From physics to insights: an overview of CERN computing pipelines

Vincenzo Eduardo Padulano (CERN EP-SFT)

Valencia, November 2025

Founded in 1954
Biggest research centre for High Energy Physics (HEP)

CERN

Accelerating Science

A collaboration of many member states, including **Spain!**

A collaboration of many member states, including Spain!

[Source]

LHC:
Large Hadron Collider

The largest scientific experiment to date

Circumference:
27 Km

Depth:
-175 m ~ -50 m

Working temperature:
-273.1 C°

Raw physics data:
1000 TB/s

Research at CERN is focused on the Standard Model, defining how the building blocks of the universe interact

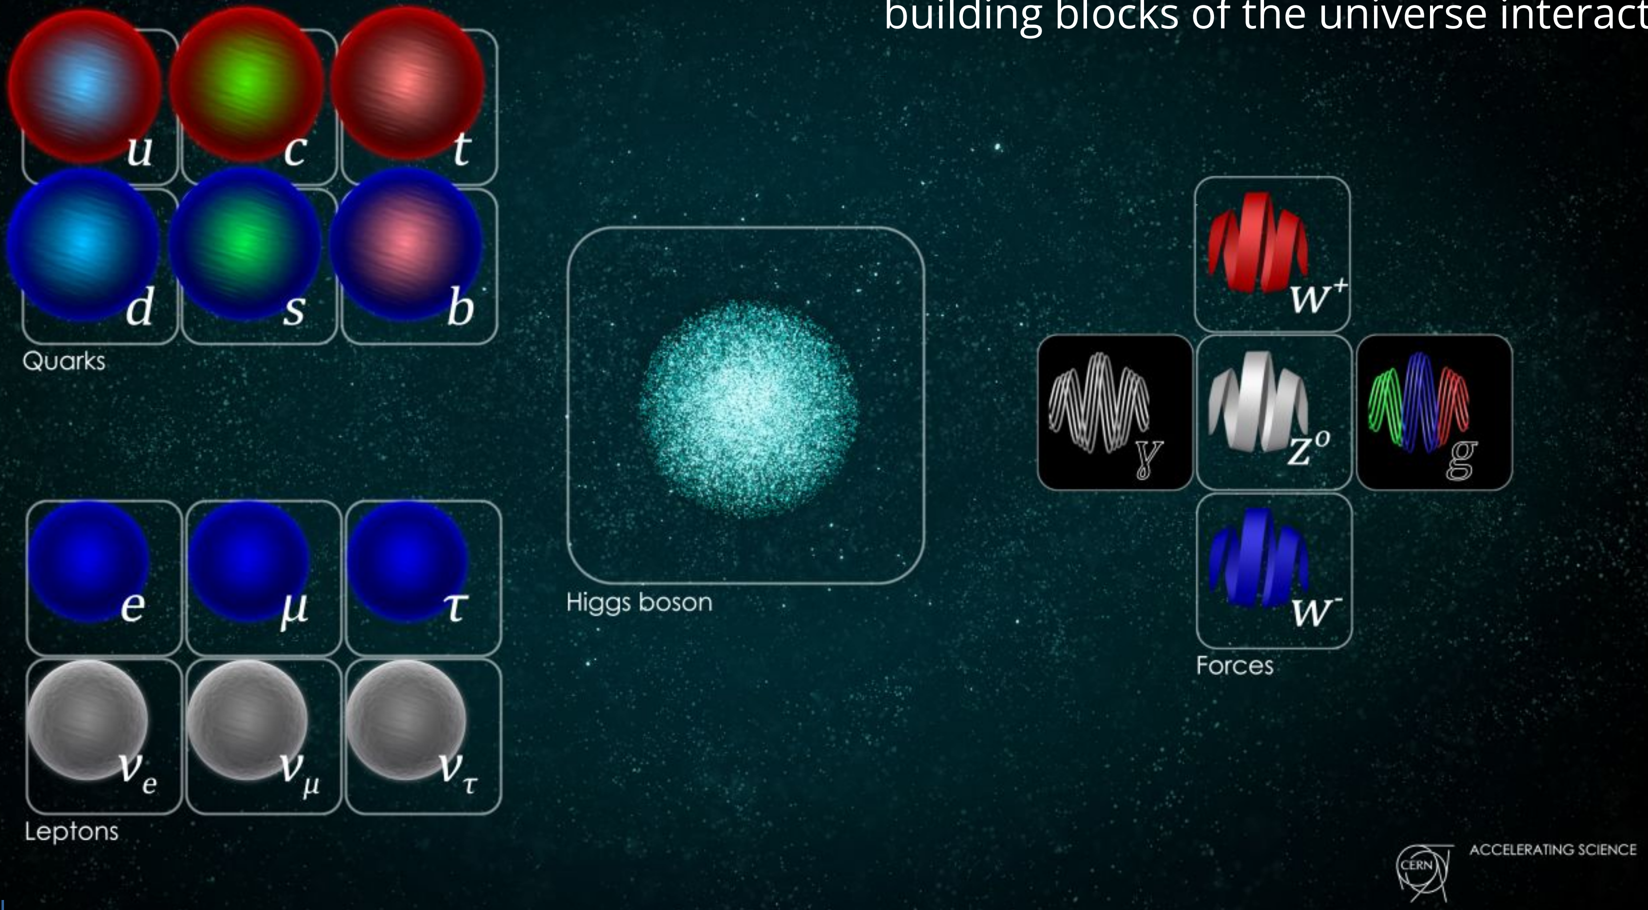

Source: CERN

Large particle detectors at CERN

ALICE

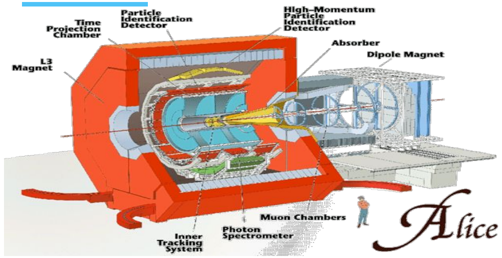

ATLAS

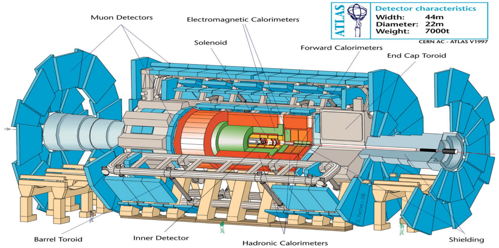

CMS

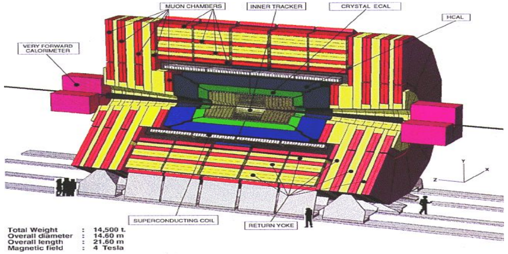

LHCb

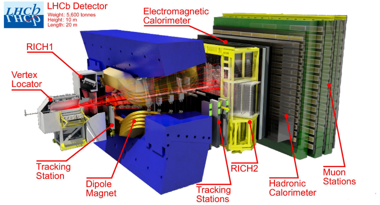

… but also large collaborations spread worldwide

Source

Source

The LHC project schedule

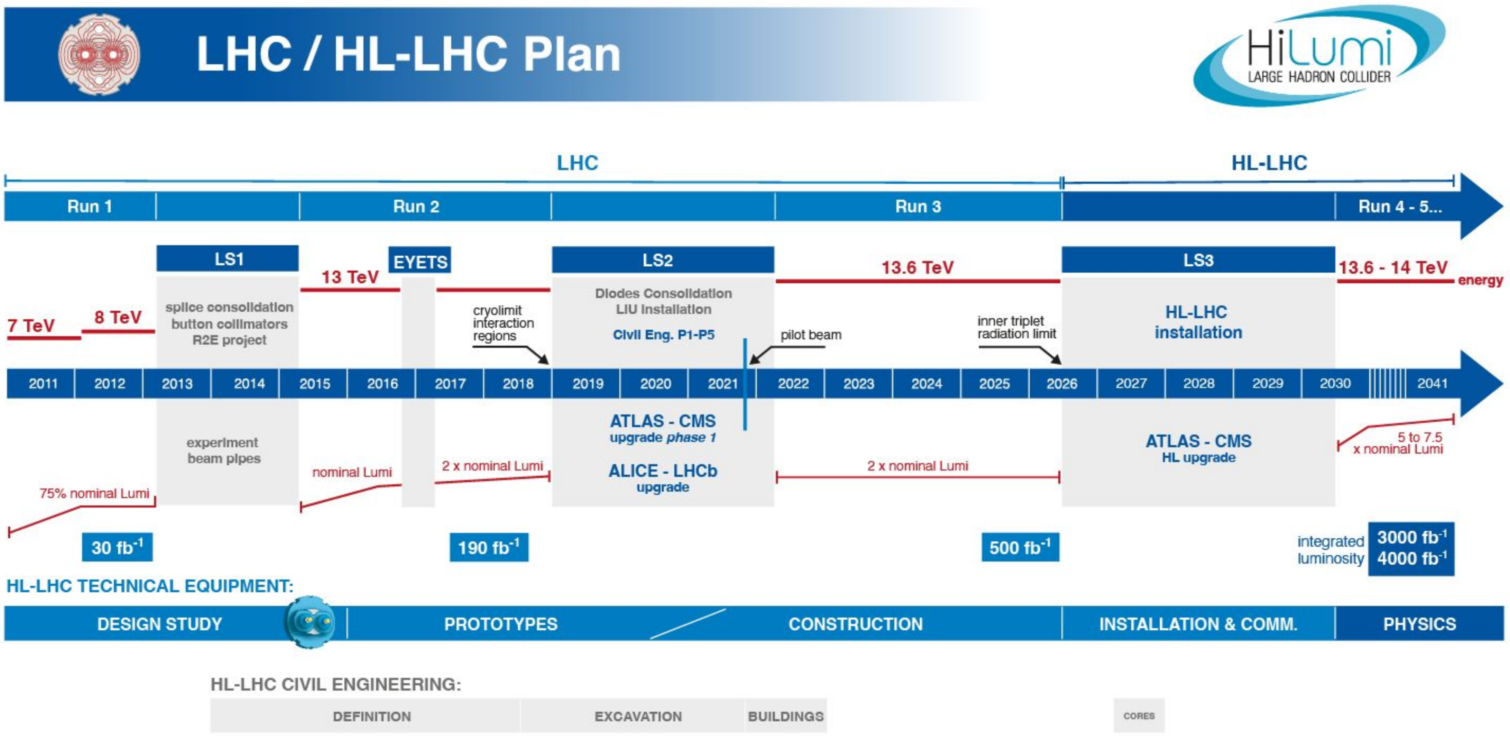

Data lifecycle at the LHC

Data Collection
1 PB/s

Trigger systems

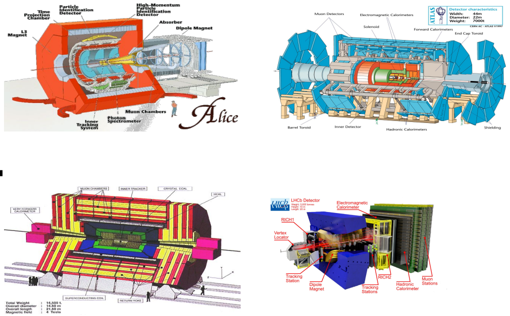

Raw data

Simulated data

Reconstruction

RECO

Processing, skimming

Analysis formats

Event selection, statistical treatment…

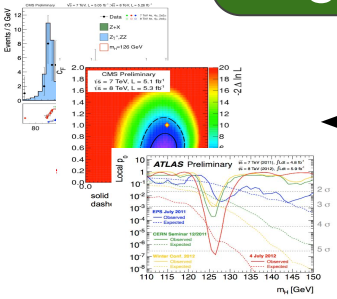

Data Analysis
O(10) GB → exploration
O(10) TB → full-scale

Offline Processing
several times a year
100 PB/y

Data lifecycle at the LHC

- Data acquisition: HEP data are statistically independent particle collisions (events)
- Analysis at experiments requires the highest possible number of collision events
  - We are looking for very, very rare events!
- One possible approach: more protons in the collisions
  - Thus, heavier computations
  - But! higher chance of uninteresting collisions

- Reconstruction (RECO): transform RAW data into understandable things
  - muons, electrons, photons...
- Further processing to create slimmer datasets
  - Usually called AOD, Analysis Datasets

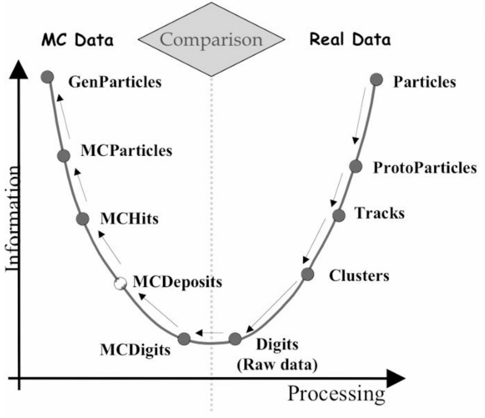

From RAW to analysis

- Reconstruction is done with many different specialized software
  - Depends on experiment (e.g. CMSSW)
- Has seen wide accelerator usage
  - Example from CMS
- AODs then see various skimming
  - experiment-dependent

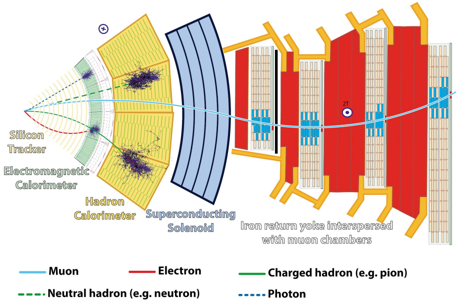

Source

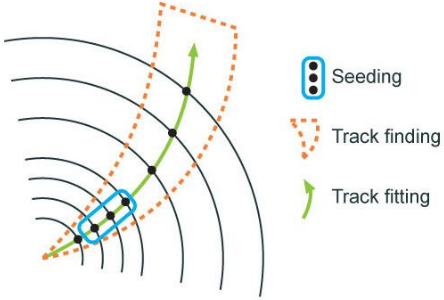

Source

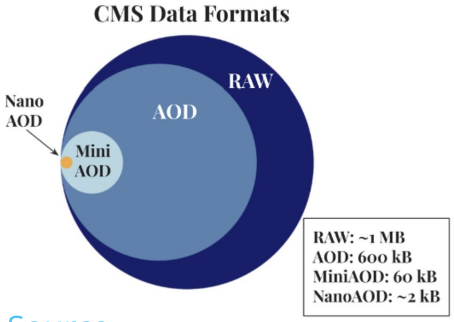

Source

- Generation (GEN): Create physics events (through Monte Carlo generators)
  - A set of particles
- Simulation: simulate response of detector based on generated events
  - Compare real physics collision with simulation from theory

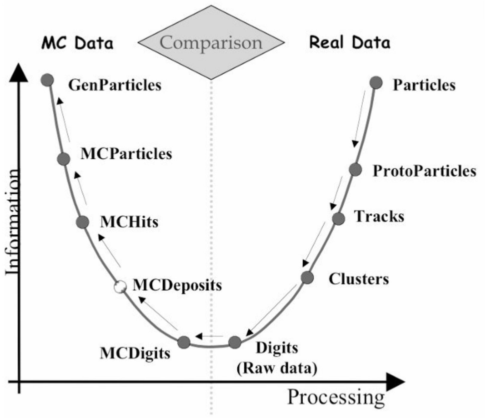

GEN and SIM software

- There are various **event generator** software
  - Pythia, MadGraph, SHERPA, …
- GEANT4: **state-of-the-art** software used for **full simulation** of particle transport through matter
- **Fast simulation**
  - Use a mix of ML techniques to run simulations
  - **Less precise** than full simulation with Geant4, **but faster**

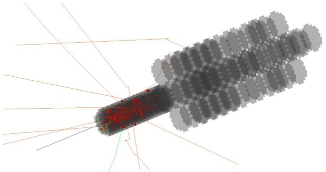

Source

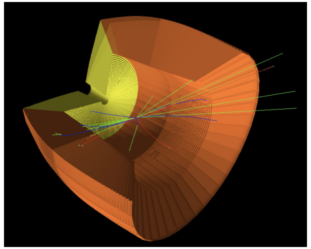

Source

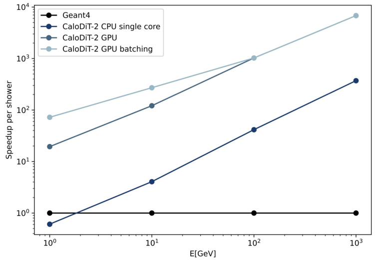

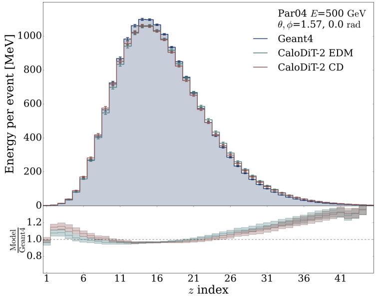

Raikwar et al.

Computing challenges

- More than **2 EB** physics data collected and stored so far
- That’s only 10% of the total dataset, **90% with HL-LHC**
- **No software R&D** means drastic **reduction** of future physics programme

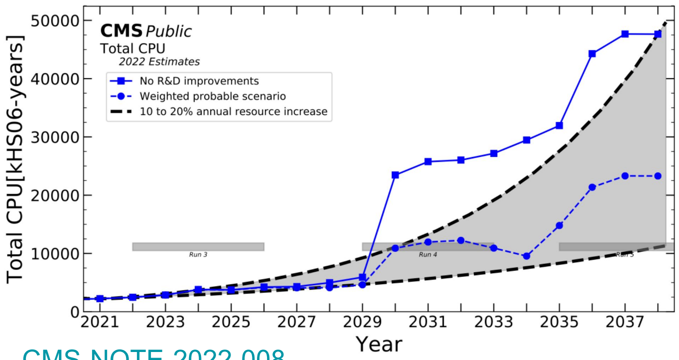

CMS-NOTE-2022-008

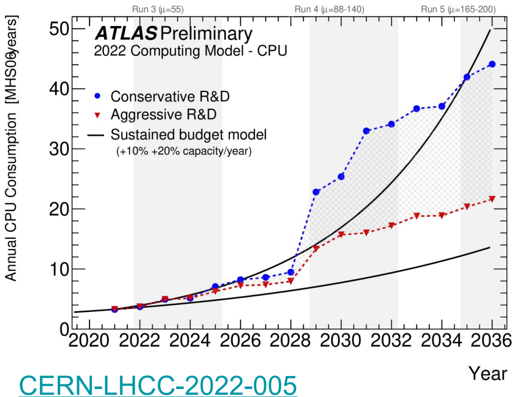

CERN-LHCC-2022-005

A computing dilemma

- Very large quantities of data collected
- Different experiments with different schemas and processing workflows
- Collaborations spread among research groups worldwide

How?

The Worldwide LHC Computing Grid

- Tier layers (Tier 0 at CERN)
- From large computing facilities to small university clusters
- All connected together.
- Physicists tap into the grid for practically all the steps of the lifecycle

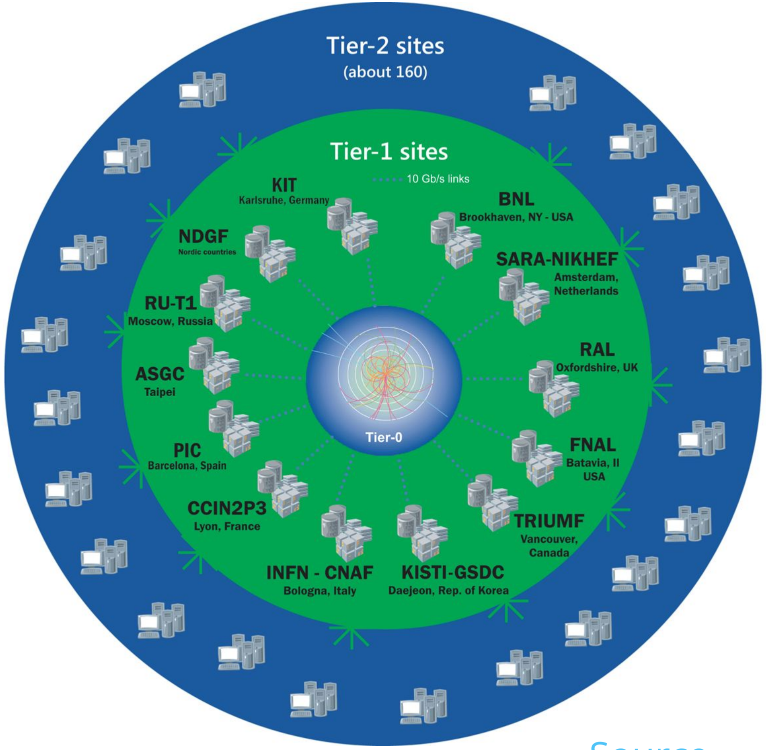

Source

Data access in HEP data analysis

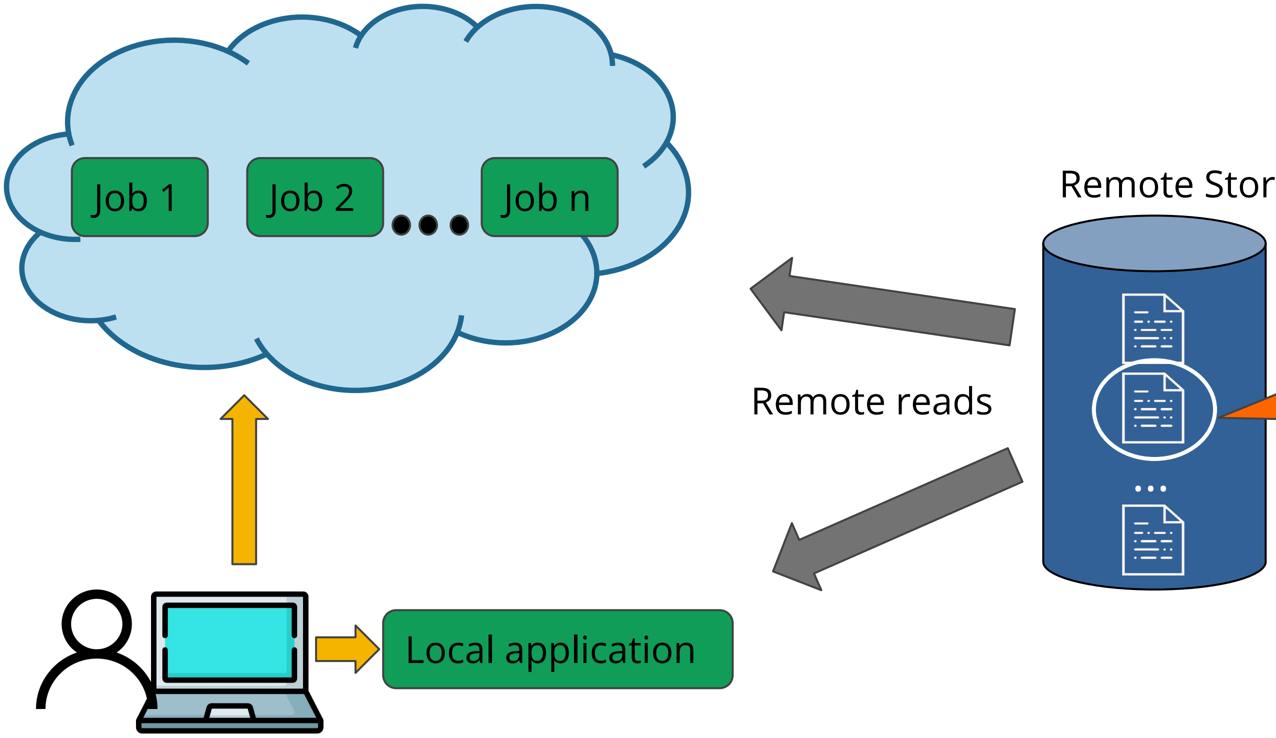

Sparse access:
One dataset can have 1000s columns

One analysis might need only few of them

Data access in HEP data analysis

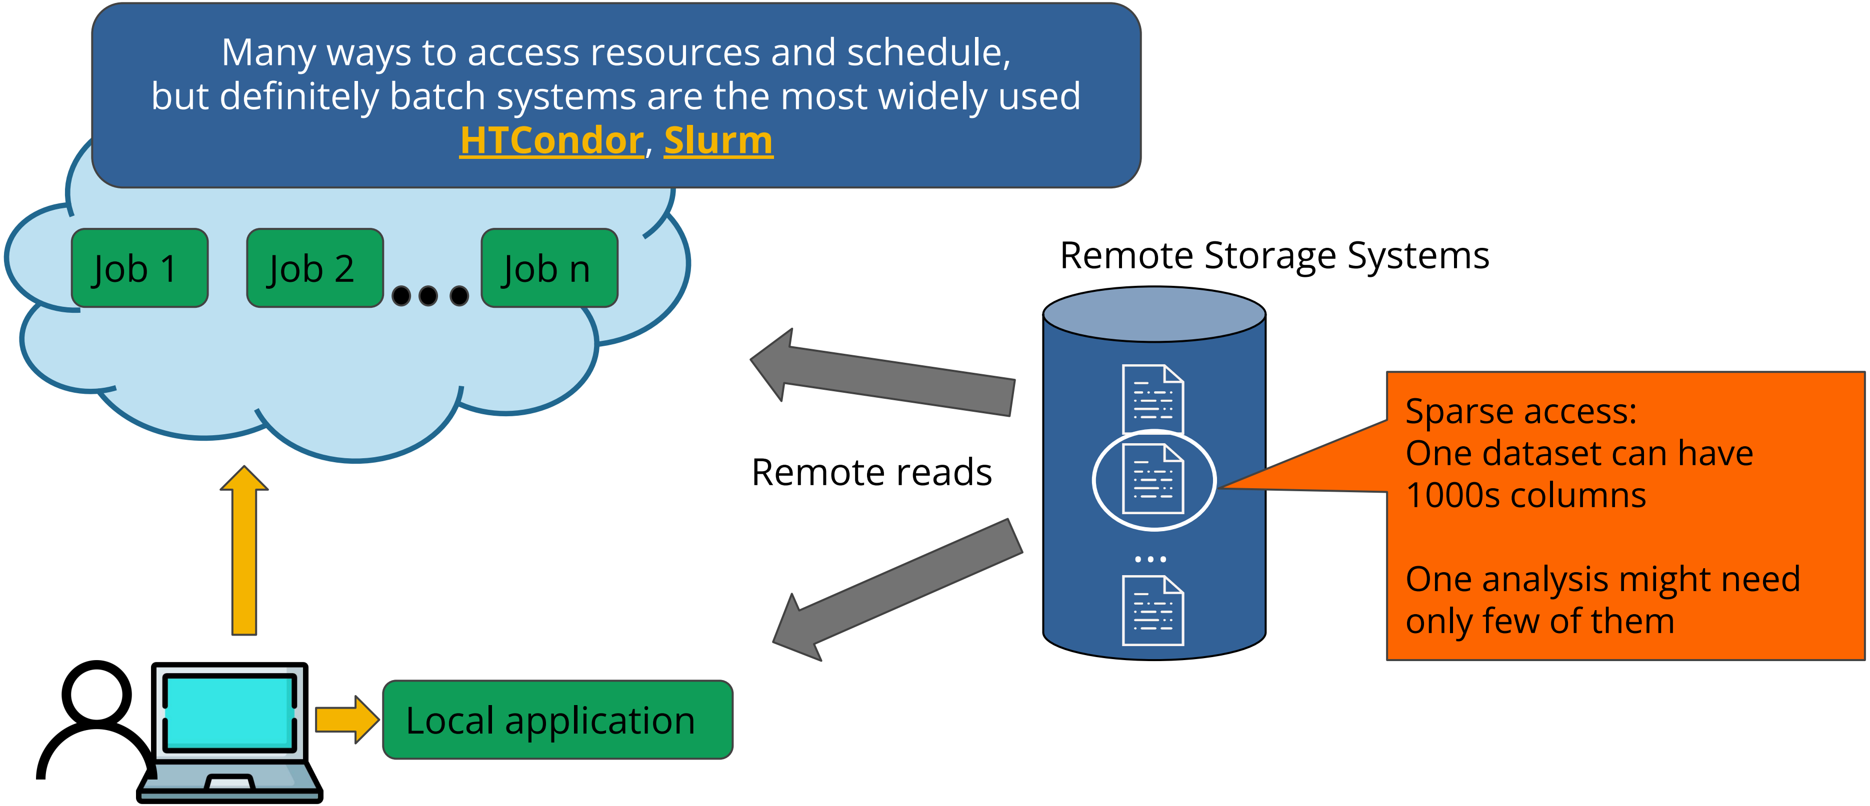

Data access in HEP data analysis

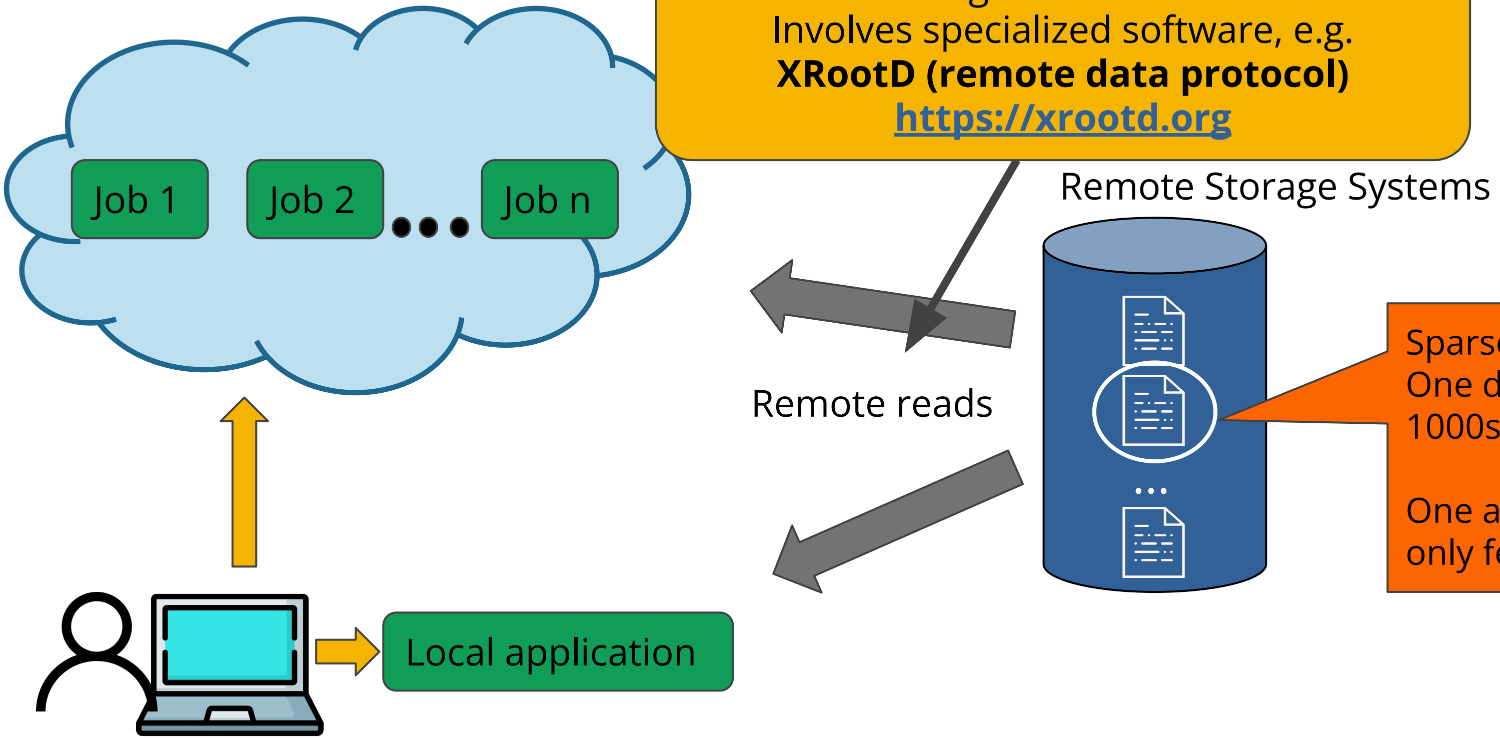

A big deal for HEP!
Involves specialized software, e.g.
XRootD (remote data protocol)
https://xrootd.org

Sparse access:
One dataset can have 1000s columns

One analysis might need only few of them

Data access in HEP data analysis

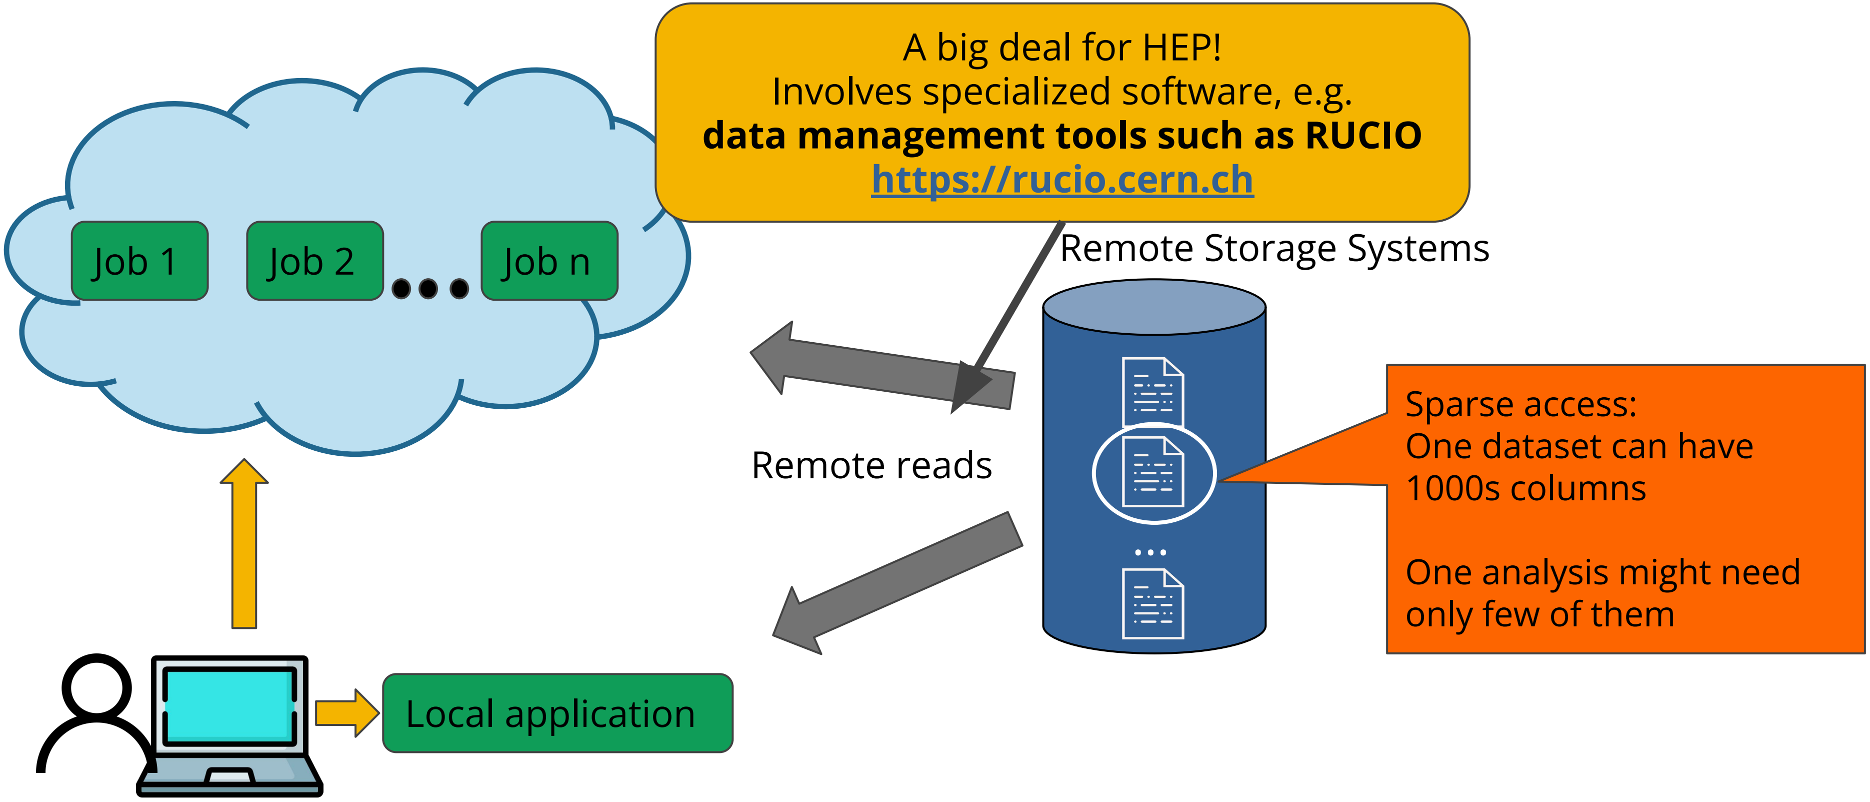

And what about software access?

- CVMFS: The **standard** HEP **software distribution** tool
- A shared filesystem with lazy on-demand access to files (libraries, scripts etc.)
- Aggressive caching for maximum performance
- Used by institutes worldwide

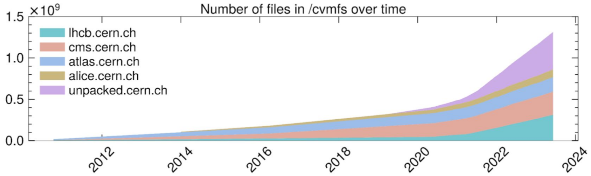

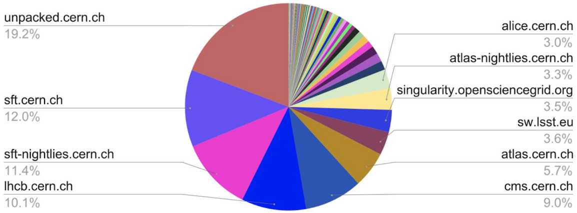

https://cernvm.cern.ch/fs/

HEP data processing

High Energy Physics analysis

Gather accelerator data already stored in analysis formats

Perform analysis:

- Select interesting events
- Transform physics measurements with complex functions
- Aggregate into relevant statistics

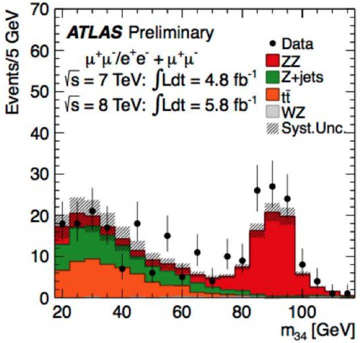

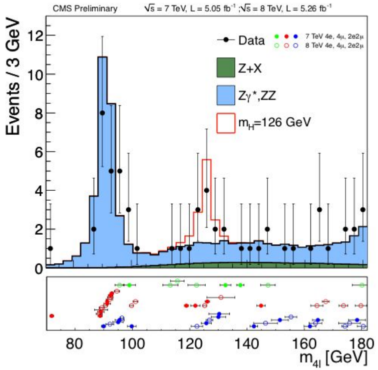

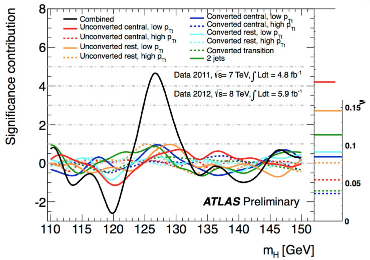

Various actors involved

Analysts:

- The end users
- (Mostly) physicists affiliated with some CERN experiment
- Interested in minimizing time between starting the application and seeing the final results
  - a.k.a "time to plot"

Framework developers:

- Develop software for the analysts
  - experiment-specific (CMSSW, ATHENA)
  - generic (ROOT, GEANT4)
- Striving for performance, user friendliness and HEP feature needs

Site admins:

- Responsible for computing clusters and infrastructures
- Worry about the software stack, resilience, hardware scalability

ROOT, a protagonist in HEP computing

- Storage, processing, analysis and visualisation of scientific data
- Widely adopted in **High Energy Physics** and in other scientific and industrial fields
  - **2 EB** data stored in ROOT data format
  - Fits and parameters estimations for discoveries (e.g. the Higgs)
  - Thousands of ROOT plots in scientific publications

root.cern

ROOT: An Open International Collaboration

CERN openlab · Google Summer of Code · iris hep · Bundesministerium für Forschung, Technologie und Raumfahrt

Open-source and Open-development

- On GitHub, LGPL 2.1
- PR-based model with a public review process
- Already 89 unique contributors in 2025!

Open-planning

- The Program of Work (PoW) can be influenced with active engagement and contributions!
- Quarterly **public** reports to check progress according to experiments' inputs → link to Q3 2025

Core ROOT team

The HEP Dataset

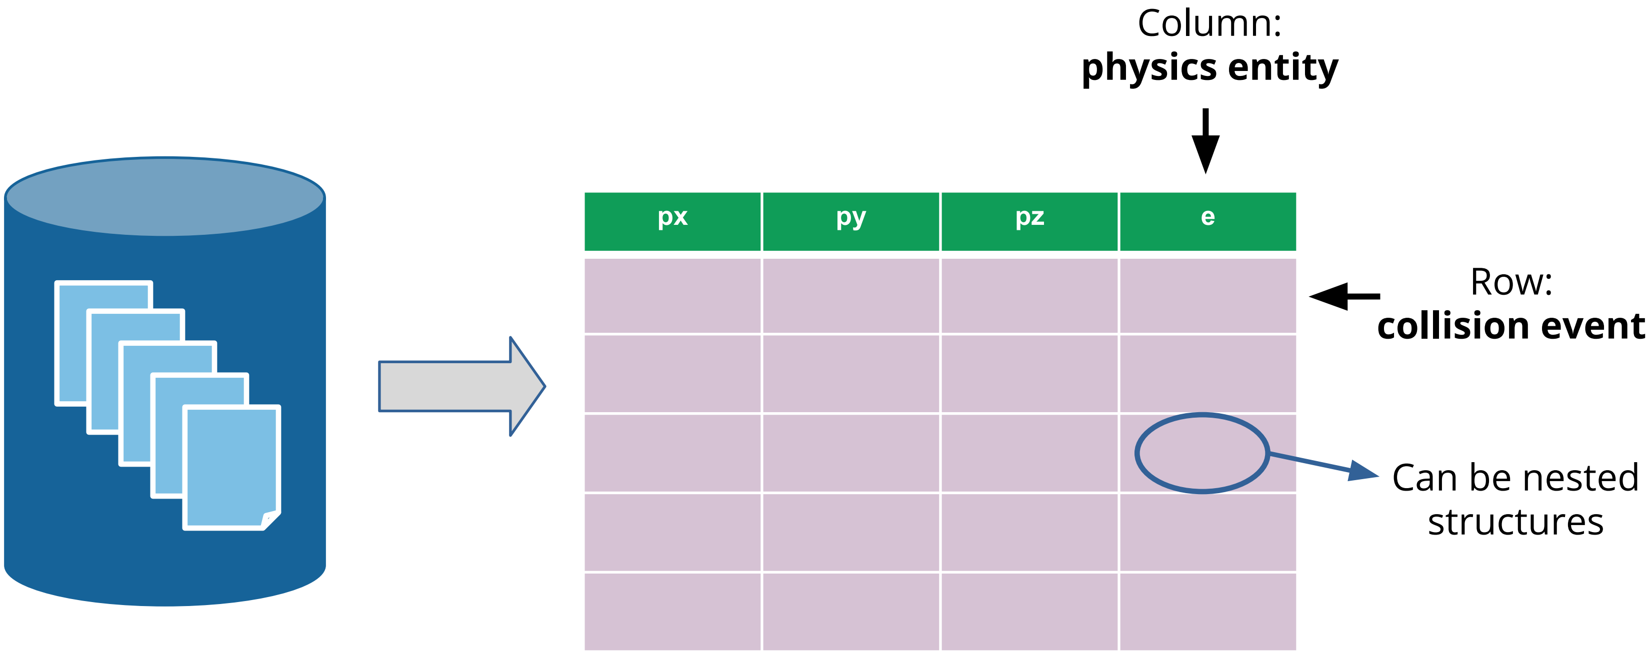

dataset stored in one or multiple files

The HEP Dataset

A relational DB, arbitrarily **nested**, with **varying sizes** per each DB entry

All of this **in each file**

Events
[ev0,ev1,ev2]
[ev0,ev1]

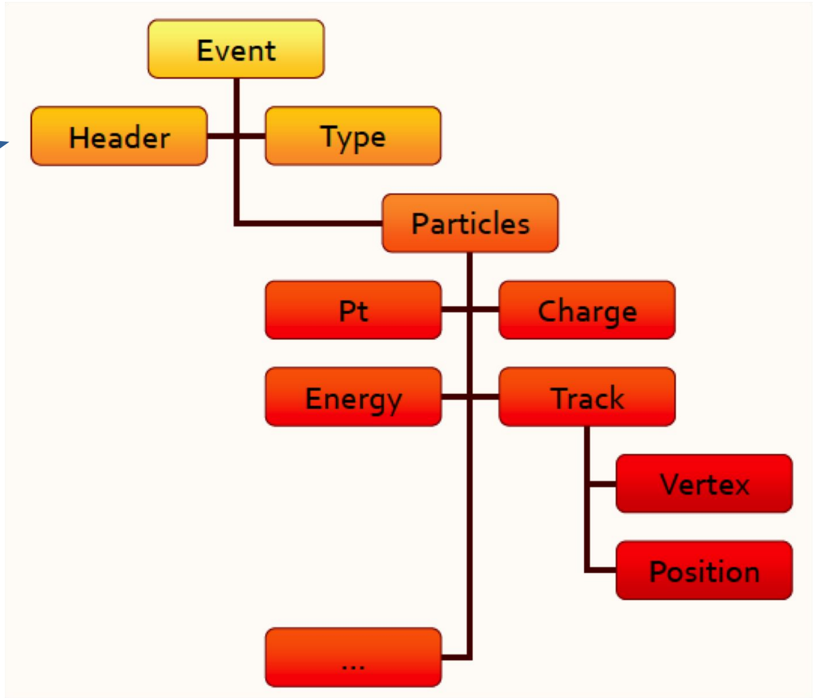

The HEP data format, today...

TTree is the data **format** for **HEP data** used in production by **all** LHC **experiments**:

- Columnar
- Extremely flexible
  - Write simple types, collections, arbitrary user-defined types
- Scalable
  - More than **2 EB** written so far

… and tomorrow

[Source]

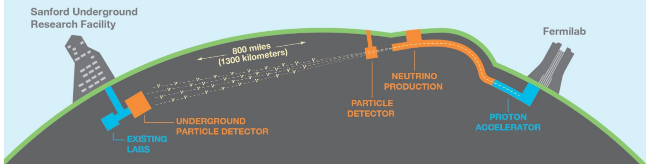

- Next-Generation Collider Experiments: **HL-LHC**, **DUNE**
  - **10-30** times the computing requirements
  - Single events in the **multi-gigabyte** range for DUNE, heavy-ion experiments
  - Real **analysis challenge** depends on several factors: number of events, analysis complexity, number of reruns, etc.

- Full exploitation of **modern storage hardware**
  - Ultra fast networks and **SSDs**: 10GB/s per device reachable (HDD: 250MB/s)
  - Flash storage is **inherently parallel** → asynchronous, parallel I/O key
  - Heterogeneous computing hardware → **GPU** should be able to load data directly from **SSD**, e.g. to feed ML pipeline
  - Distributed storage systems move from **POSIX** to **object stores**

RNTuple

Redesigned **I/O** subsystem, based on **25+ years** of **TTree** experience

- **Less** disk and CPU **usage**
- **Efficient** support of **modern hardware**
- Transparent **file-less** storage
- **Covering** all of today's **TTree** use cases
- **Binary format** defined in a dedicated specification

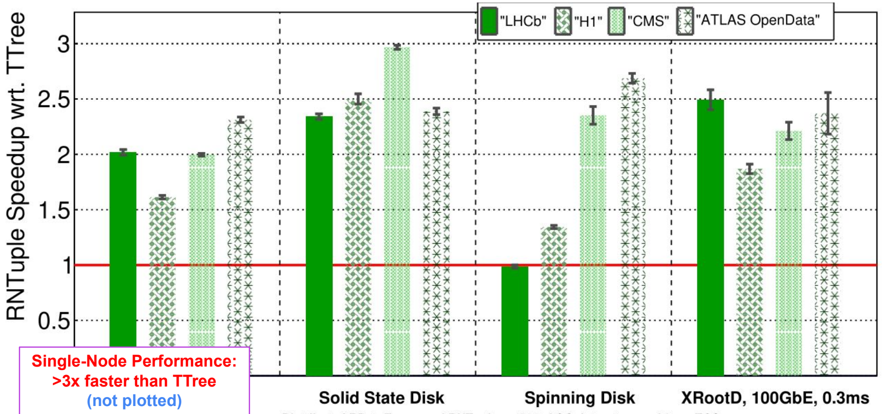

Distributed RDataFrame and RNTuple on 200xAGC dataset served from EOS

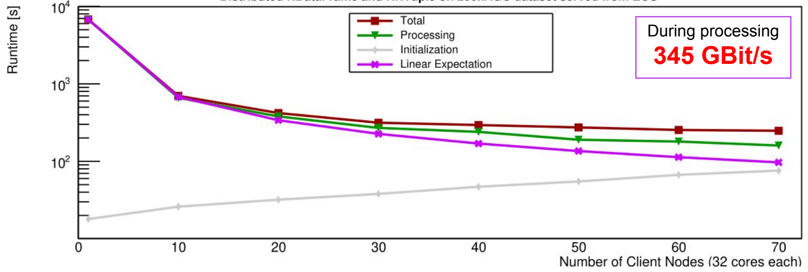

RDataFrame: ROOT's declarative interface for analysis

- ROOT's high level interface for analysis

- Supports native parallel execution
- multi-threaded and distributed
- Exploit CMS compute, today

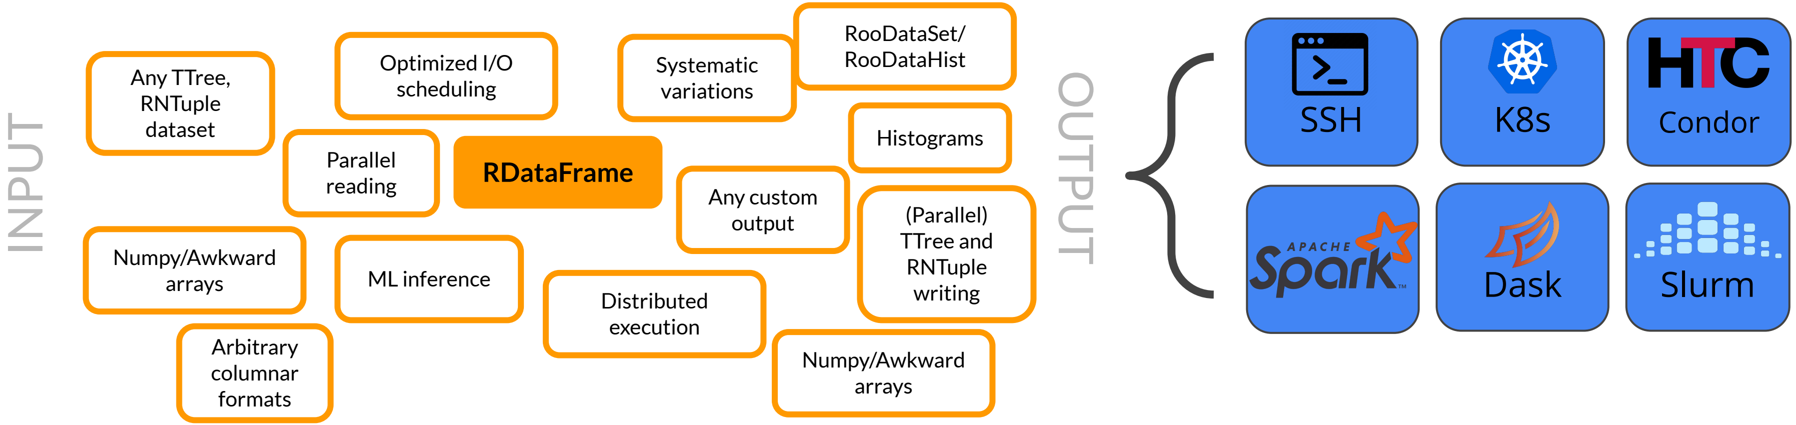

RDataFrame: ROOT's declarative interface for analysis

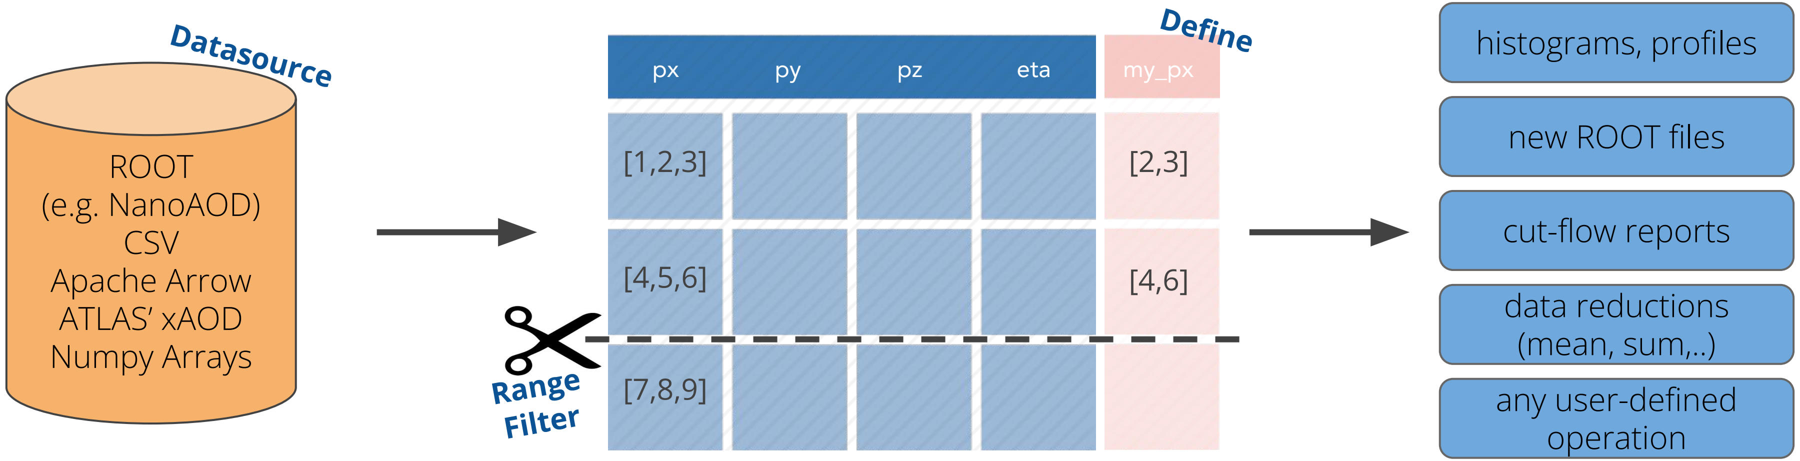

A small RDataFrame analysis

```python
import ROOT
df = ROOT.RDataFrame('Events', 'Run2012BC_DoubleMuParked_Muons.root')

df_mass = df.Filter('nMuon == 2')\
            .Filter('Muon_charge[0] != Muon_charge[1]')\
            .Define('Dimuon_mass', 'InvariantMass(pt, eta, phi, mass)')

h = df_mass.Histo1D(
('Dimuon_mass', 'Dimuon mass;m_{#mu#mu} (GeV);N_{Events}', 30000, 0.25, 300), 'Dimuon_mass')

h.Draw()
```

A small RDataFrame analysis

```
import ROOT
df = ROOT.RDataFrame('Events', 'Run2012BC_DoubleMuParked_Muons.root')

df                                    )\
                                      [0] != Muon_charge[1]')\
                                      , 'InvariantMass(pt, eta, phi, mass)')


h
('Dimuon_mass', 'Dimuon mass;m_{#mu#mu} (GeV);N_{Events}', 30000, 0.25, 300), 'Dimuon_mass')

h.Draw()
```

Open the dataset

A small RDataFrame analysis

```
import ROOT
df = ROOT.RDataFrame('Events', 'Run2012BC_DoubleMuParked_Muons.root')

df_mass = df.Filter('nMuon == 2')\
            .Filter('Muon_charge[0] != Muon_charge[1]')\
            .Define('Dimuon_mass', 'InvariantMass(pt, eta, phi, mass)')

h = df_mass
('  uon ma                    );N_{Events}', 30000, 0.25, 300), 'Dimuon_mass')

h.Draw()
```

Analysis steps:

1. Select events with exactly two muons
2. Select muons with opposite charge
3. Compute invariant mass

```
import ROOT
df = ROOT.RDataFrame('Events', 'Run2012BC_DoubleMuParked_Muons.root')

df_mass = df.Filter('nMuon == 2')\
            .Filter('Muon_charge[0] != Muon_charge[1]')\
            .Define('Dimuon_mass', 'InvariantMass(pt, eta, phi, mass)')

h = df_mass.Histo1D(
('Dimuon_mass', 'Dimuon mass;m_{#                              uon_mass')

h.Draw()
```

Aggregate result into histogram

A small RDataFrame analysis

```
import ROOT
df = ROOT.RDataFrame('Ever

df_mass = df.Filter('nMuor
            .Filter('Muon_
            .Define('Dimuc

h = df_mass.Histo1D(
('Dimuon_mass', 'Dimuon_ma                                  'Dimuon_mass')

h.Draw()
```

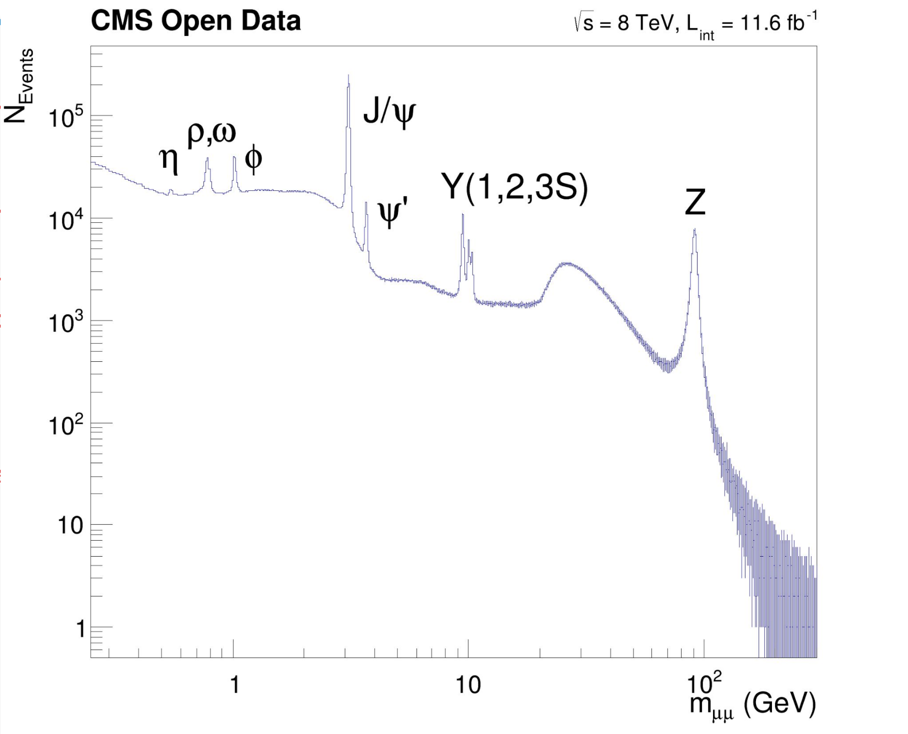

Outreach and collaborations

We teach ROOT data analysis:

- At CERN, for the summer students
  - Around 200 students per year
- Online on various occasions
- At universities, on demand in collaboration with interested groups

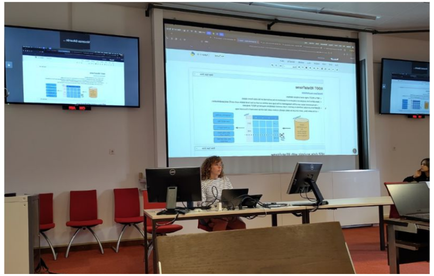

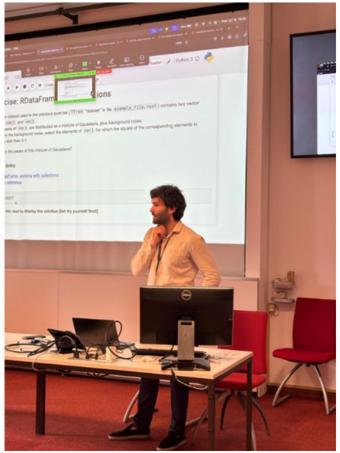

Collaborate with Us

- ~100 unique contributors per year, 10% of the commits in `24 authored outside of the core dev team
- **A single, motivated individual can make the difference, irrespective of own level of experience!**

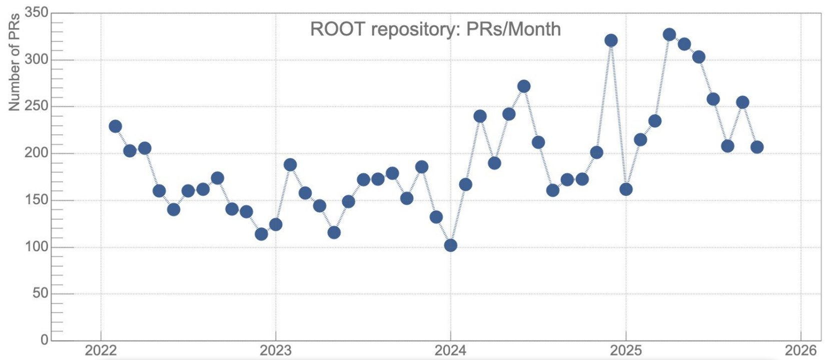

http://root.cern/open_projects

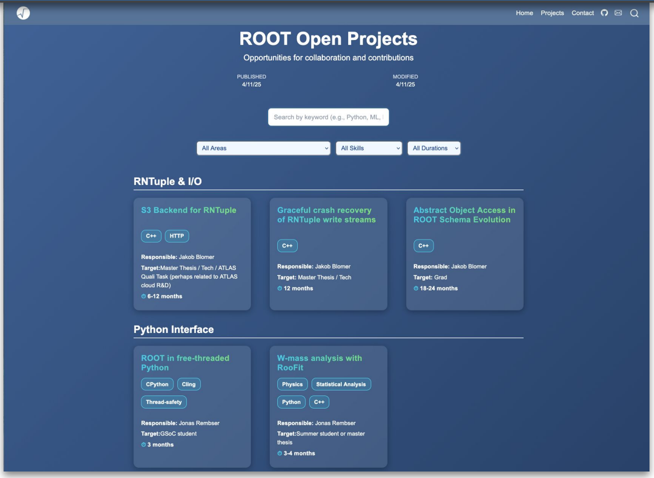

How to start?

- We mark the easiest issues in the GH tracker with good first issue
- Selecting from the list of open projects
- Simple contribution guidelines are available

Conclusions

- CERN is a leading laboratory for High Energy Physics
- A rich scientific programme, with hard computing challenges
- ROOT is a tool for storage, processing, analysis and visualisation of physics data
  - A protagonist in the software landscape of the field
- Collaboration is possible and encouraged!
  - Contact us if you're interested in one of our student projects