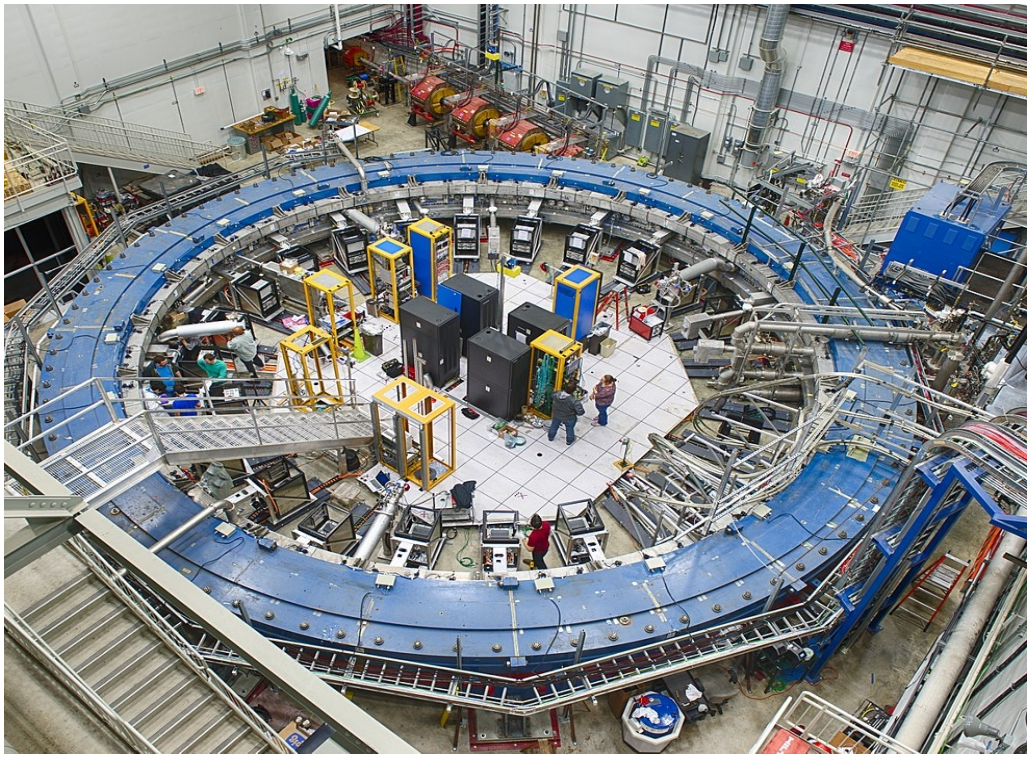

Computing challenges and prospects for the g-2 experiment

Software and Computing
for Small HEP Experiments
Workshop - 15 Nov 2021

Paolo Girotti (INFN Pisa)
on behalf of the Muon g-2 collaboration

Collaboration

- Muon g-2 is a relatively small HEP experiment
- Part of Fermilab Muon Campus
- ~200 collaborators from 7 countries
- Very young community (40% under 35 yo)

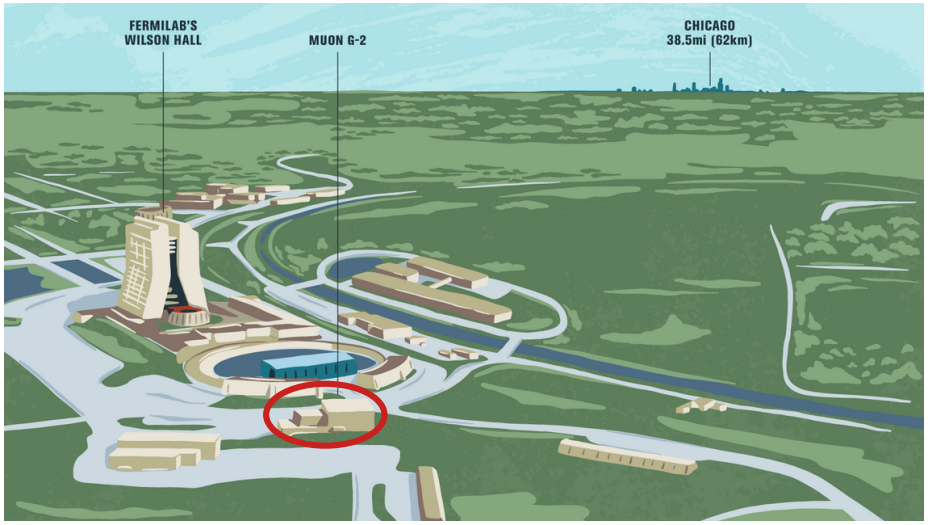

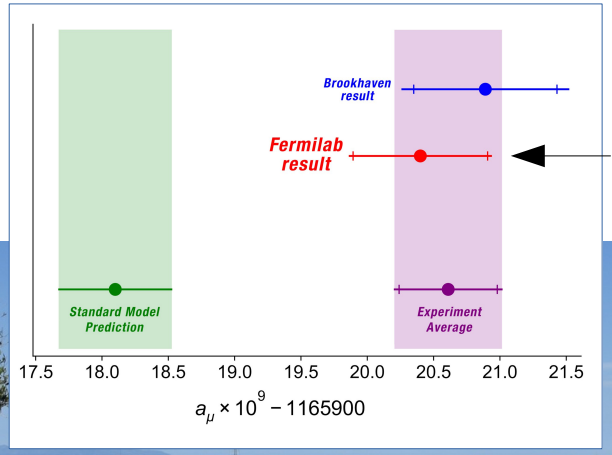

First physics result
7 April 2021

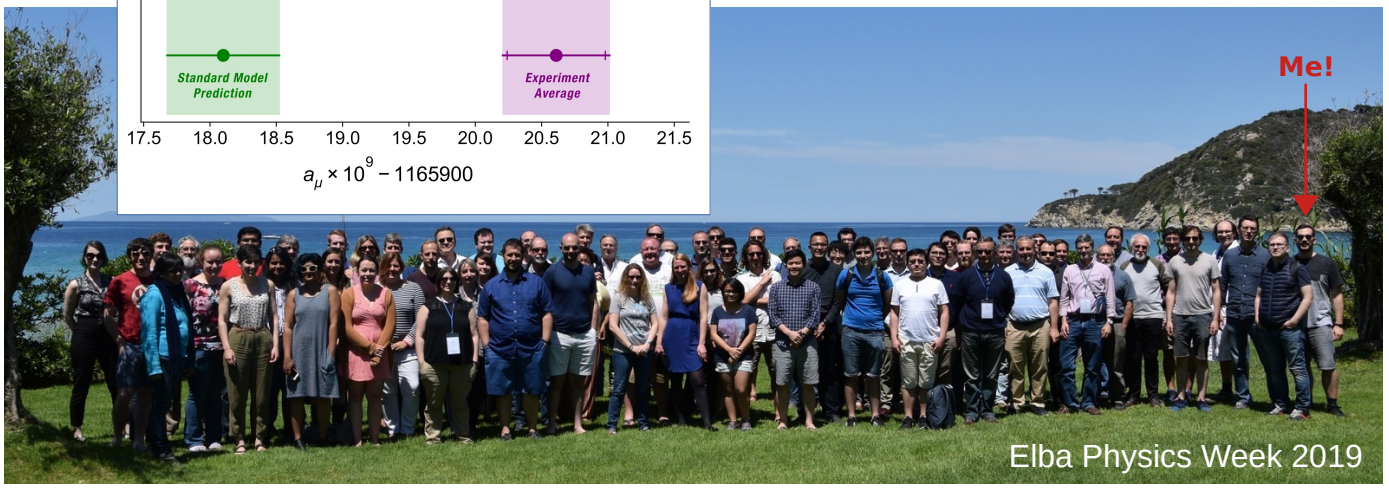

Muon g-2 primer

- The Muon g-2 Experiment (E989) measures the anomalous precession moment of the muon **very precisely**
  - Goal: repeat and improve BNL (2001) measurement with 4x precision
- A beam of polarized muons circulates inside a storage ring
- The magnetic field is measured by NMR probes
- Decay positrons carry the precession signal and are detected by calorimeters

$$\vec{\omega}_a = \vec{\omega}_s - \vec{\omega}_c = -\left(\frac{g-2}{2}\right)\frac{e\vec{B}}{m} \equiv -a_\mu \frac{e\vec{B}}{m}$$

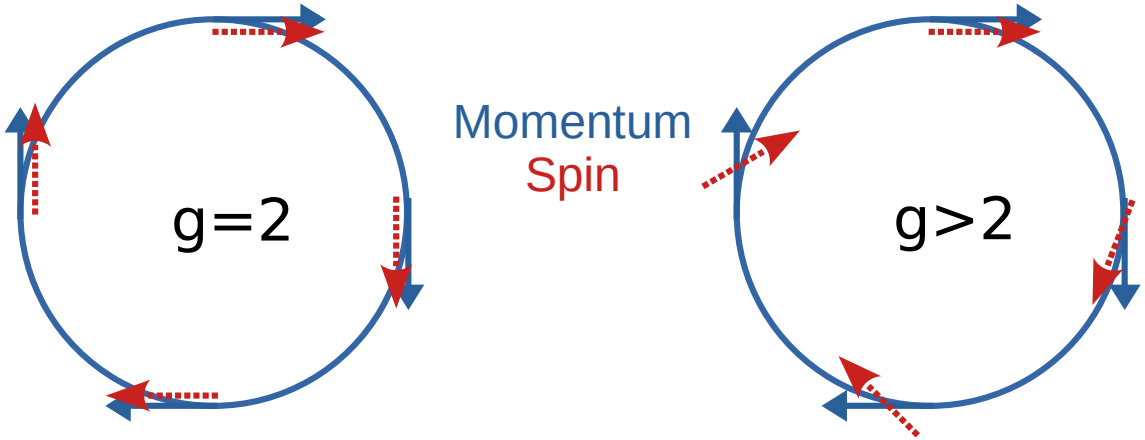

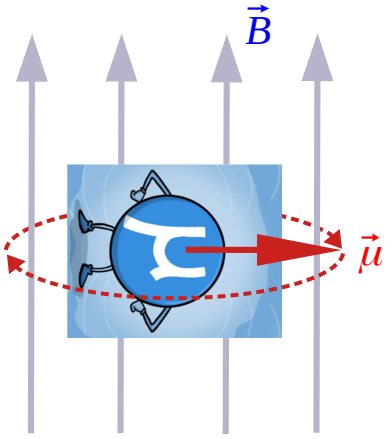

The muon precession frequency is influenced by the interactions with **all** possible virtual particles

Master formula

$$a_\mu = \frac{\omega_a}{\tilde{\omega}'_p(T_r)} \frac{\mu'_p(T_r)}{\mu_e(H)} \frac{\mu_e(H)}{\mu_e} \frac{m_\mu}{m_e} \frac{g_e}{2}$$

Constants known from other experiments

ω_a

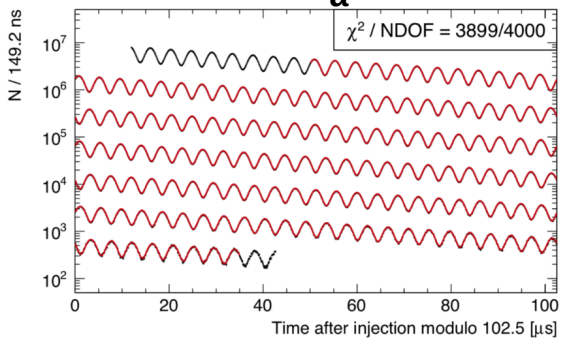

ω_p

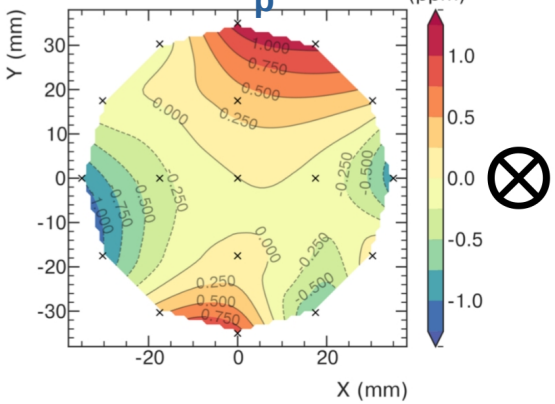

$\rho(r)$

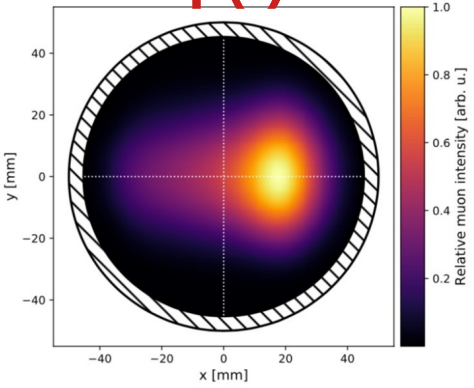

Three quantities measured:

- ω_a: Muon anomalous precession frequency
- ω_p: Larmor precession frequency of protons in water (mapping B)
- ρ_r: Muon distribution in the storage ring

Goal: measure a_μ with 140 ppb accuracy (100 stat + 100 syst)

The Muon g-2 Experiment

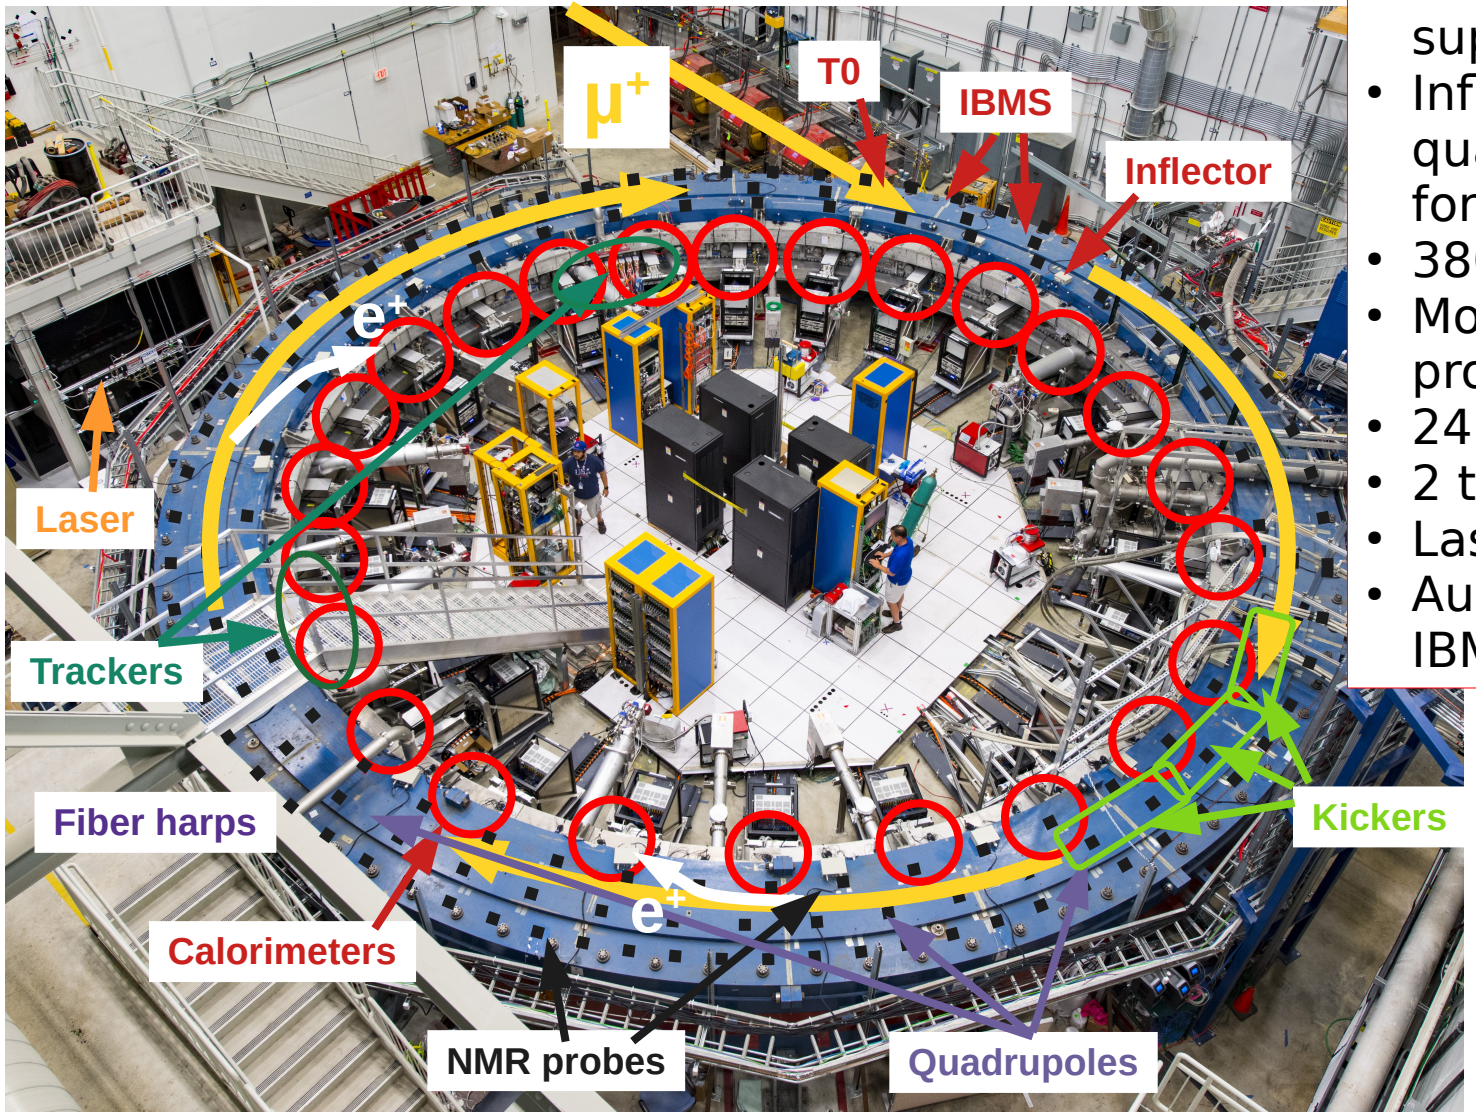

- 15 meter wide dipole superconducting magnet
- Inflector, kickers, quadrupoles, collimators for beam insertion
- 386 NMR probes
- Moving trolley with 17 probes
- 24 calorimeters
- 2 tracker stations
- Laser calibration system
- Auxiliary detectors: T0, IBMs, Fiber harps

A lot of components and detectors
→ a lot of **data**!

How much data?

- 12 bunches of muons every second
- Muons circulate for 700 μs, generating ~2000 positron hits on calorimeters
- 1296 calorimeter channels @800 MHz
- 2048 tracker channels @400 MHz
- One subrun (2 GB) every 5 seconds
- 24/7 running for 4-7 months each year
- Currently running 5th year of production

  → Collected **~6 PB** of raw data so far

  → Plus simulation ~1.5 PB

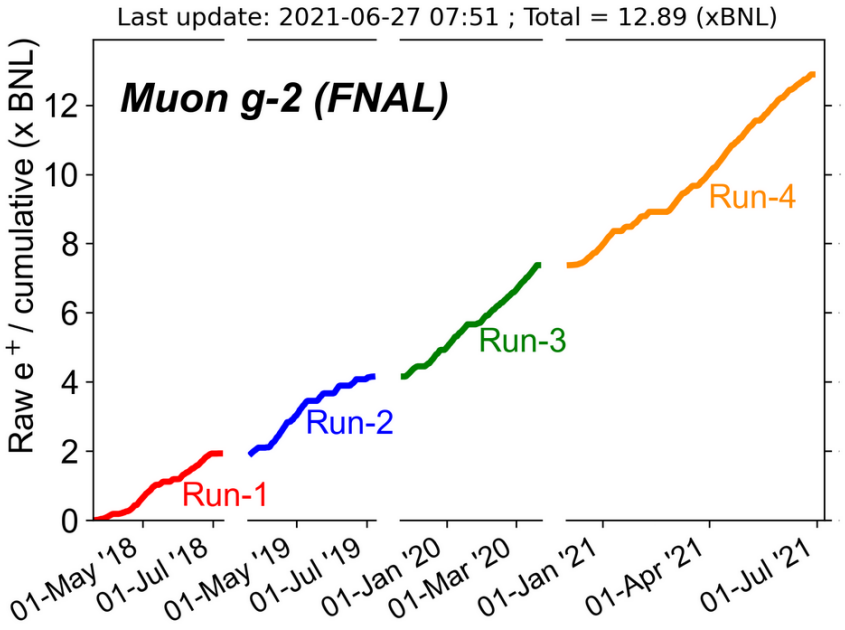

	Raw	Nearline	Offline	Field
Run1	1.0 PB	-	2.2 PB	29 TB
Run2	1.1 PB	11 TB	2.5 PB	48 TB
Run3	1.3 PB	18 TB	2.0 PB*	87 TB
Run4	2.3 PB	35 TB	0.2 PB*	98 TB

*Under production

How do we *see* data?

- **Online** with Midas DAQ / DQM
- Instantaneous feedback from all components
- Fraction of data gets displayed with **Plotly**

MIDAS

DQM

Beam data Nov 14th 2021

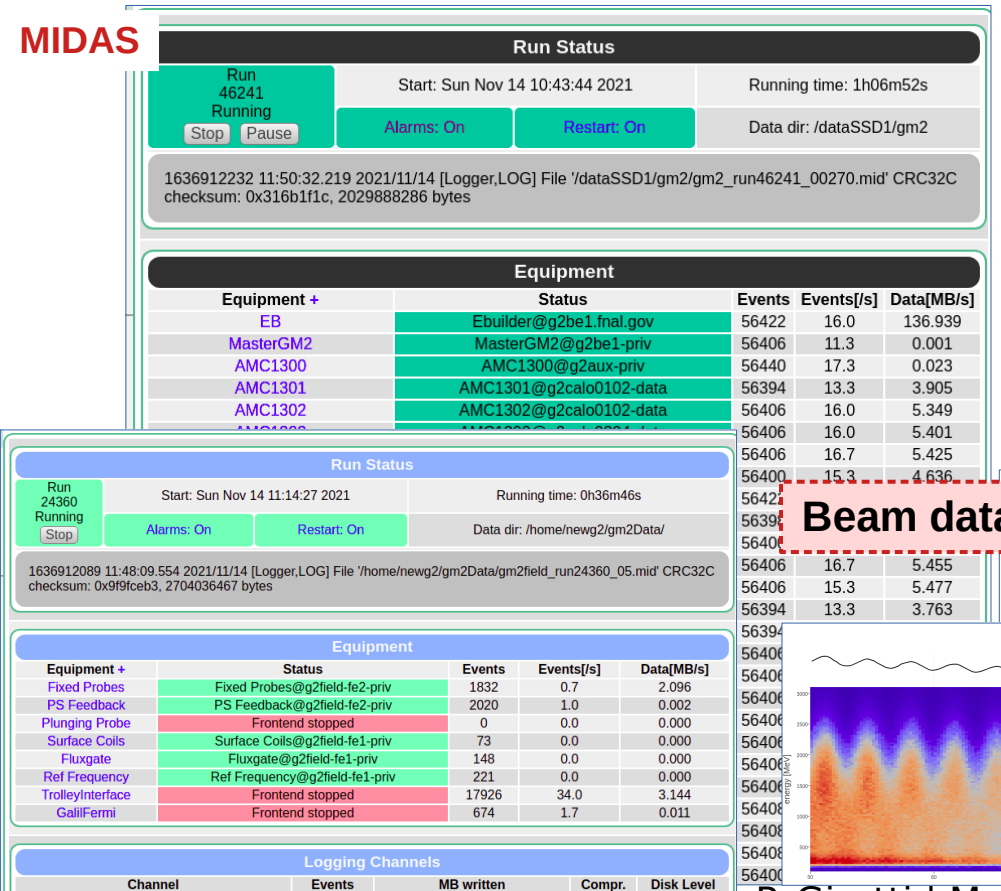

How do we *see* data?

- **Nearline** for fast turnaround analyses
- Simplified reconstruction of the data ~30 minutes after acquisition
- Interactive plots with Bokeh
- Web visualization of ROOT trees

Run number	time	comment	nEvents	
Search	Search	Search	Search	Search
46188	2021-11-10 01:32:11.796052	production	103545.0	Y
46187	2021-11-09 23:32:05.258684	production	102702.0	Y
46186	2021-11-09 21:02:09.319796	idle run	0.0	T
46185	2021-11-09 20:47:04.427447	idle run	0.0	T
46184	2021-11-09 19:42:12.897917	alternative mode run	100461.0	T
46183	2021-11-09 18:07:05.400655	alternative mode run	2287.0	T
46182	2021-11-09 17:52:10.623573	standard mode run	0.0	T
46181	2021-11-09 17:42:04.655508	Double Pulse Odd=0 Even=60	0.0	L
46180	2021-11-09 17:37:10.919716	Double Pulse Even=0 Odd=65000	0.0	S
46179	2021-11-09 17:27:04.378694	Trolley run	0.0	T

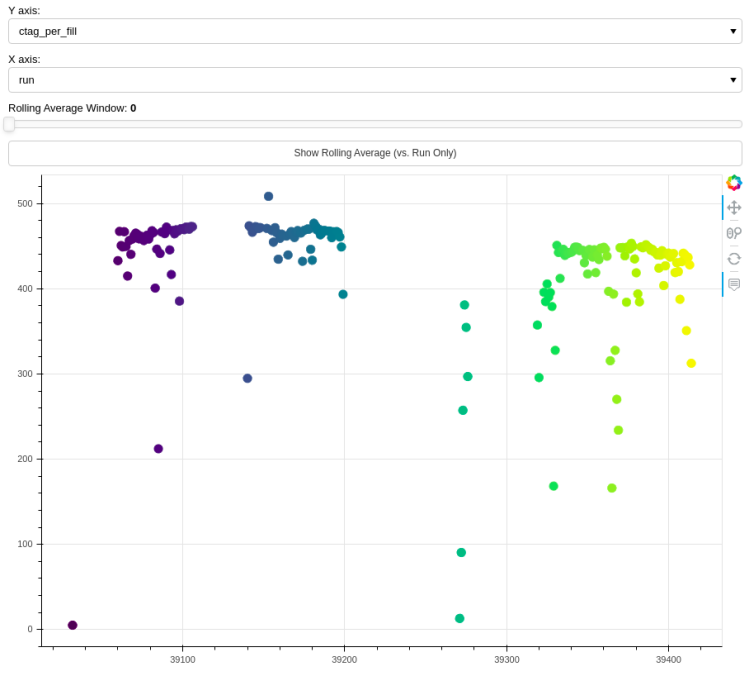

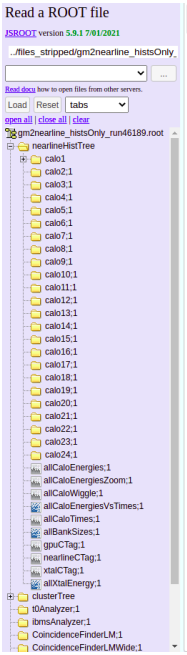

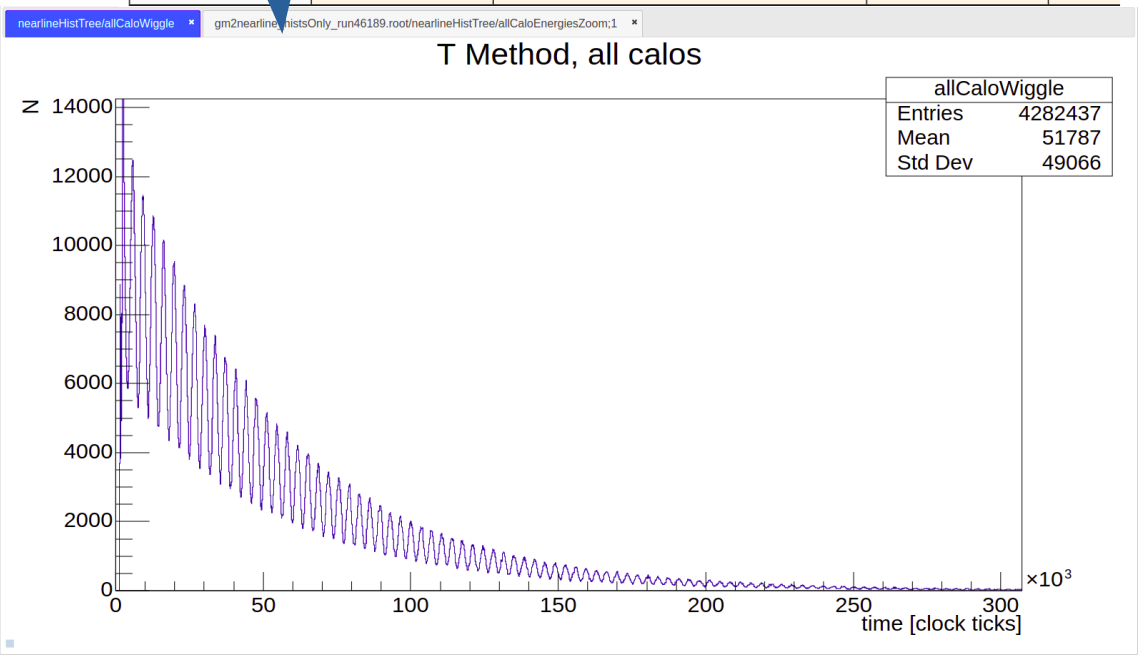

Reconstruction flow

- With Muon g-2 we get one value, but <u>many pieces</u> are needed
- Many analyses but all interconnected

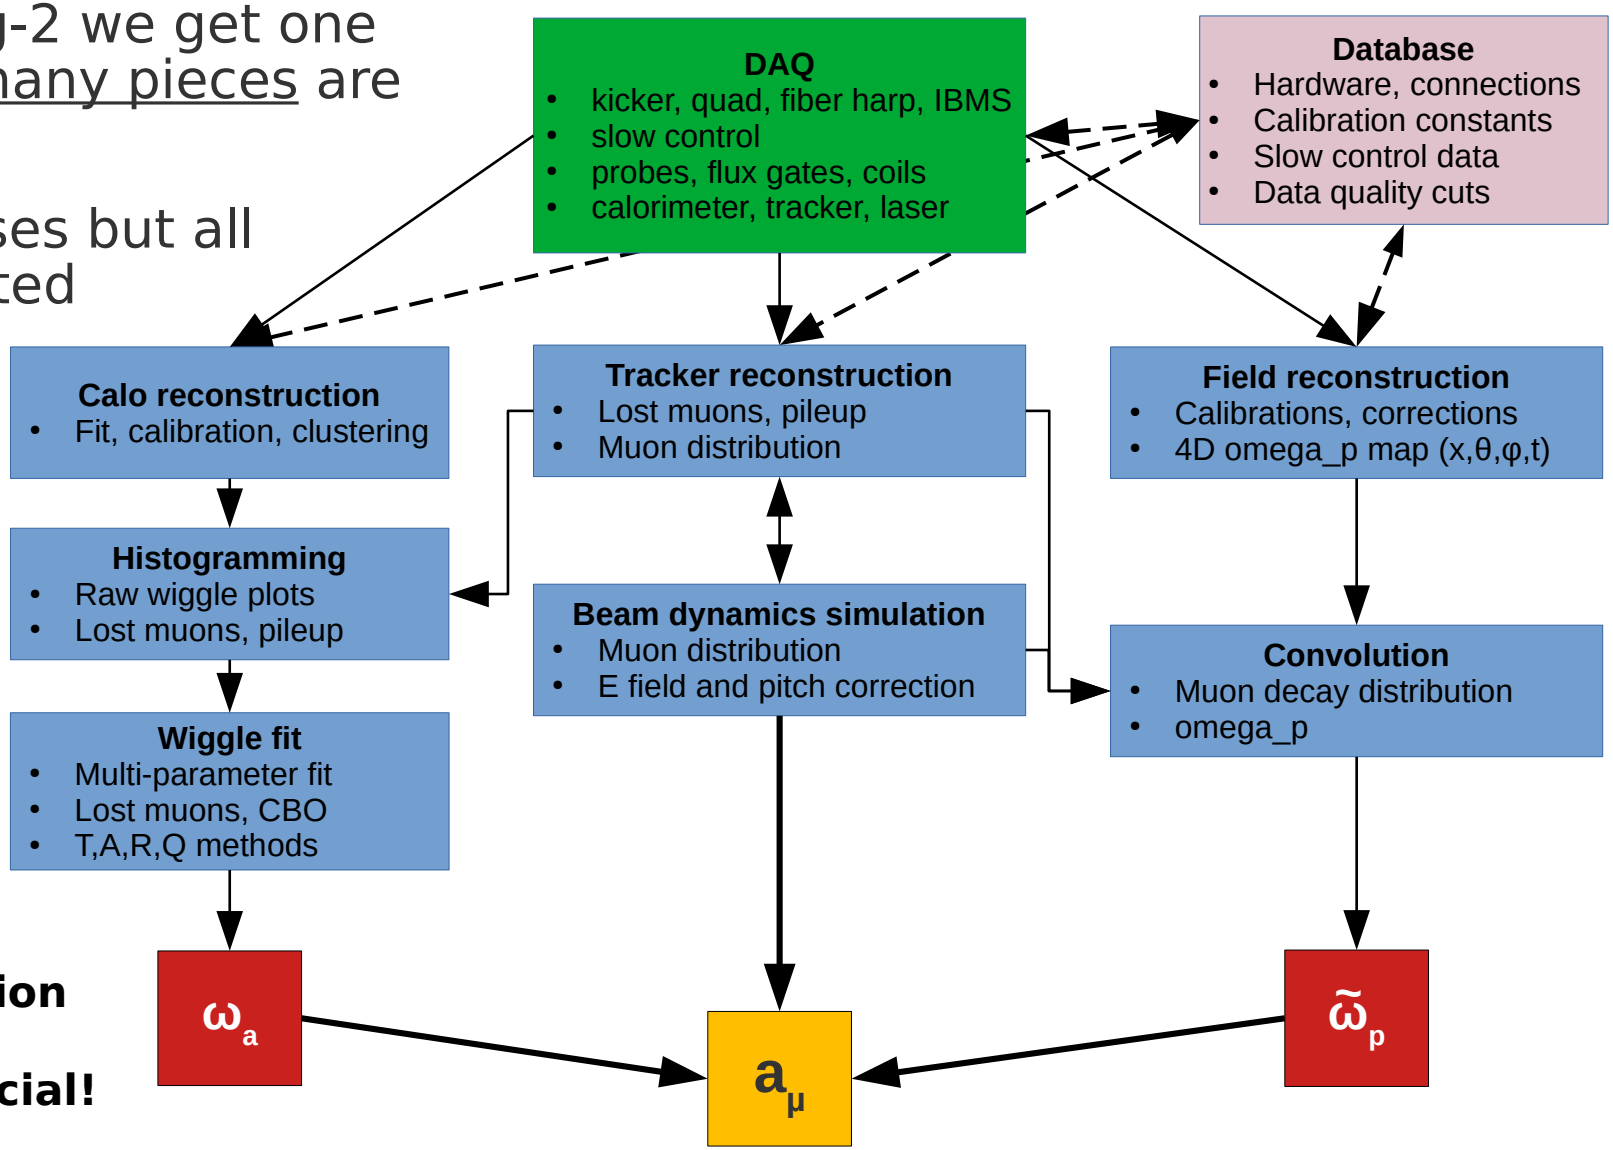

Solid reconstruction software and framework are crucial!

Data reconstruction

- **Offline** with full reconstruction of the data for full-fledged analyses
- We use ***art*** framework
  - Highly modular
  - Made for HEP physics
  - Seamless transition between simulation and real data
  - Non trivial learning-curve to climb
- **ROOT/C++** based software
- Code repository with version control
  - One of the first Fermilab experiments to use **git**
  - Hosted by Redmine

https://art.fnal.gov/

art

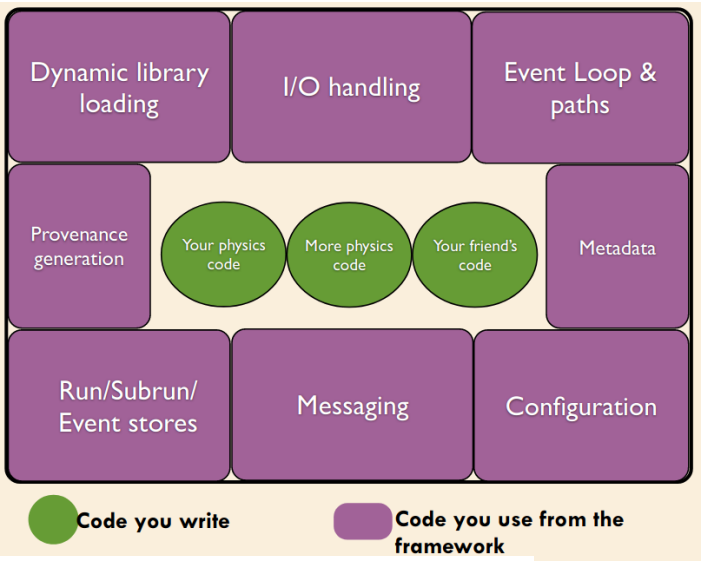

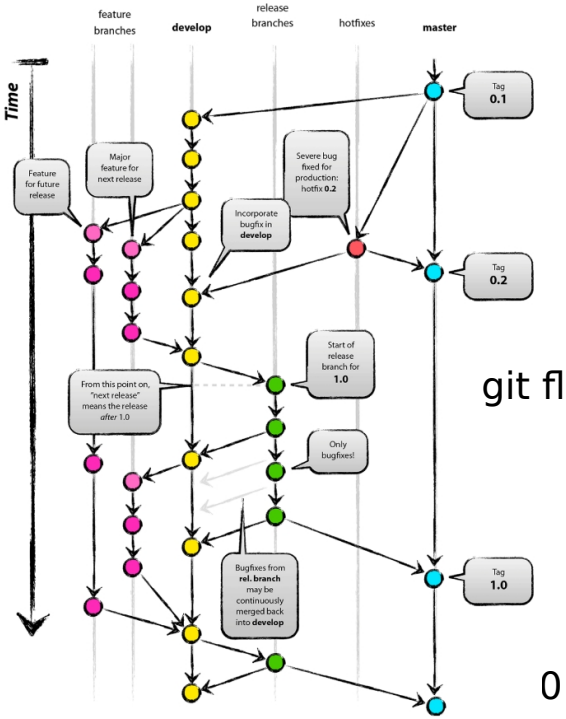

git flow

Database

- Database estensively used for both online DAQ configuration and offline data reconstruction
- Big effort during last year to transition critical reconstuction parameters and calibrations to a condition database supported by Fermilab SCD
- **PostgreSQL** interface
- **I**nterval **O**f **V**alidity based on run, subrun, event numbers

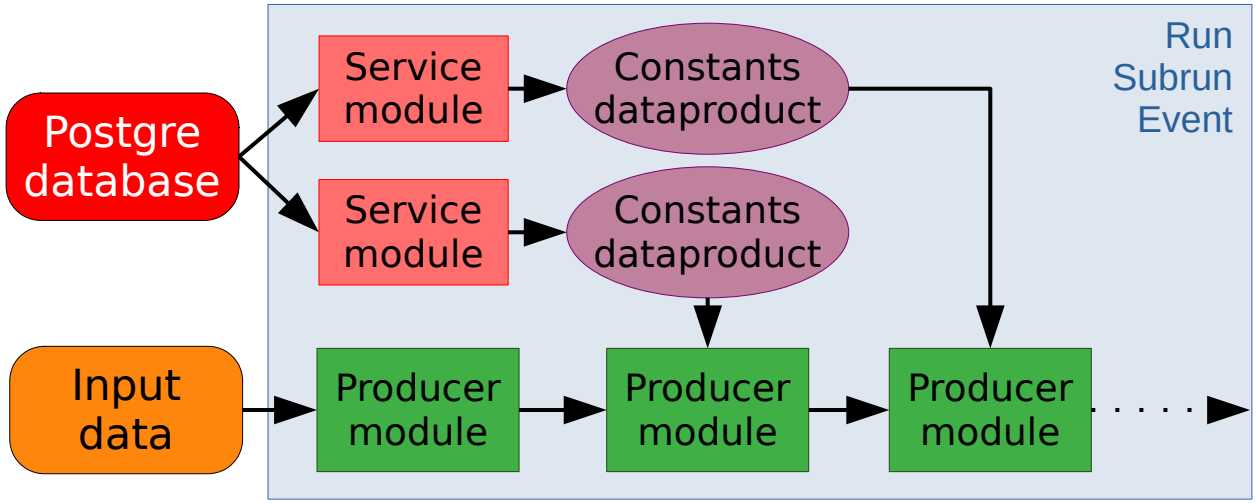

Terminal interface with **psql**

```
gm2_conditions_dev=> select calonum,xtalnum,oofcorrec
tion from oof_correction_data,oof_correction_iovs whe
re begin_time = 243370196 and calonum=1 and iov_id=10
 limit 10;
 calonum | xtalnum | oofcorrection
---------+---------+---------------
       1 |       0 |      0.932184
       1 |       1 |       0.94373
       1 |       2 |      0.933066
       1 |       3 |      0.928901
       1 |       4 |      0.930891
       1 |       5 |      0.918331
       1 |       6 |      0.940355
       1 |       7 |      0.901765
       1 |       8 |       0.93363
       1 |       9 |      0.939446
(10 rows)
```

Running on GRID

- Production jobs run on **Grid**
  - Both onsite and offsite
  - 5000 reserved slots for g-2 at all times
  - Data I/O is handled by SAM, IFDH and FTS
- **POMS** interface
  - Useful tool to keep track of which version a dataset has been reconstructed with
  - Automatic slicing of dataset, recovery of failed jobs and grid submissions with crontabs
  - Frequent monitoring is useful to ensure job success

Production Operations Management System

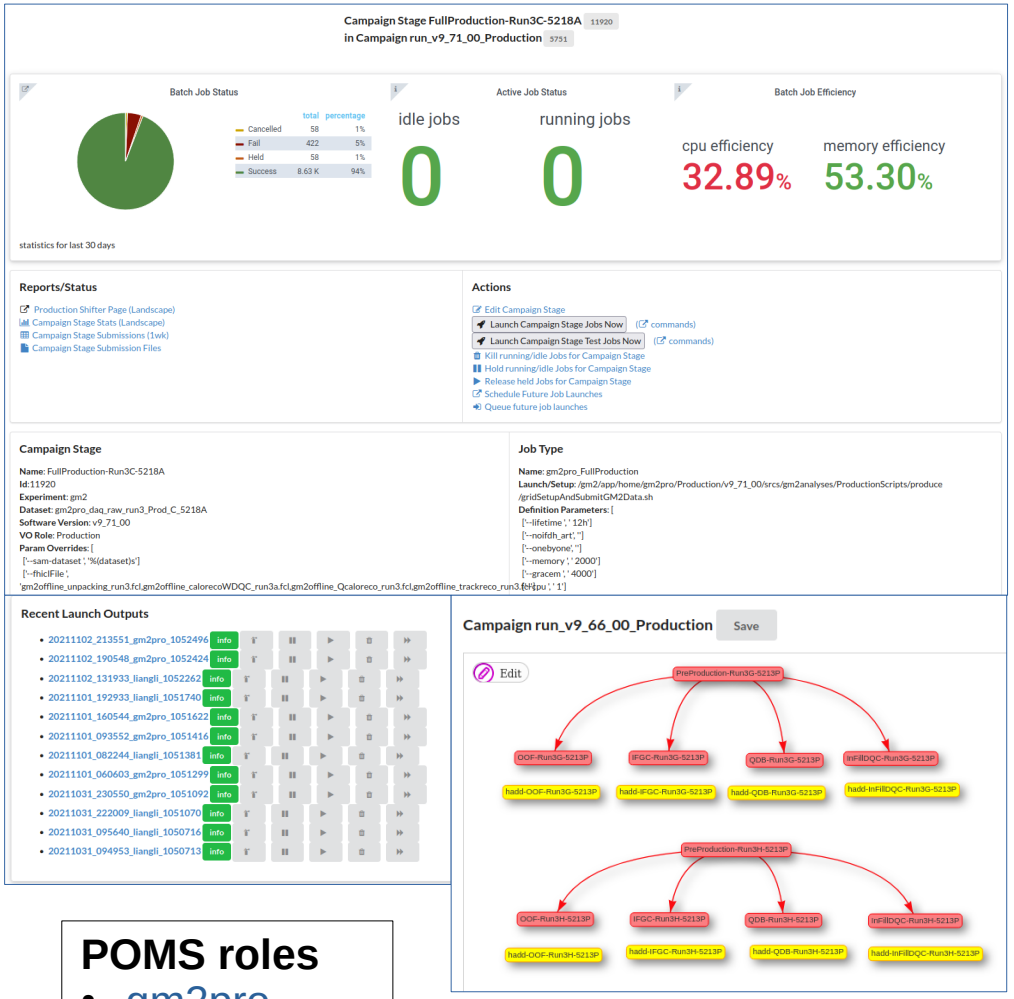

POMS roles
- gm2pro
- gm2shifter
- gm2analysis

Running on GRID

- Production job monitoring
  - **Grafana** for user efficiency and trend plots
  - More than 5 million jobs over the last year

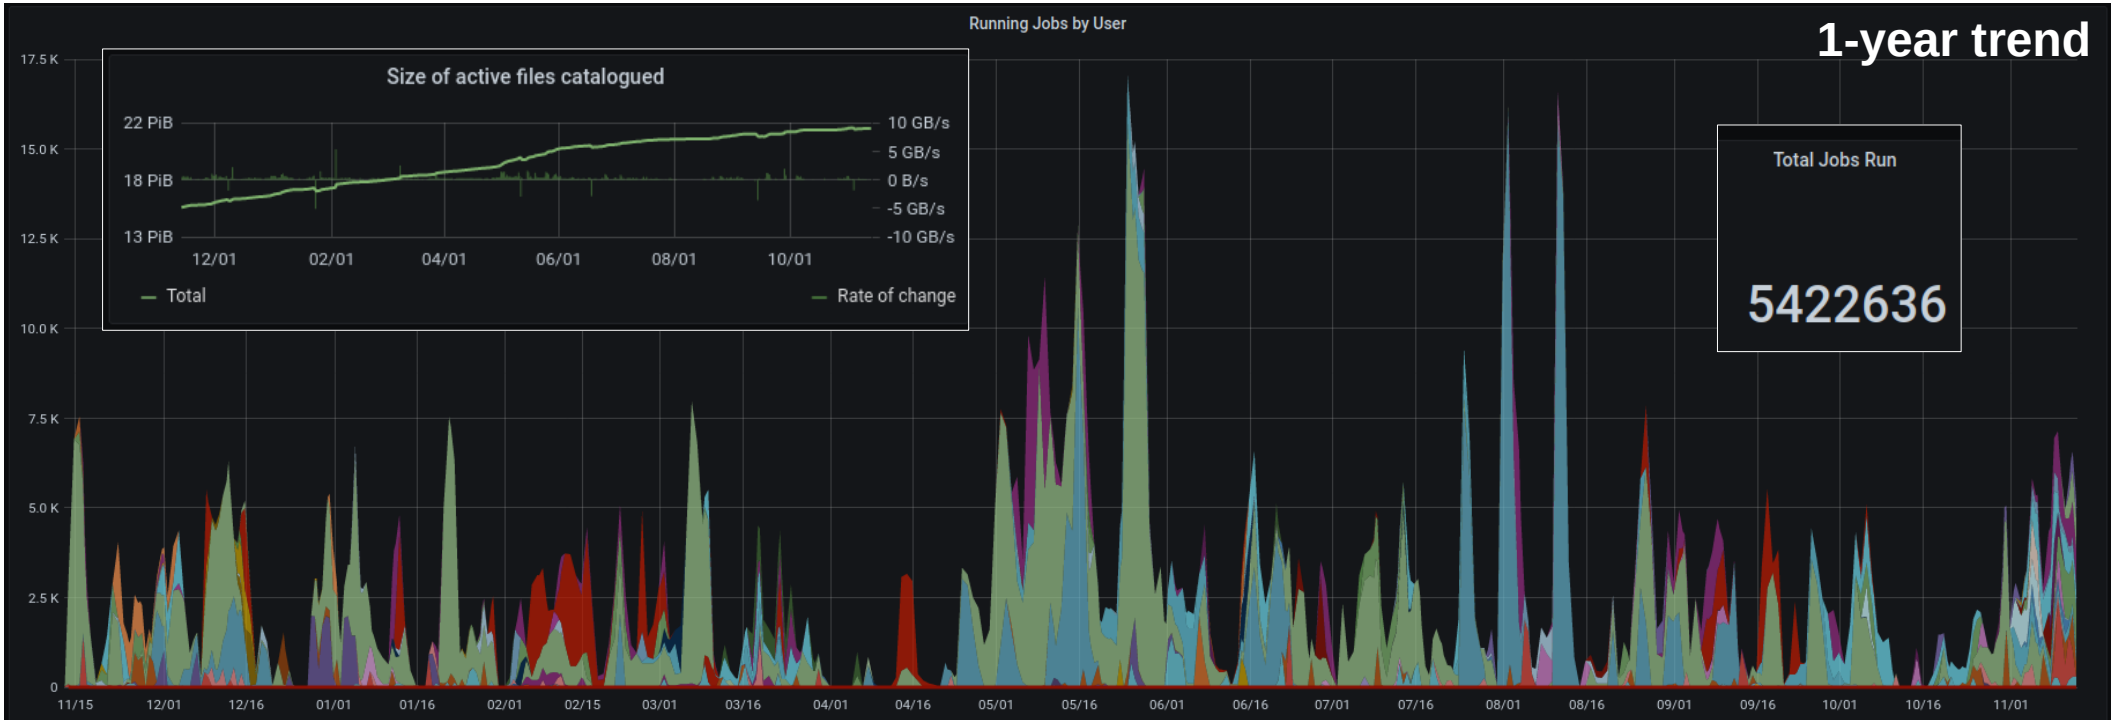

Simulation

- Important piece for g-2 physics, parallel to data taking
- Many simulation packages for the various parts of the beamline and the storage ring
  - MARS (Proton target)
  - BMAD (beamlines & g-2)
  - G4beamline (beamlines)
  - Gm2ringsim (injection & g-2)
  - COSY (g-2 storage ring)

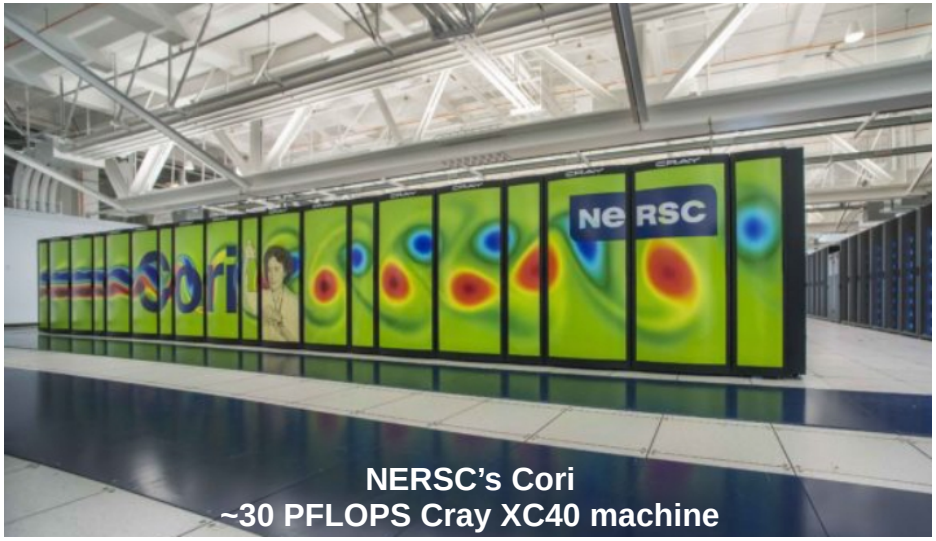

NERSC's Cori
~30 PFLOPS Cray XC40 machine

- Complete simulation of fields, beam dynamics, muon decays, and detector interactions in the storage ring
- Making use of HPC computing **@NERSC** **New!**
  - Via HEPCloud
  - First production completed with 10 B events

Challenges (1)

- **Challenge**:
  - Many datasets (~50), each of them has to pass through a 4-step procedure
- **Solution**:
  - Rolling production, with well defined steps and parallel execution of them
  - Requires good coordination and many people looking at different datasets every time

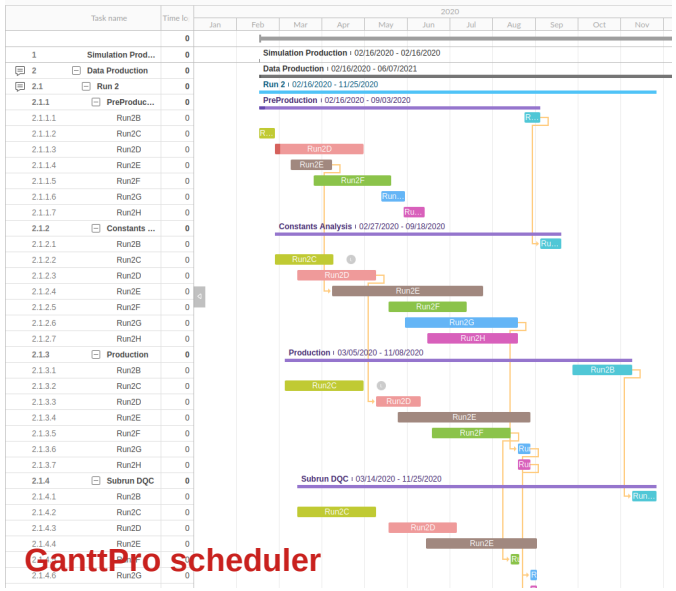

GanttPro scheduler

Rolling production scheme

Pre-production	Subset A	Subset B	Subset C	Subset D
Calibration		Subset A	Subset B	Subset C
Full production			Subset A	Subset B
DQC analysis				Subset A

Challenges (2)

- **Challenge**:
  - Too few production experts and too much data to process
- **Solution**:
  - Implemented a shift calendar similar to online-DAQ shifts
  - Whole collaboration is now involved with institution quotas
  - Big effort in making documentation and checklist user-friendly and optimized!

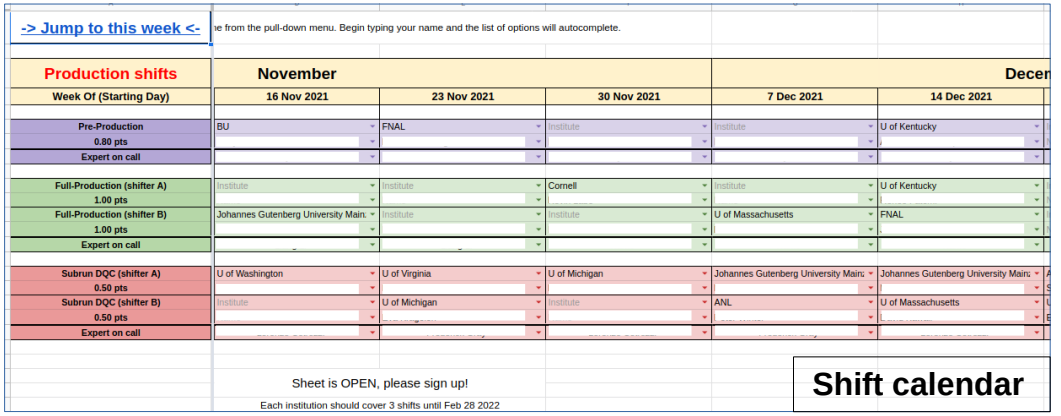

-> Jump to this week <-

Production shifts	November				Decer
Week Of (Starting Day)	16 Nov 2021	23 Nov 2021	30 Nov 2021	7 Dec 2021	14 Dec 2021
Pre-Production	BU	FNAL			U of Kentucky
0.80 pts					
Expert on call					
Full-Production (shifter A)			Cornell		U of Kentucky
1.00 pts					
Full-Production (shifter B)	Johannes Gutenberg University Main			U of Massachusetts	FNAL
1.00 pts					
Expert on call					
Subrun DQC (shifter A)	U of Washington	U of Virginia	U of Michigan	Johannes Gutenberg University Main	Johannes Gutenberg University Main
0.50 pts					
Subrun DQC (shifter B)		U of Michigan		ANL	U of Massachusetts
0.50 pts					
Expert on call					

Sheet is OPEN, please sign up!

Each institution should cover 3 shifts until Feb 28 2022

Shift calendar

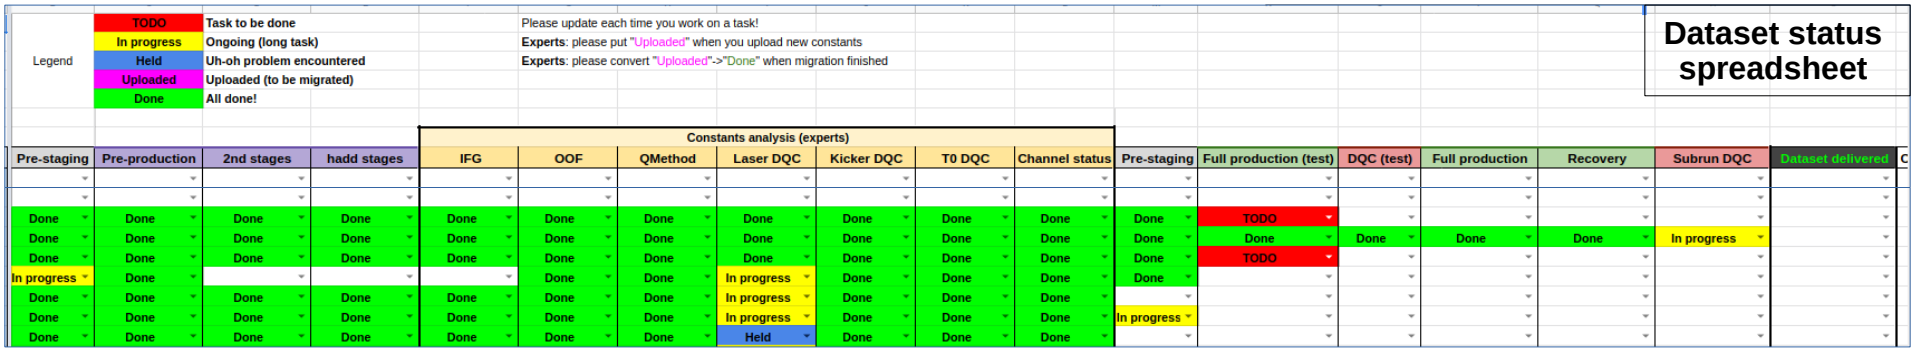

Legend		
	TODO	Task to be done
	In progress	Ongoing (long task)
	Held	Uh-oh problem encountered
	Uploaded	Uploaded (to be migrated)
	Done	All done!

Please update each time you work on a task!

Experts: please put "Uploaded" when you upload new constants

Experts: please convert "Uploaded"->"Done" when migration finished

Dataset status spreadsheet

				Constants analysis (experts)													
Pre-staging	Pre-production	2nd stages	hadd stages	IFG	OOF	QMethod	Laser DQC	Kicker DQC	T0 DQC	Channel status	Pre-staging	Full production (test)	DQC (test)	Full production	Recovery	Subrun DQC	Dataset delivered
Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	TODO					
Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	In progress	
Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	TODO					
In progress	Done				Done	Done	In progress	Done	Done	Done	Done						
Done	Done	Done	Done	Done	Done	Done	In progress	Done	Done	Done							
Done	Done	Done	Done	Done	Done	Done	In progress	Done	Done	Done	In progress						
Done	Done	Done	Done	Done	Done	Done	Held	Done	Done	Done							

Challenges (3)

- **Challenge**:
  - When an error is uncovered, a dataset is reprocessed using the full reconstruction
- **Solution**:
  - Expanded ***art*** to include data merger mechanism
  - Tracker data (which is computing intensive) will be processed separately and then merged with the production files
  - Ongoing progress

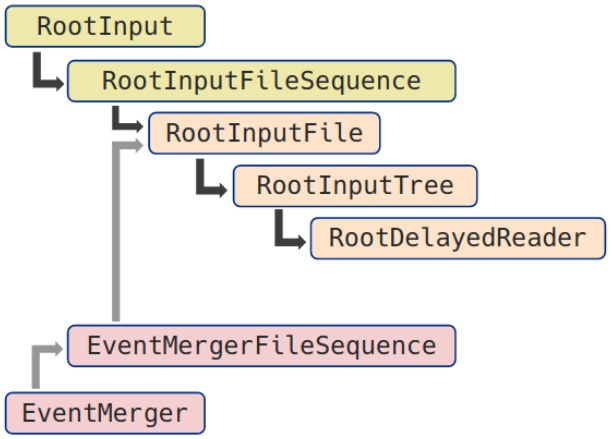

Challenges (4)

- **Challenge**:
  - All production data is backed up on tape, and fast-acces disks can’t hold all the data (many PBs!)
  - When files are needed, pre-staging process copies them to disks (dCache).
  - Experiments typically compete in a queue to get access to the common read disk pools
  - The system is not designed for peak loads
  - Can experience pre-staging delays during these peak periods
- **Solution**:
  - Working with Fermilab Scientific Computing Division to mitigate

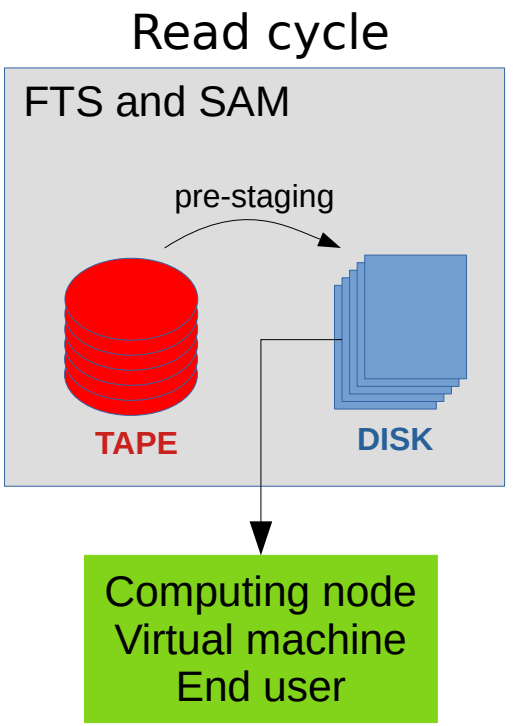

With the rolling production scheme, we need 5-6 datasets prestaged at all times!

Challenge (5)

- **Challenge**:
  - File bookkeeping and metadata is handled by SAM (Sequential Access via Metadata)
  - It's an evolution of the Fermilab Tevatron Run 2 storage system
  - SAM is unlikely to scale in the long term future of huge datasets
- **Solution** (proposed):
  - Upgrade to RUCIO
  - Modern, file-based, quotas and lifecycle management
  - "Manages multi-location data in a heterogeneous distributed environment"
  - No metadata catalog, working on integrating metadata into RUCIO to eventually completely replace SAM

Summary

- Muon g-2 is a relatively small experiment, but all the data must be produced and analyzed with precision and reliability
- A lot of improvements in the software since the first data acquisition (2017)
- Highlights:
  - ✔ Database integration
  - ✔ Data reconstruction shifts and Grid interface with POMS
  - ✔ Simulation running at HPC centers
- Some challenges with I/O and data handling but we are working on them
- Run 2-3-4 data under reconstruction and analysis

Thank you!

https://muon-g-2.fnal.gov/
pgirotti@fnal.gov