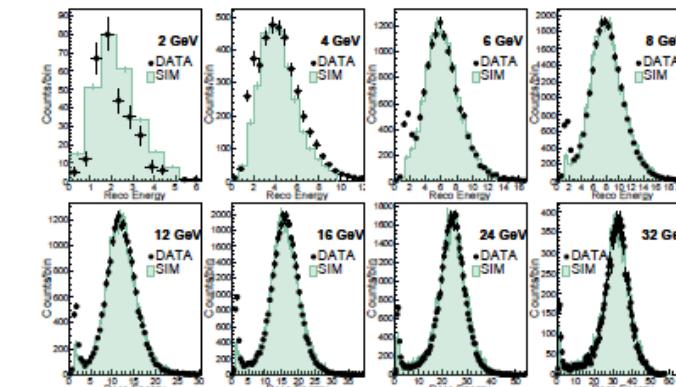
More than just pretty event displays



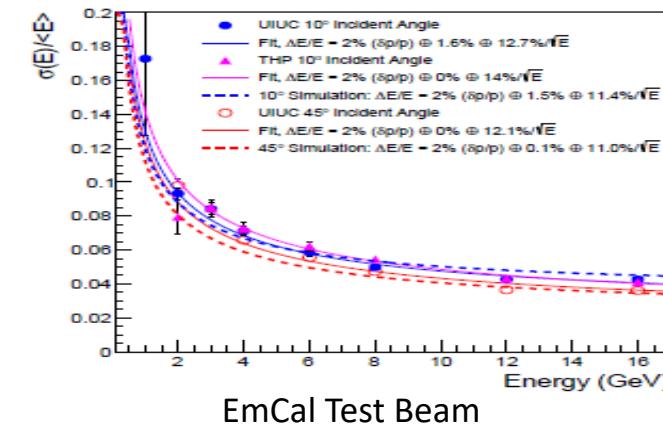
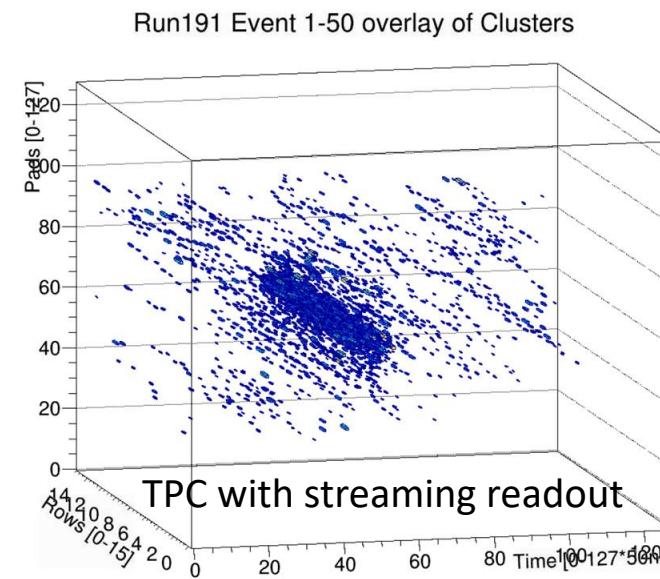
Upsilon Reconstruction



EmCal Hadron Rejection



Hadronic Calorimeter Test Beam



EmCal Test Beam

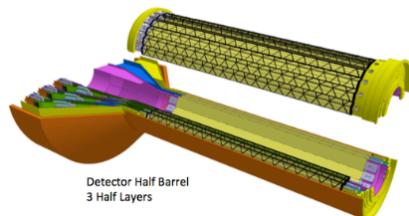
sPHENIX Run Plan

Year	Species	Energy [GeV]	Phys. Wks	Rec. Lum.	Samp. Lum.	Samp. Lum. All-Z
Year-1	Au+Au	200	16.0	7 nb ⁻¹	8.7 nb ⁻¹	34 nb ⁻¹
Year-2	p+p	200	11.5	—	48 pb ⁻¹	267 pb ⁻¹
Year-2	p+Au	200	11.5	—	0.33 pb ⁻¹	1.46 pb ⁻¹
Year-3	Au+Au	200	23.5	14 nb ⁻¹	26 nb ⁻¹	88 nb ⁻¹
Year-4	p+p	200	23.5	—	149 pb ⁻¹	783 pb ⁻¹
Year-5	Au+Au	200	23.5	14 nb ⁻¹	48 nb ⁻¹	92 nb ⁻¹

Only events in a \pm 10cm vertex range have the full tracking information
 Event vertex range is \pm 30cm (right column)

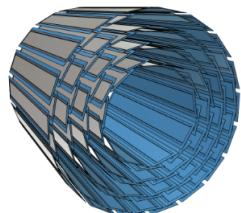
Species	Energy [GeV]	Rec. Lum.	Samp. Lum.	Samp. Lum. All-Z
Au+Au	200	35 nb ⁻¹ (239 billion)	80 nb ⁻¹ (550 billion)	214 nb ⁻¹ (1.5 trillion)
p+p	200	—	197 pb ⁻¹ (8.3 trillion)	1.0 fb ⁻¹ (44 trillion)
p+Au	200	—	0.33 pb ⁻¹ (0.6 trillion)	1.46 pb ⁻¹ (2.6 trillion)

The large data producers in 200 GeV Au+Au (worst case)



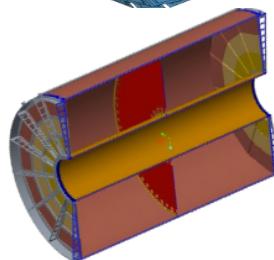
Monolithic Active Pixel Sensors (MAPS)

~ 35Gbit/s



Intermediate Silicon Strip Tracker (INTT)

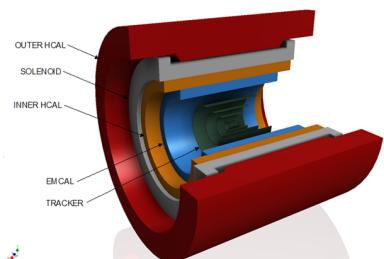
~ 7Gbit/s



Compact Time Projection Chamber (TPC)

~ 80Gbit/s

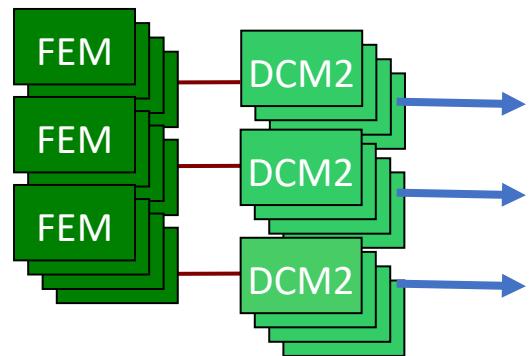
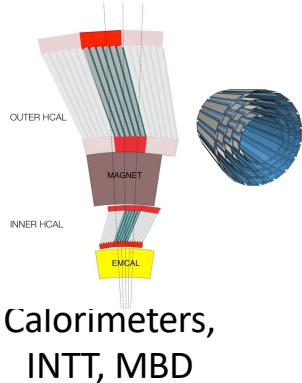
Calorimeters (primarily Emcal, hadronic cal.) **~ 8Gbit/s**



~ 130Gbit/s

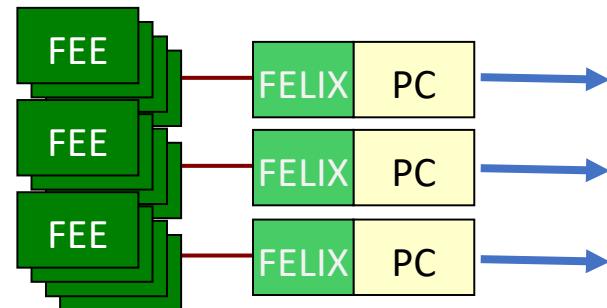
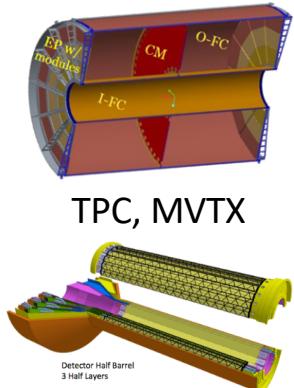
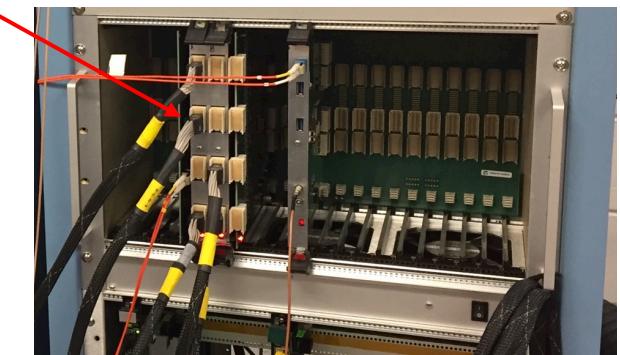
After applying RHIC x sPHENIX duty factor **~ 100Gbit/s**

Two Classes of Front-end Hardware



The calorimeters, the INTT, and the MBD re-use the PHENIX “Data Collection Modules” (v2)

Triggered readout



The TPC and the MVTX are read out through the ATLAS “FELIX” card directly into a standard PC

Streaming readout



ATLAS FELIX Card



Installed in a PC

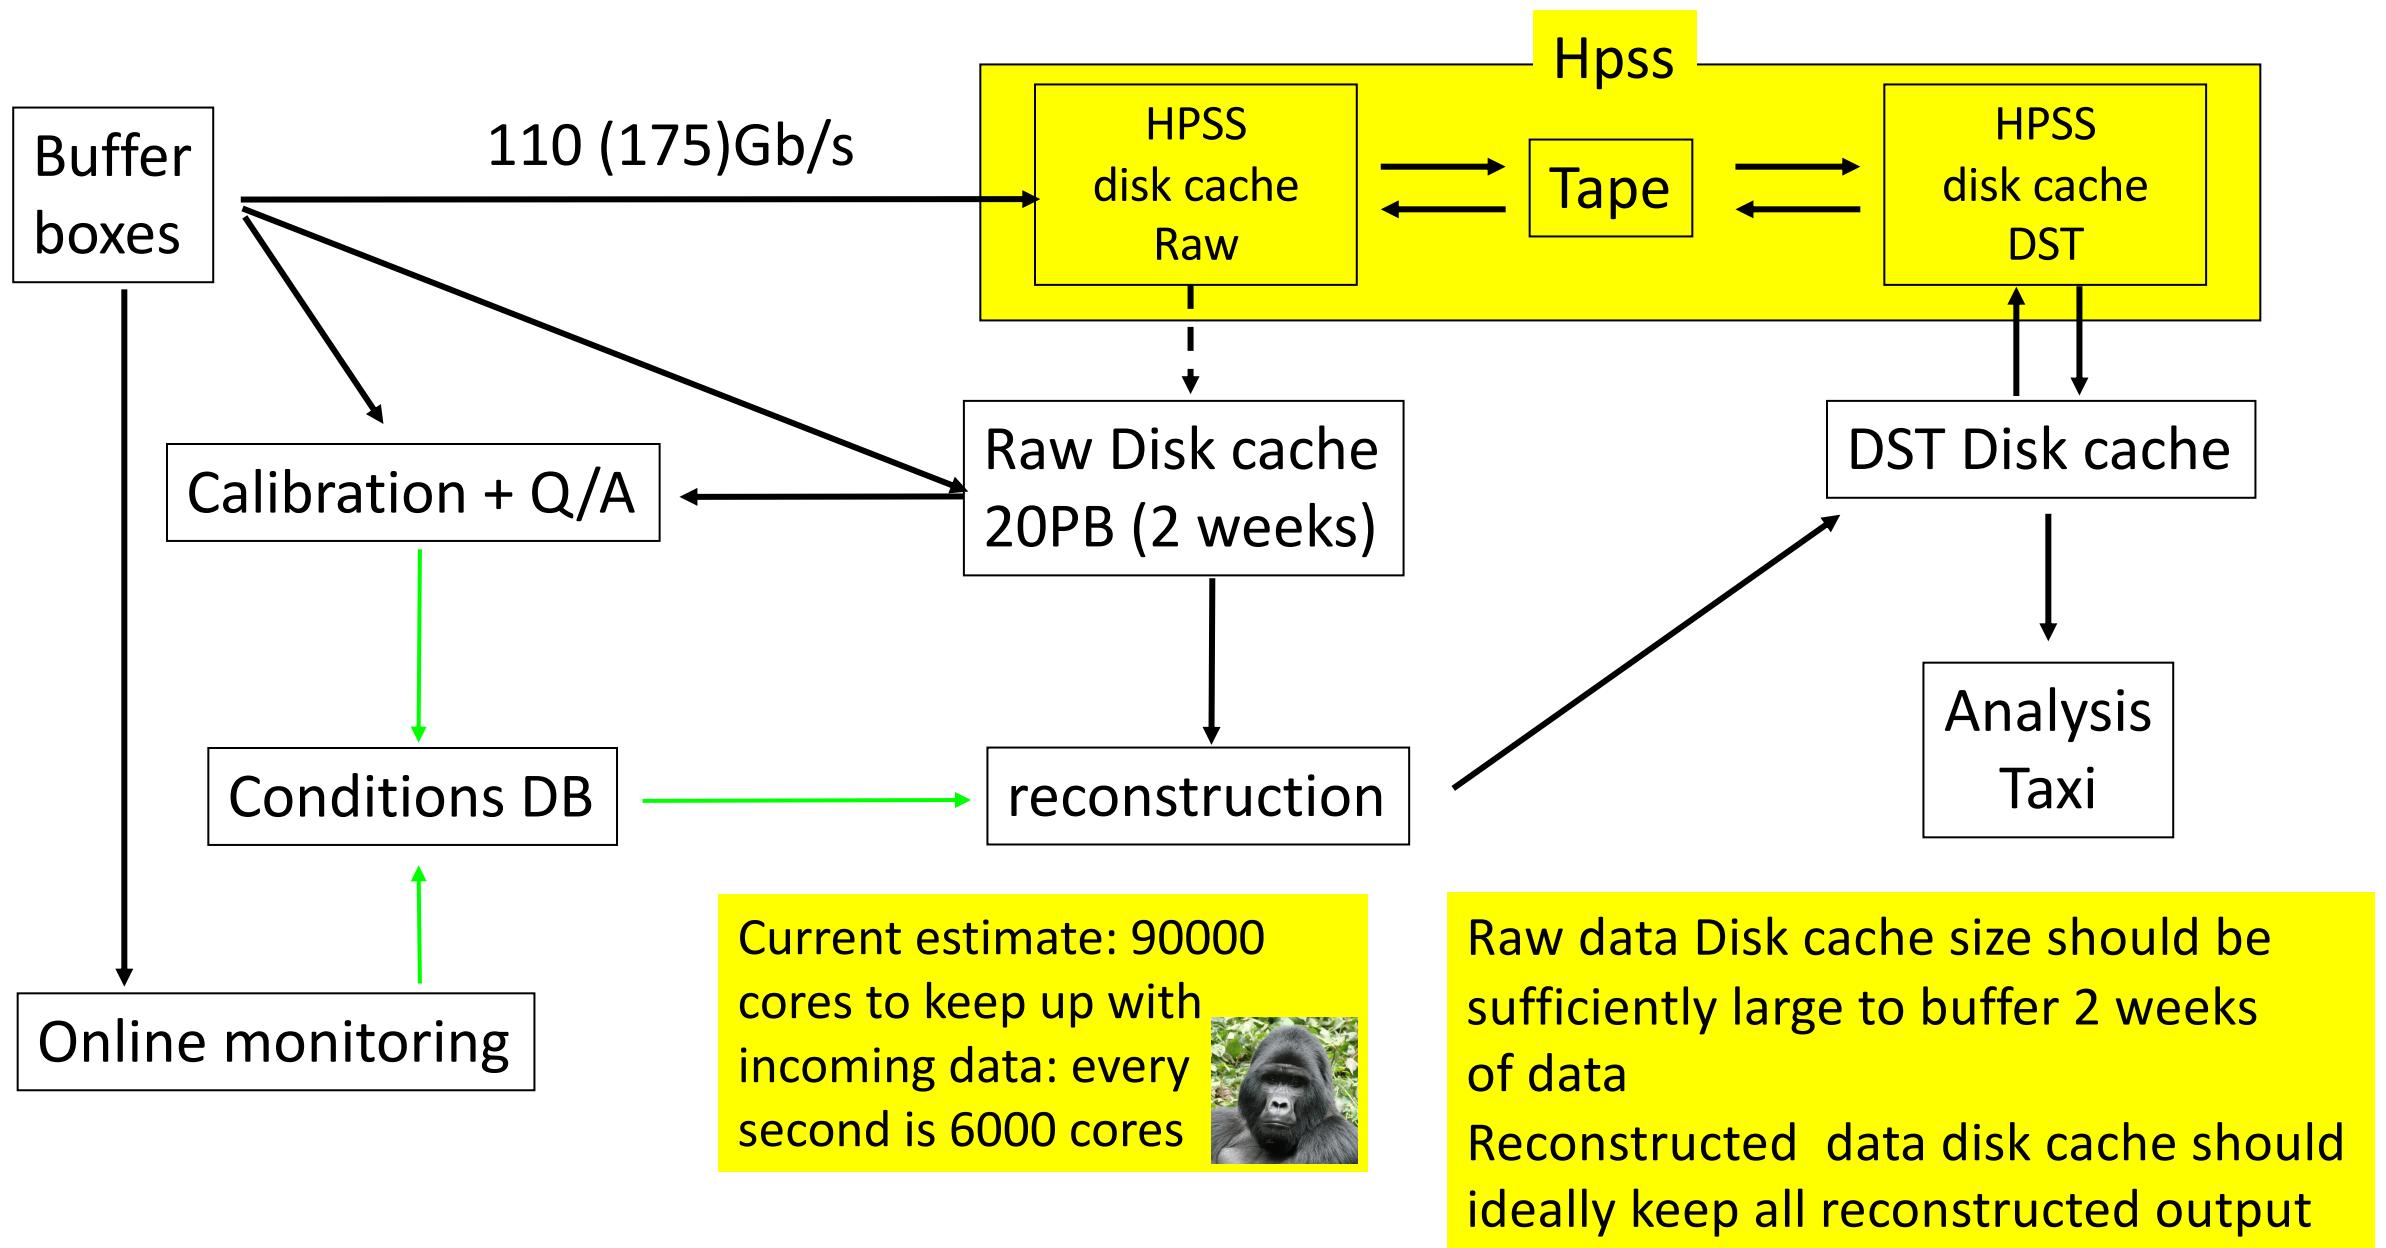
Event rates

- The run plan is to acquire 15 kHz of collisions (subtract 1.5 weeks for rampup):
 - Run 1: Au+Au: 14.5 weeks · 60% RHIC uptime · 60% sPHENIX uptime → 47 billion events, 1.7 Mbyte/event → 75PB (110Gb/s)
 - Run 2 and 4: p+p, p+A: 22 weeks · 60% RHIC uptime · 80% sPHENIX uptime → 96 billion events, 1.6 Mbyte/event → 143PB
 - Run 3 and 5: Au+Au: 22 weeks · 60% RHIC uptime · 80% sPHENIX uptime → 96 billion events, 1.6 Mbyte/event → 143PB
- The DAQ system is designed for a sustained event rate of 15kHz
- We cannot “trade” smaller event sizes for higher rates.
- The new detectors (TPC, MVTX) will not have the ultimate data reduction factors until at least Year-4 (based on ALICE and others’ experience), with lzo compression in daq only minor change in data volume

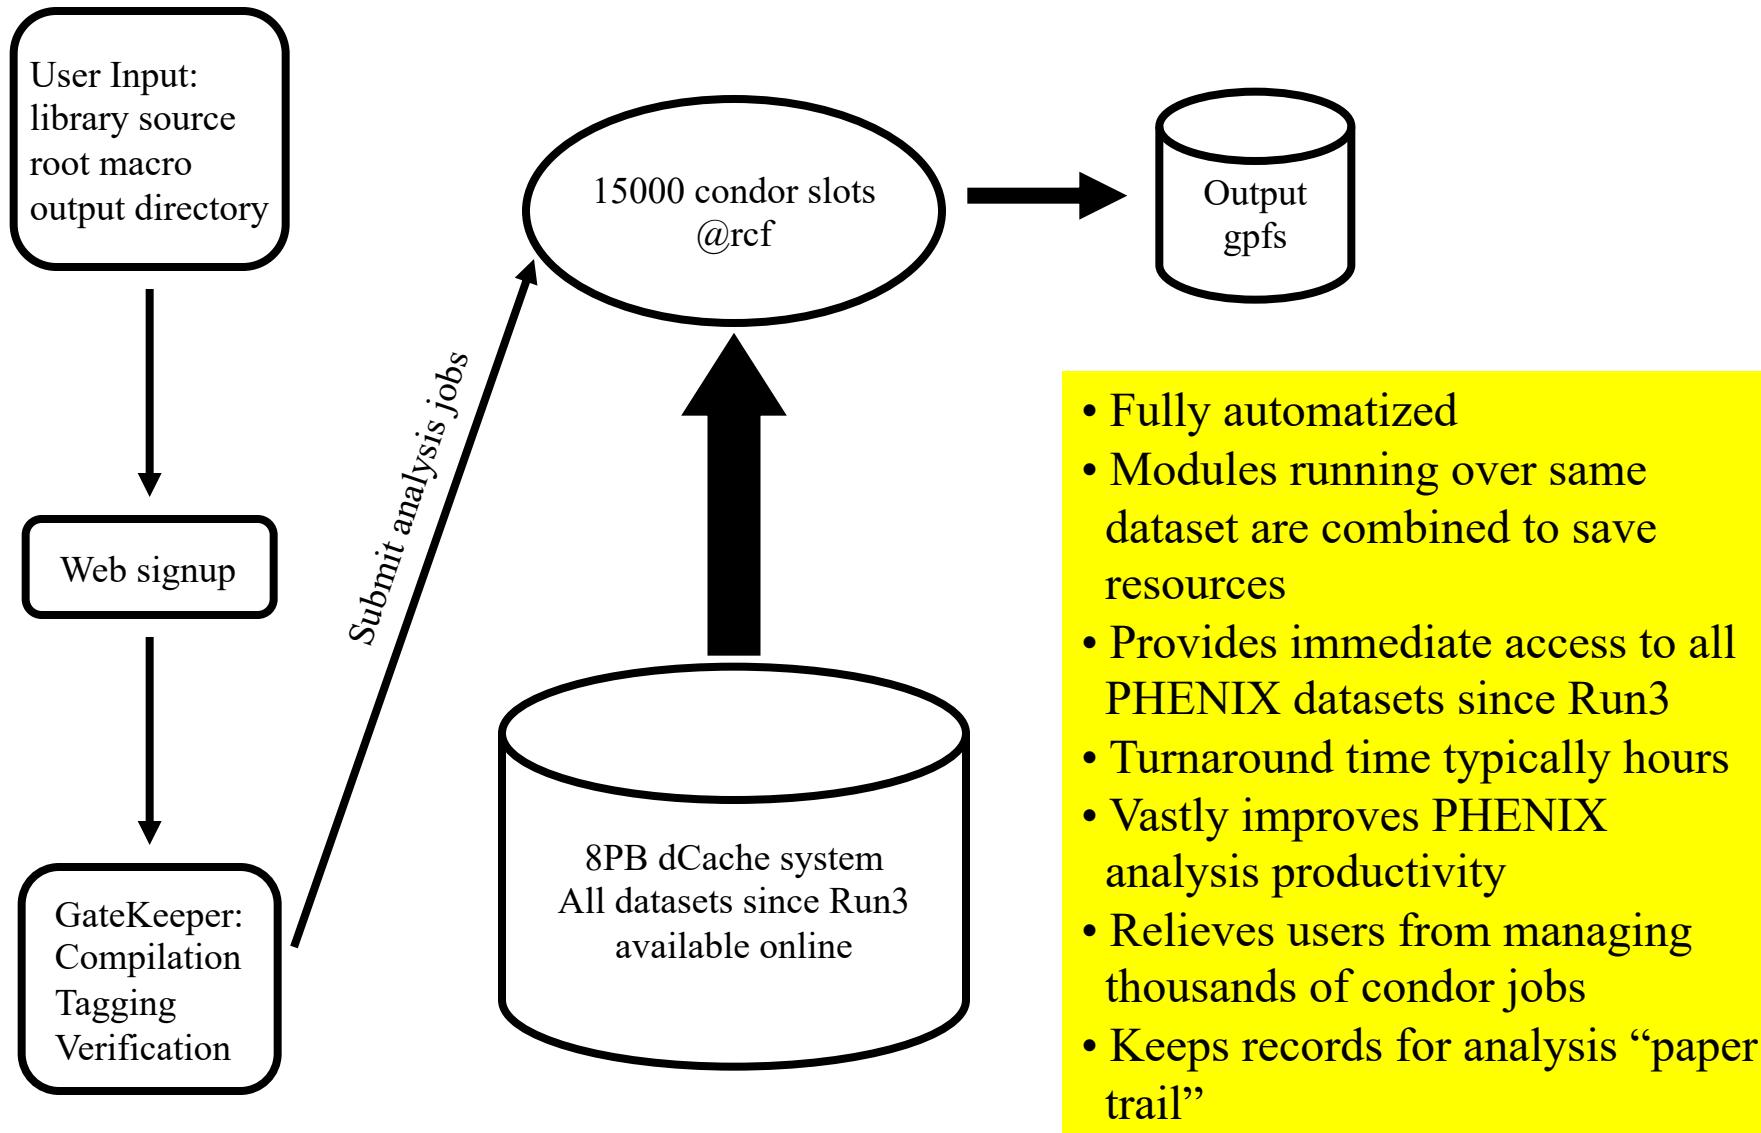
Event building

- Nothing in online requires the assembling of **all** events (we do not employ level2 triggers, all events selected by our level1 triggers are recorded)
- Moving the event builder to the offline world makes it a lot simpler
 - The offline event builder does not have to keep up with peak rates
 - In offline we have many more cpus at our disposal
 - Crashes can be easily debugged
 - No loss of data due to event builder issues
 - Subevents are ordered in raw data files
- Disadvantage: Need to deal with ~60 input files in data reconstruction
- Still need to build a fraction of the events for monitoring purposes
- Combining triggered with streamed readout is going to be fun

Reconstruction + analysis flow



The Analysis Taxi, mother of all trains



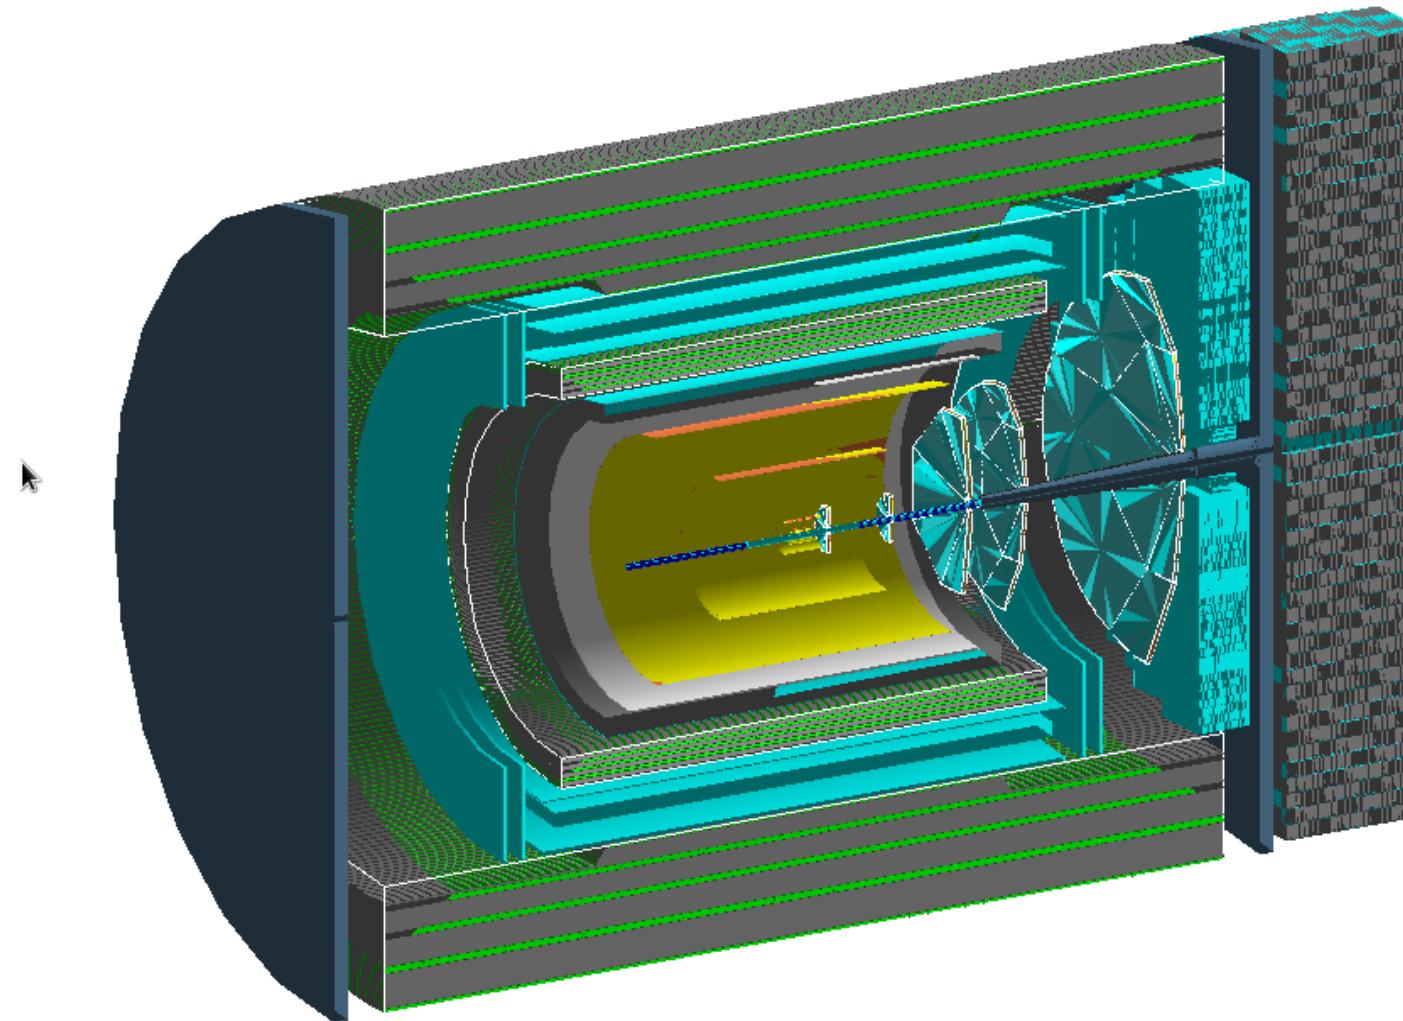
Towards EIC

The challenges of heavy ion event reconstruction dwarfs whatever the EIC will throw at us
sPHENIX might (will) evolve into the day 1 EIC detector by adding forward instrumentation
In any case - the central parts of the proposed EIC detectors are very similar to sPHENIX
→ Any improvement made to sPHENIX software will benefit the EIC program

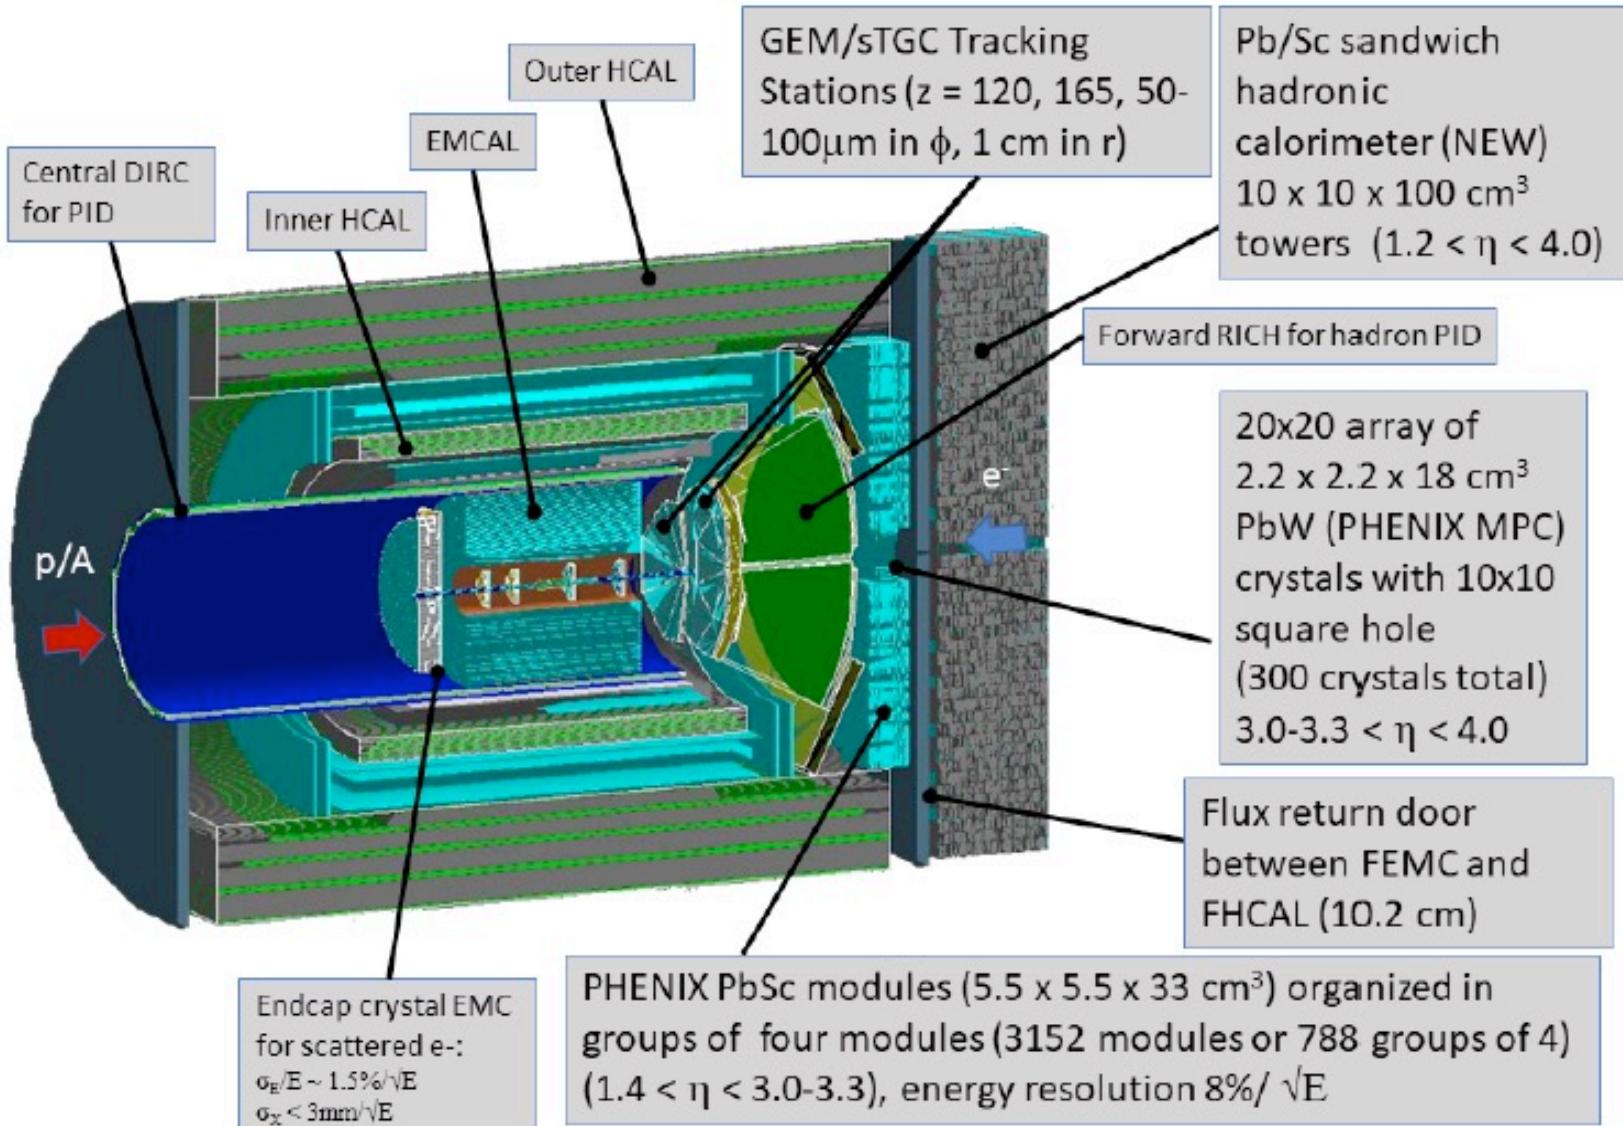
Our software is containerized (singularity), the code in github, our libraries are in cvmfs
<https://github.com/sPHENIX-Collaboration/Singularity>
→ The OSG can provide the computing resources needed by non sPHENIX EIC users

We have some tutorials how to put together simple detectors from basic shapes:
<https://github.com/sPHENIX-Collaboration/tutorials>

fsPHENIX – forward instrumentation for cold QCD

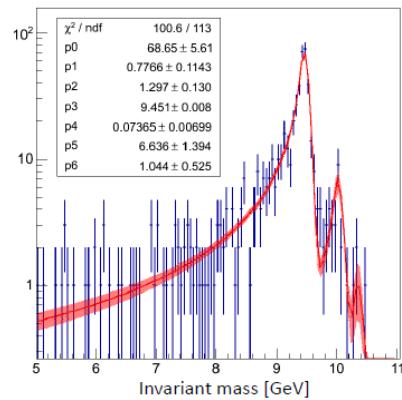
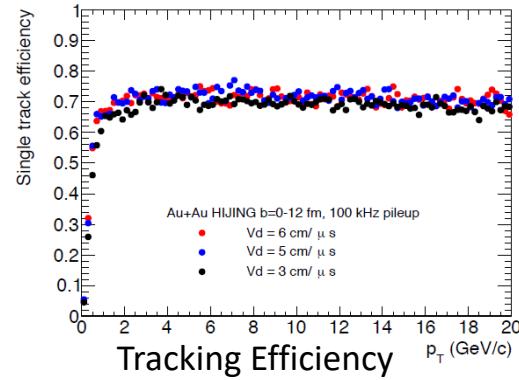


ePHENIX out of the box

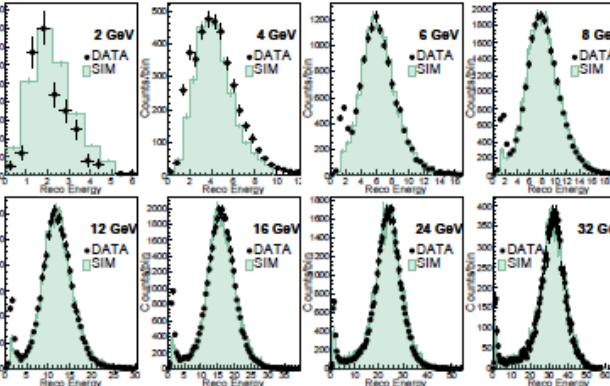


Fully developed G4 model, including digitization and reconstruction

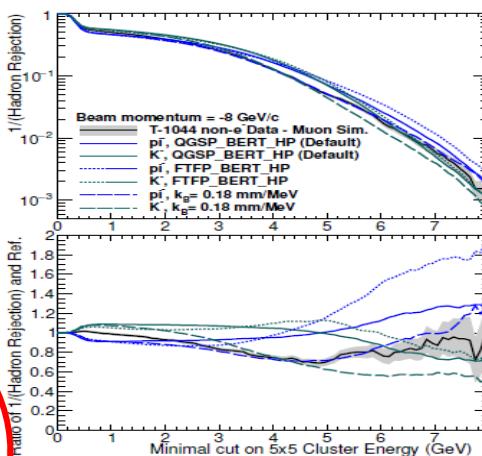
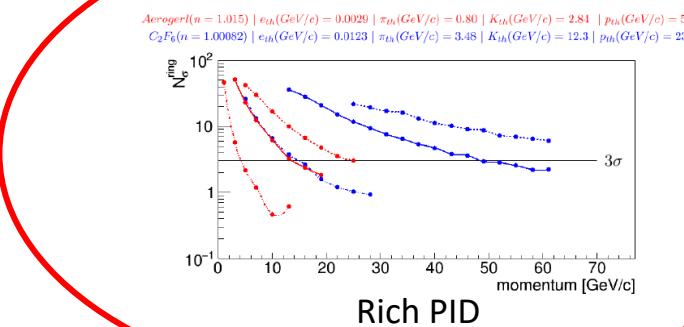
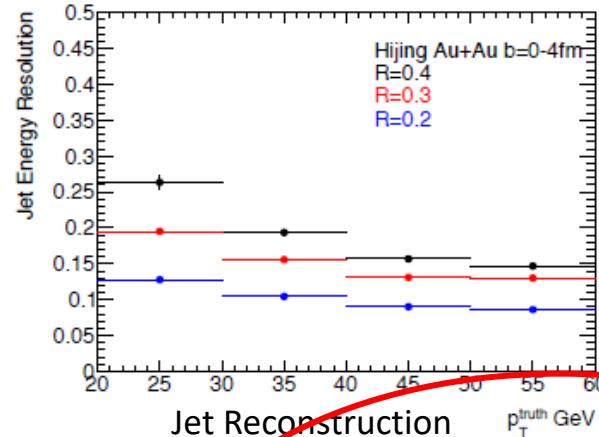
More than just pretty event displays



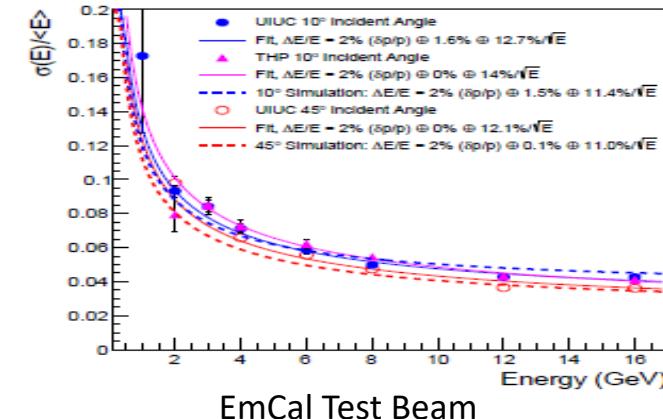
Upsilon Reconstruction



Hadronic Calorimeter Test Beam



EmCal Hadron Rejection



EmCal Test Beam

Contains all you need to
simulate and analyze data NOW
Stonybrook is looking into the
detection of leptoquarks, hopefully
first results in time for the EIC users
meeting July 22-26