Sprucing and Analysis Productions: Offline data processing in LHCb without the pain

CHEP 2024

N. Skidmore on behalf of the LHCb experiment

Oct 2024

The LHCb data challenge

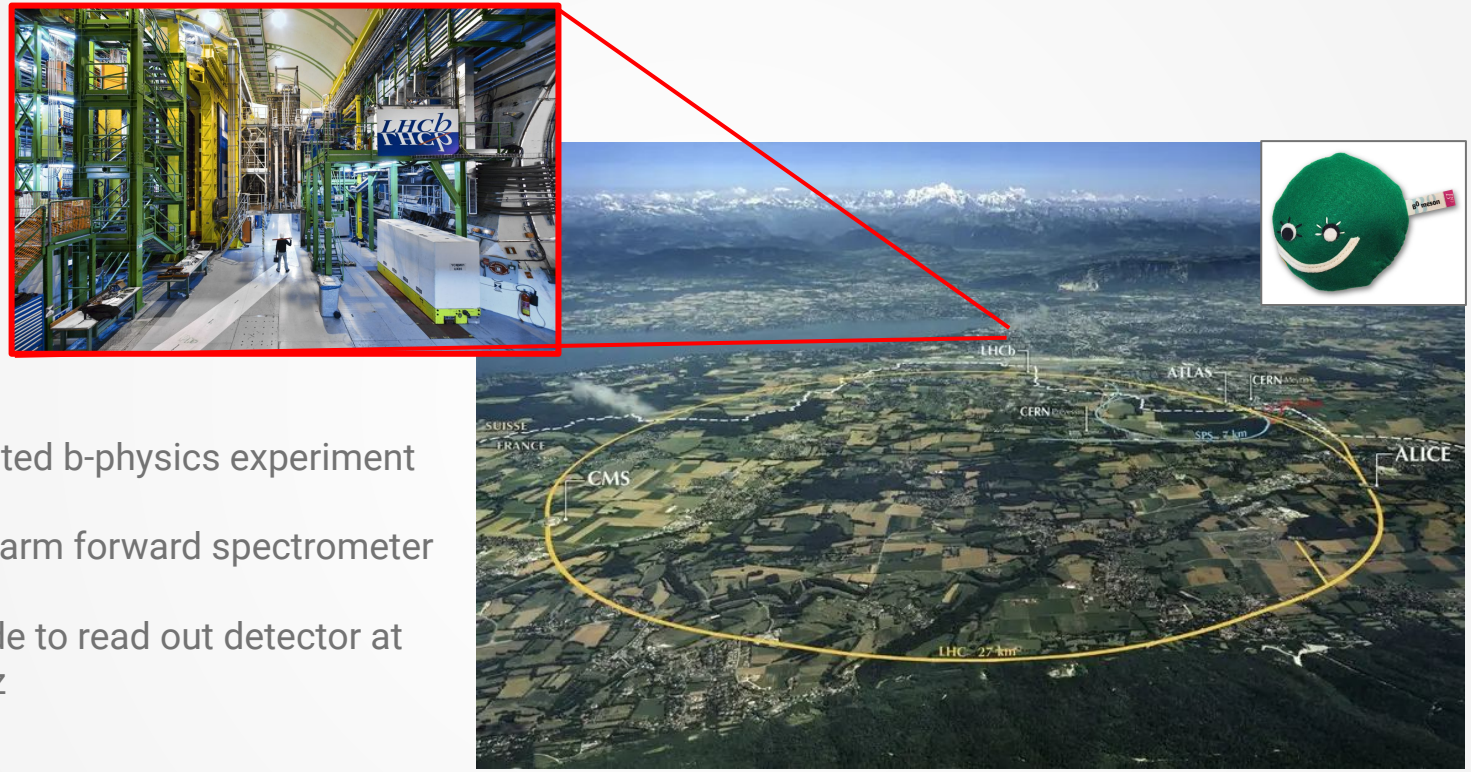

- Dedicated b-physics experiment
- Single-arm forward spectrometer
- Upgrade to read out detector at 30MHz

The LHCb data challenge

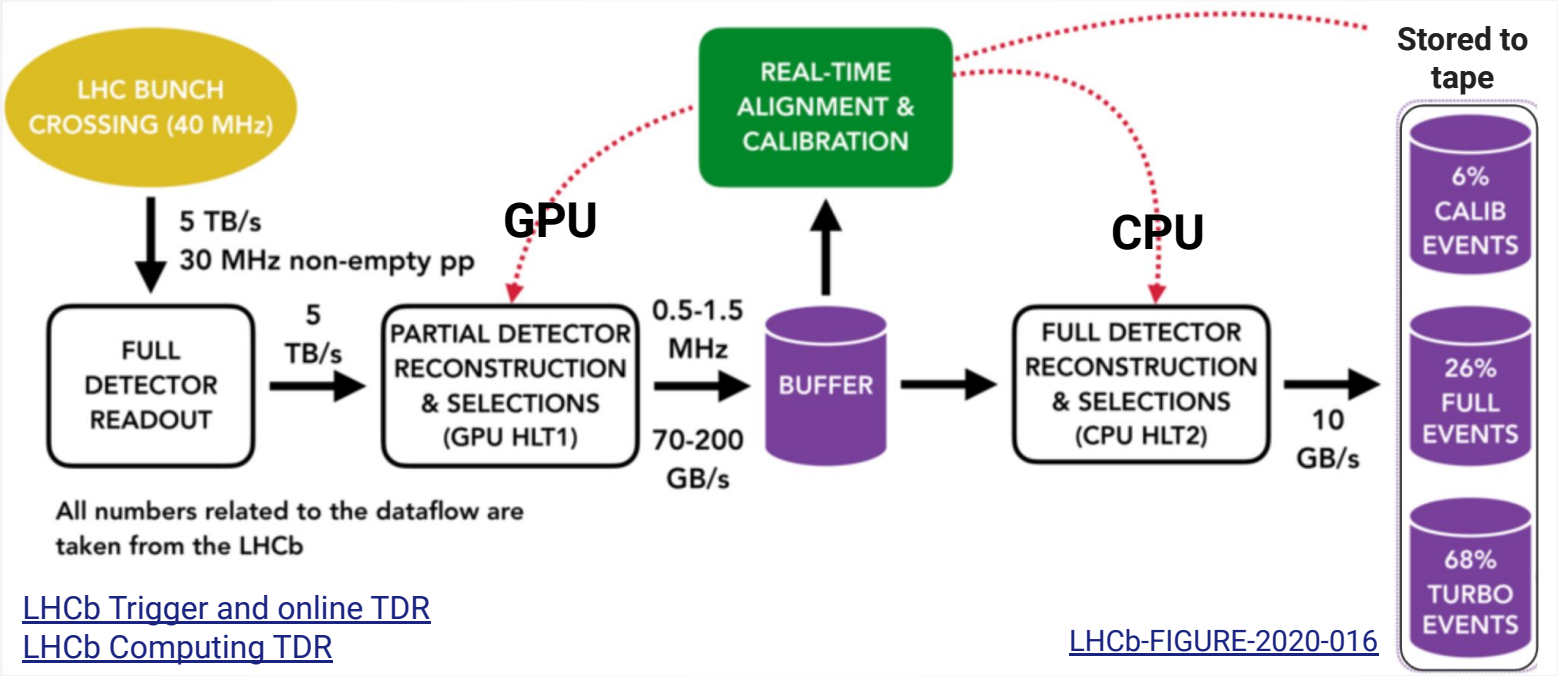

LHCb Trigger and online TDR
LHCb Computing TDR

LHCb-FIGURE-2020-016

See Alessandro's talk for more on LHCb's HLT

The LHCb data challenge

As of last week…

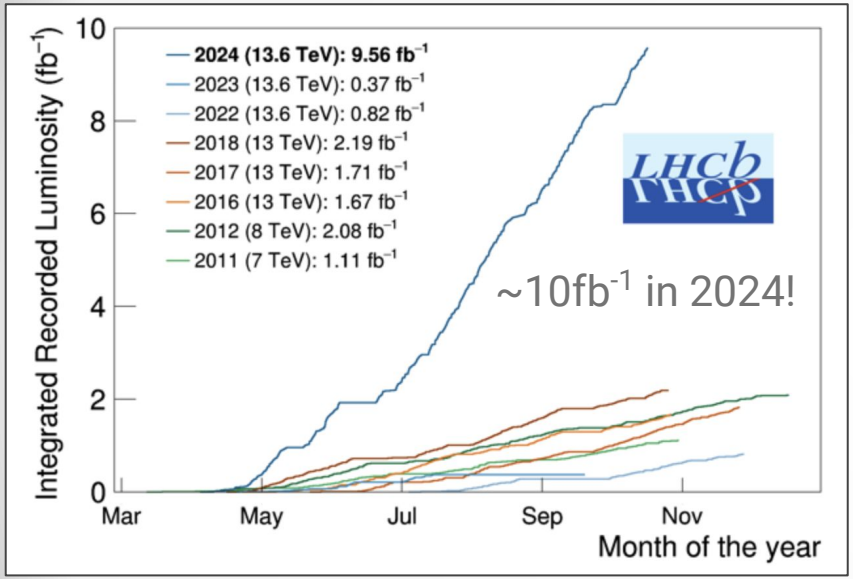

COLLABORATION DETECTOR LATEST POSTS OTHER PHYSICS RESULTS

End of successful proton-proton collision data taking period

By pietrzyk

OCT 18, 2024 #data analysis, #data taking, #LHCb detector, #RTA

The 2024 proton-proton collision period at the LHC ended this week. This period was particularly successful for LHCb, also owing to the excellent performance of the LHC. The rate at which the experiment was able to acquire integrated luminosity was spectacularly higher than in previous years, as shown in the luminosity plots presented in this page. LHCb has largely exceeded its target of an integrated luminosity of 9 fb^{-1}, which is higher than the sum of luminosities from Run 1 and Run 2. LHCb operated at a record instantaneous luminosity of $2\times10^{33}cm^{-2}s^{-1}$, corresponding to μ=5.3, the average number of visible pp collisions per bunch crossing. Not only did LHCb accumulate a record amount of data, but the efficiency of beauty and charm particle selection for hadronic channels has increased by a factor between 2 and 3 compared to previous years, due

Public announcement

The LHCb data challenge

The offline LHCb data challenge

- LHCb has been saving 10GB of data to tape every LHC second
- Data necessarily has a complex journey before it reaches analysts
- For Run 3+ LHCb revolutionised this process for the benefit of the analyst

Typical LHCb analyst circa 2018

Modern day LHCb analyst

The offline LHCb dataflow

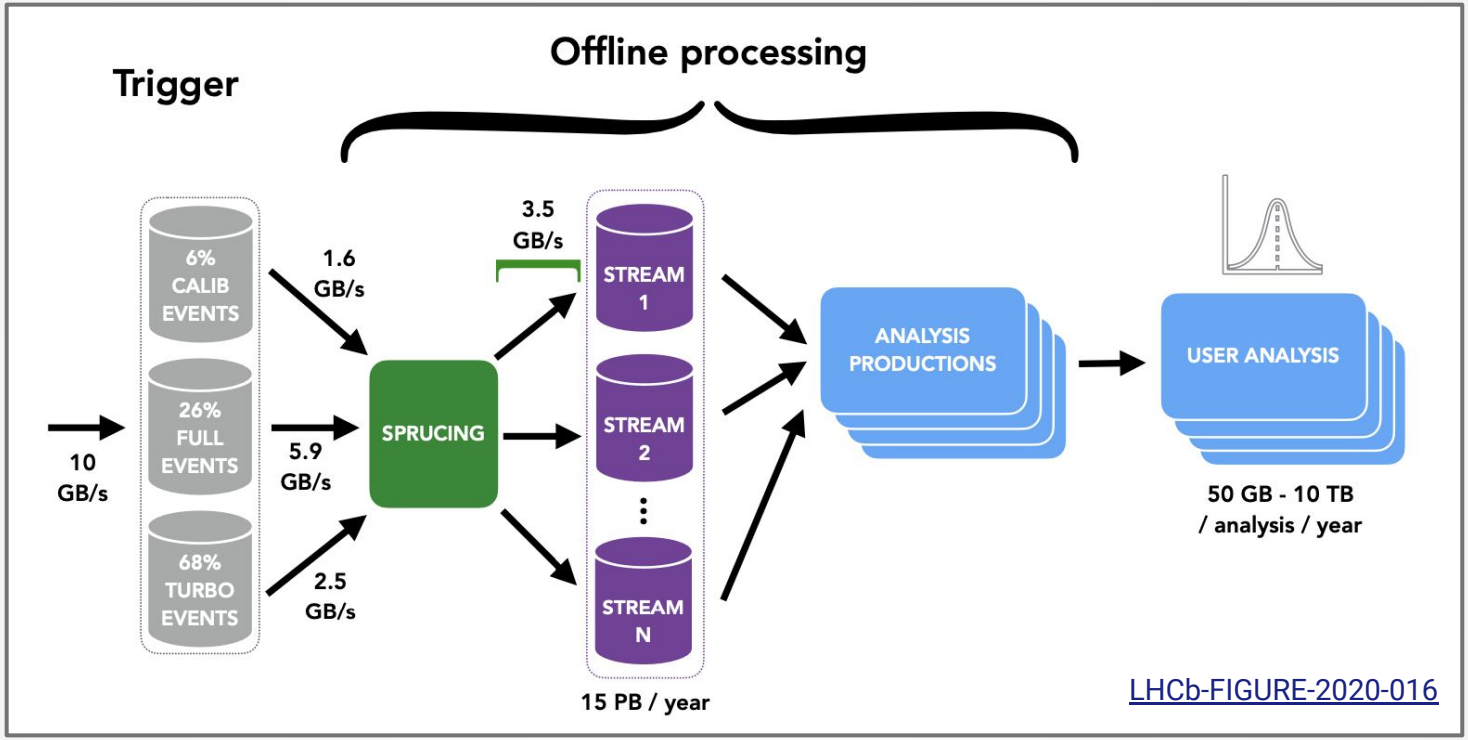

Sprucing - managing tape/disk persistency

Pure Turbo

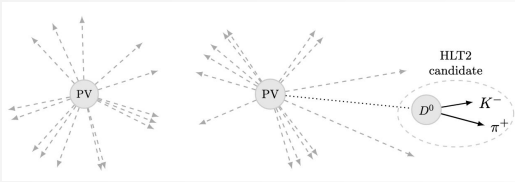

Full - Persistreco + RAW

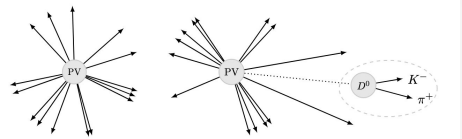

Rich Muon ...

Few kB ~ **Selective persistence** ~ **250kB**

Pure Turbo — **Full event**

Turbo SP

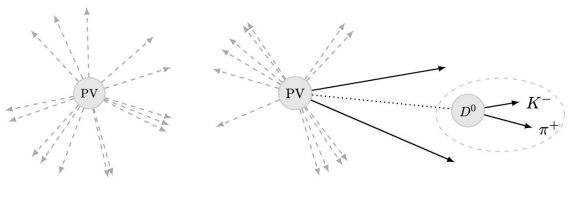

Full - Persistreco

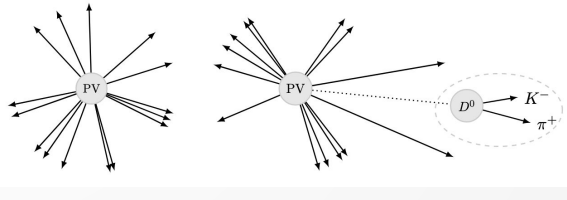

Sprucing - how to manage tape/disk persistency

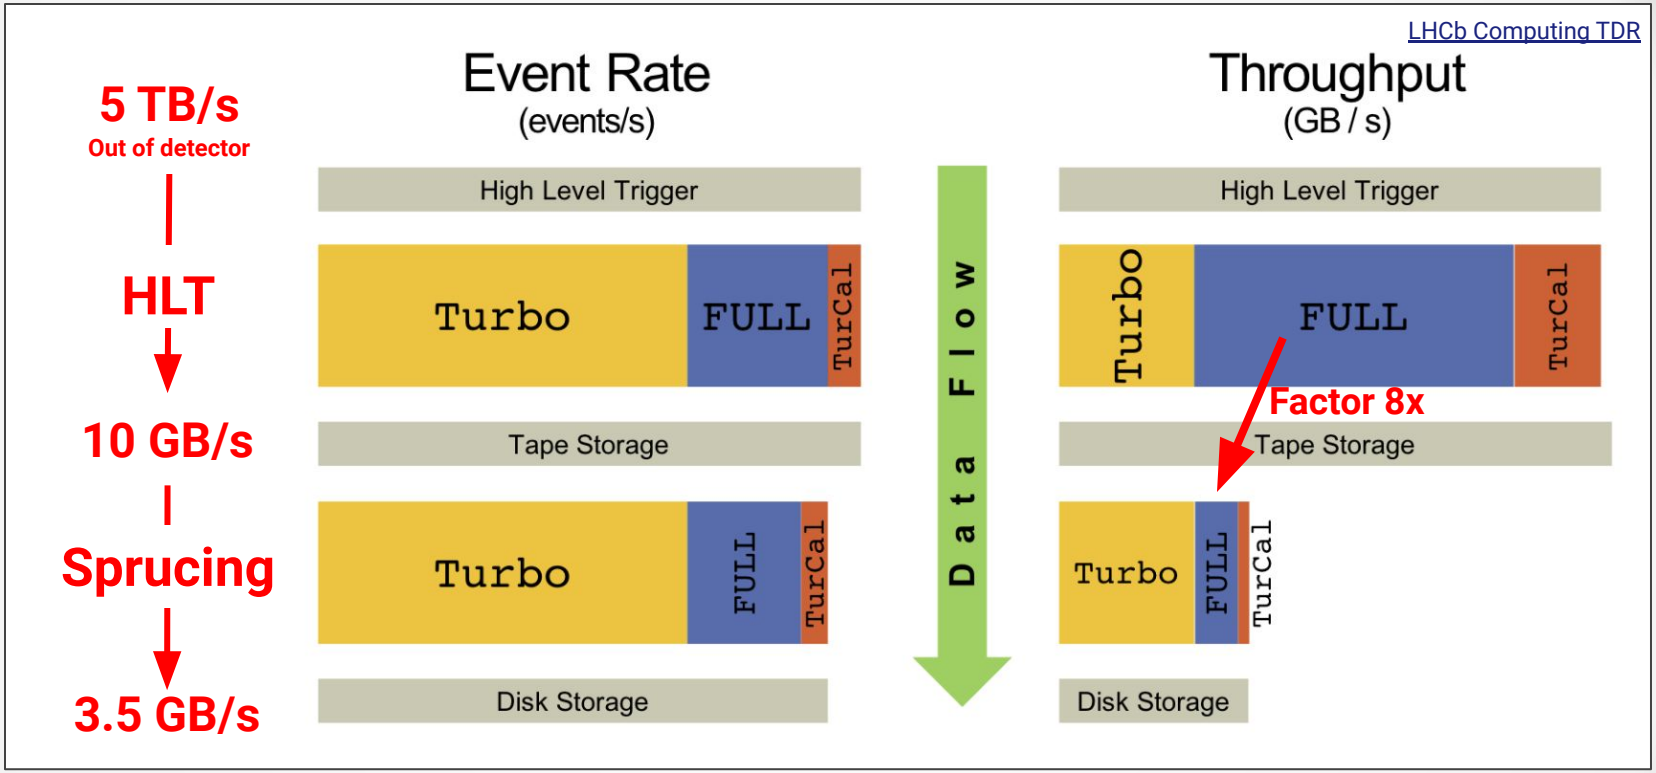

Can keep inclusively selected full events on tape for future exploitation in yearly re-sprucing campaigns

Sprucing - managing tape/disk persistency

LHCb spruced **35PB** of data in 2024 over 3 physics streams and 2 technical streams

- Bandwidth reductions met
- Sprucing on FULL stream requires **5x less CPU per event** compared to the Run 2 analogue

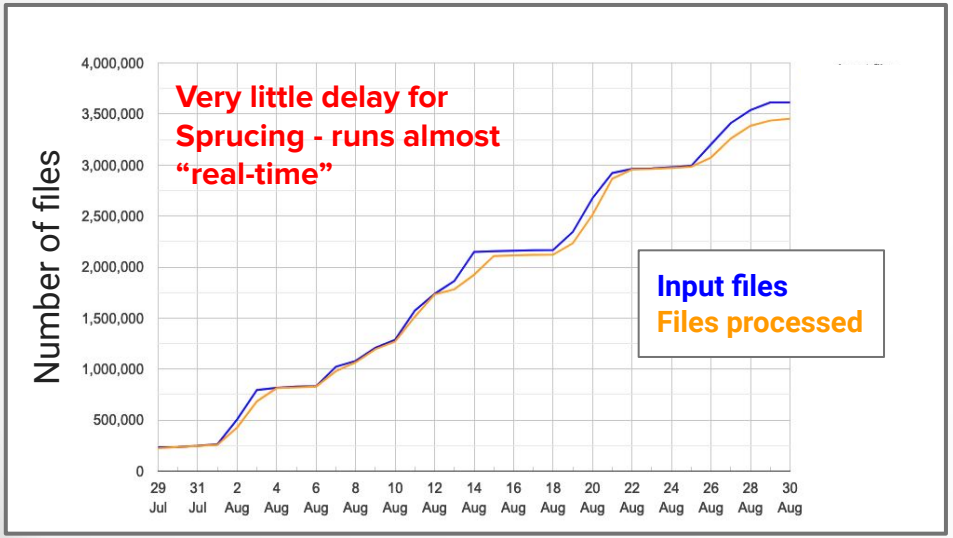

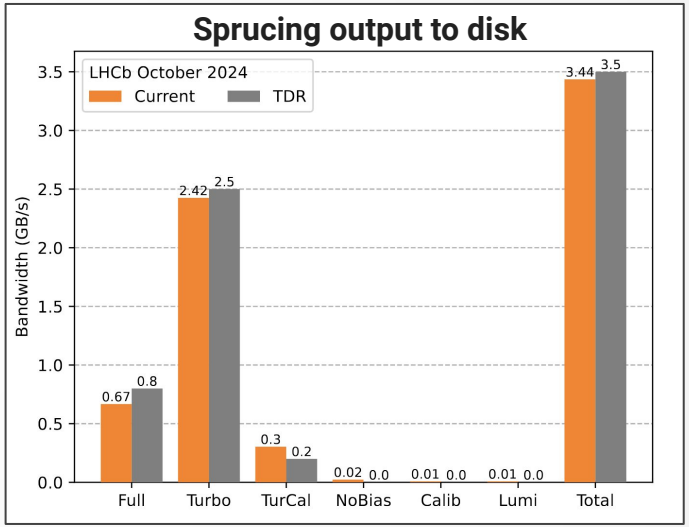

See Ross' talk for more on LHCb BW

Analysis productions - declarative ntupling

Pre - run 3 analysts made their own ntuples

THE PROBLEM

- Submitting, monitoring and error handling O(10,000) grid jobs
- No data provenance
- Thousands of failing grid jobs

⇒ BIG barrier between analysts and data

Analysis productions - declarative ntupling

THE SOLUTION ⇒ Analysis productions

- Centralise and automate ntuple creation
  ⇒ Saves countless analyst-hours
- Exploit DIRAC transformation system
  ⇒ Full data provenance
- Full job testing on GitLab CI
  ⇒ No buggy jobs on grid

Simple yaml job configuration

```yaml
defaults:
  application: DaVinci/v64r10@x86_64_v2-el9-clang16-opt
  output: DATA.ROOT
  options:
    entrypoint: bs2dspi_run3.dv_simple:alg_config
    extra_options:
      input_raw_format: 0.5
      input_type: ROOT
      simulation: False
      data_type: "Upgrade"
      geometry_version: run3/trunk
      conditions_version: master
      input_process: "TurboPass"
      input_stream: "b2oc"
  inform:
    -
  wg: B2OC

{%- set datasets = [
 ('2024Data', 'MagDown', '24c2'),
 ('2024Data', 'MagUp', '24c2'),
```

What application to run

Job options

Data to run on

Use Jinja templating to “render” the YAML

See Abhijit’s talk for DaVinci’s Funtuple that runs in APs

Analysis productions - declarative ntupling

Comprehensive job testing through GitLab pipelines

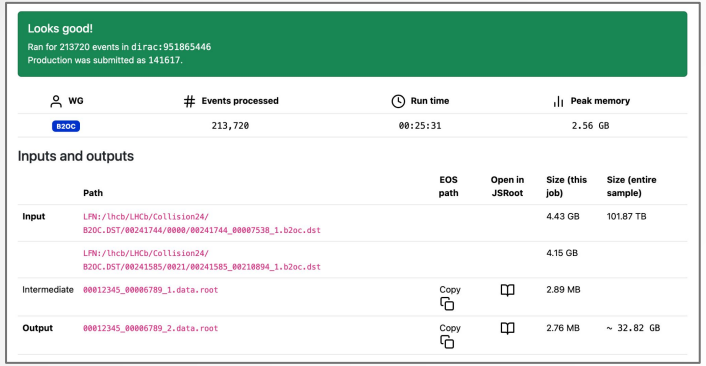

Reporting on memory usage

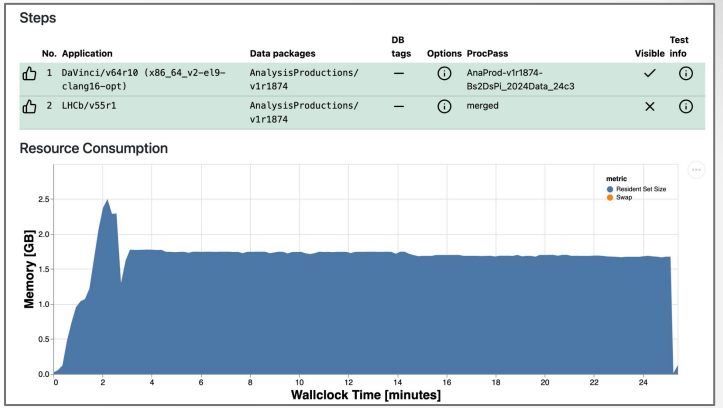

Reporting on estimated output size

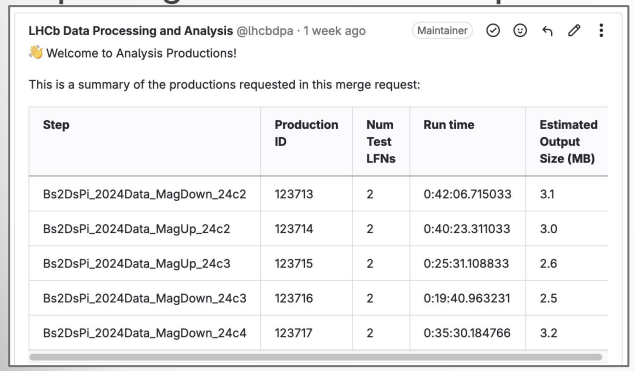

LHCb Data Processing and Analysis @lhcbdpa · 1 week ago Maintainer

👋 Welcome to Analysis Productions!

This is a summary of the productions requested in this merge request:

Step	Production ID	Num Test LFNs	Run time	Estimated Output Size (MB)
Bs2DsPi_2024Data_MagDown_24c2	123713	2	0:42:06.715033	3.1
Bs2DsPi_2024Data_MagUp_24c2	123714	2	0:40:23.311033	3.0
Bs2DsPi_2024Data_MagUp_24c3	123715	2	0:25:31.108833	2.6
Bs2DsPi_2024Data_MagDown_24c3	123716	2	0:19:40.963231	2.5
Bs2DsPi_2024Data_MagDown_24c4	123717	2	0:35:30.184766	3.2

Interactive logs with warning/error highlighting

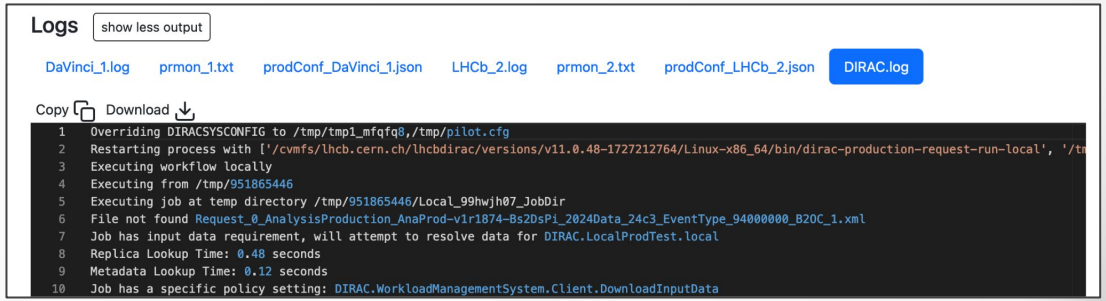

Analysis productions - declarative ntupling

Full data provenance with datasets tagged by analysis

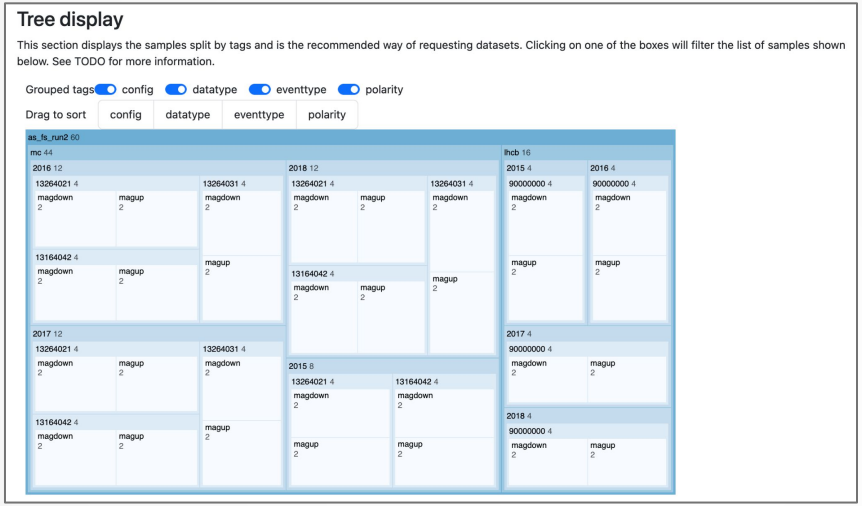

`apd` python packages allows for easy data file retrieval. Snakemake integrations!

```python
from apd import AnalysisData

datasets = AnalysisData("b2oc", "bs2dspi_run3")
bs2dspi_2024data_magdown_24c2_pfns = datasets(polarity="magdown", eventtype="94000000", datatype="2024")
```

Analysis productions - declarative ntupling

Full adoption of analysis productions at LHCb

- Over 1200 Run 3 APs have been submitted so far
- 700+ “live” APs picking up data as it was Spruced
  - Analysts have been looking at data tuples days after it was recorded by detector
- We are making amazing use of the WLCG!

Analysis productions also facilitate MC productions
See Emir’s talk

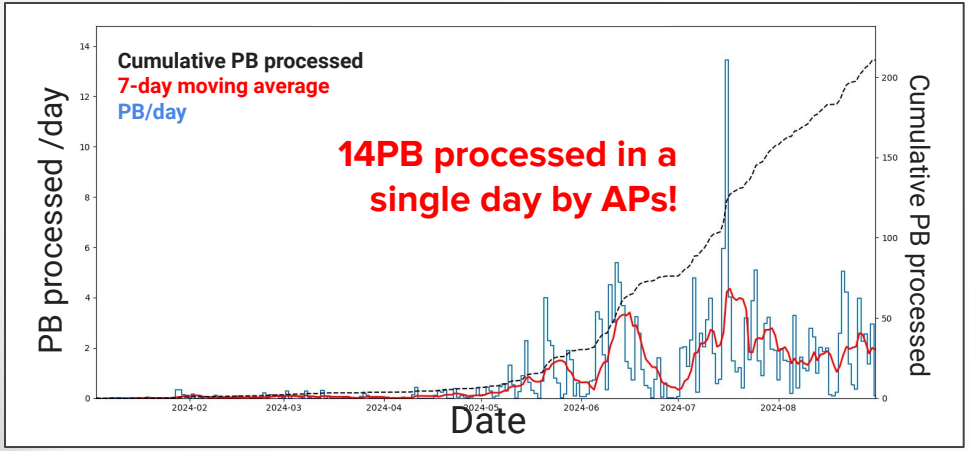

PB processed per day by APs with moving average

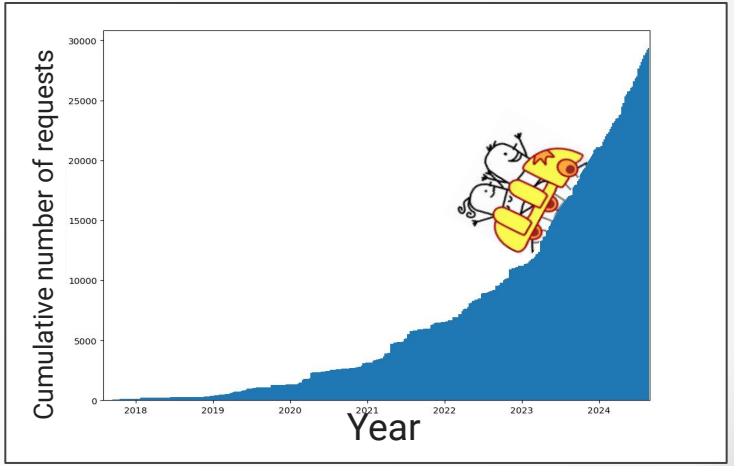

Cumulative APs submitted since invention

Analysis productions and open data

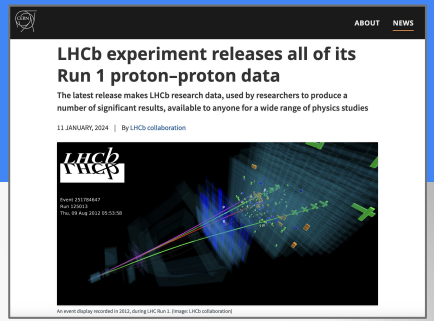

LHCb released its full Run 1 dataset ~ 800TB

- Need scalable solution going forward - **NTuple Wizard!**

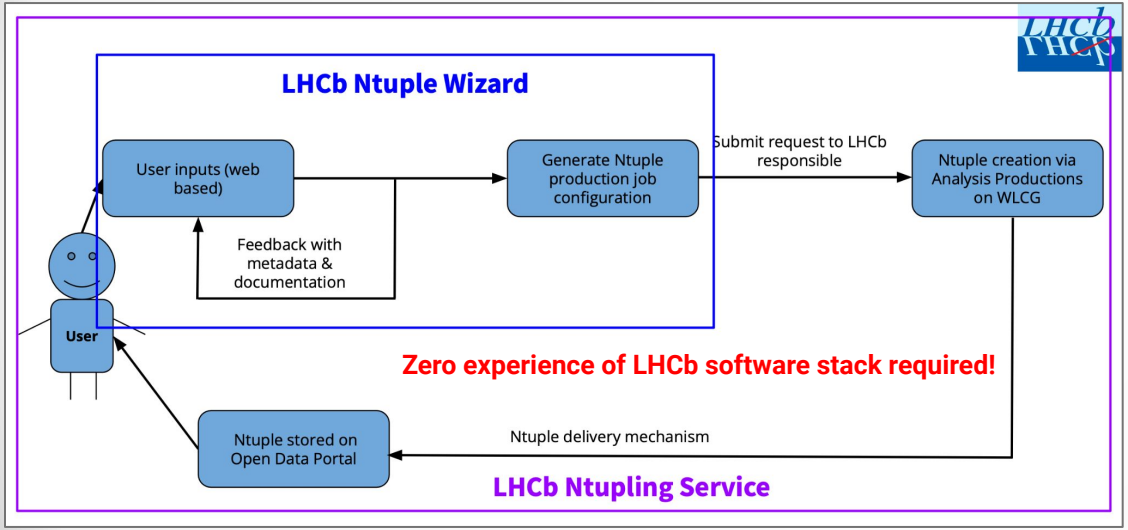

Typical non-LHCb analyst

See Piet's talk on NTuple Wizard

See Nathan's talk on continued exploitation of legacy data

Helping LHCb members new and old

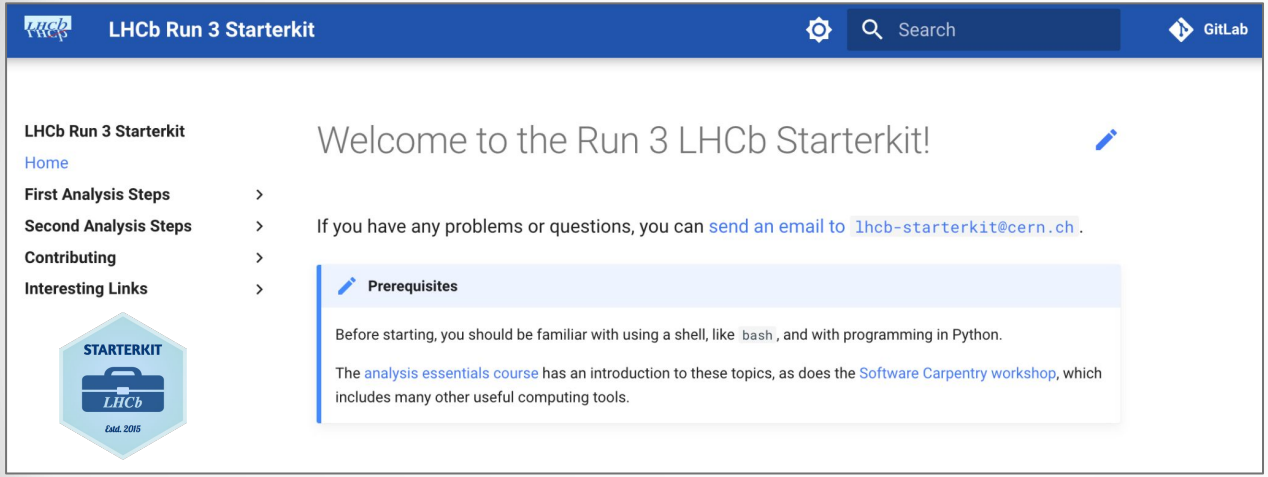

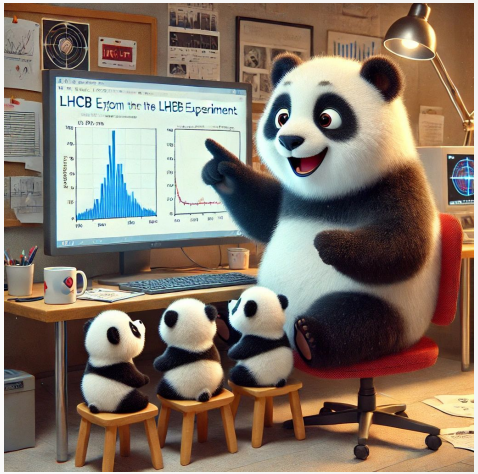

Not code documentation but a teaching resource following best pedagogical practices

[Cross experiment paper on training initiatives in HEP](#)
(submitted to Frontiers)

Summary

- LHCb Run 3 necessitated an overhaul of the offline dataflow
  - Sprucing provides a method for keeping high event persistency on tape for future exploitation but also manageable disk requirements
  - Analysis productions for ntuple creation are one of the single-most transformative changes at LHCb saving countless person-hours

- Analysis productions also allow LHCb data to be “open” via the NTuple Wizard
- With such a change in the dataflow LHCb is creating the Run 3 StarterKit for onboarding and continuous reference material

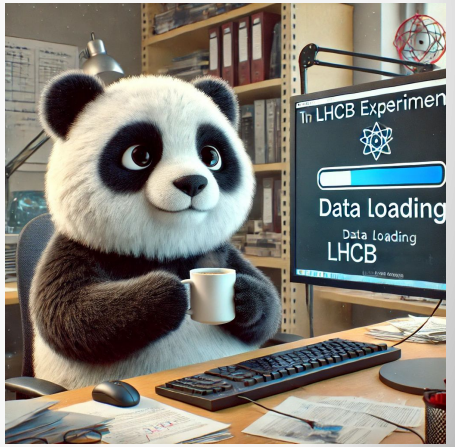

What then?

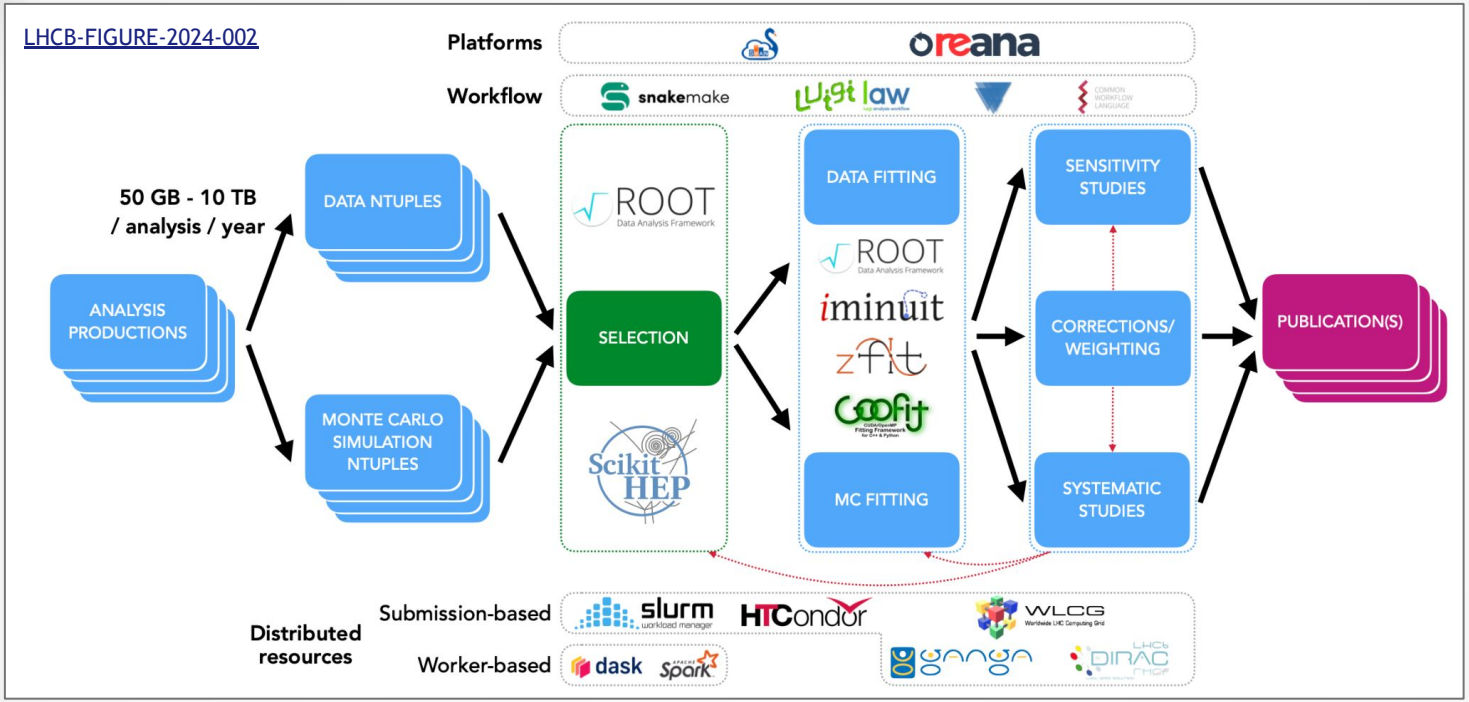

apd tool has snakemake integrations

LHCb DPA

See Mindaugas' talk on analysis scalability and reproducibility