LHCb status

Andrei Tsaregorodtsev on behalf of the LHCb Computing Project,

Journées LCG-France

26th July 2025, CC/IN2P3, Lyon

LHCb Data Processing activities

- Re-processing of data from 2024 and previous years
  - to improve the selections in view of the recorded data, or add new ones
- 2025 data processing
- Analysis Productions
- MC Productions

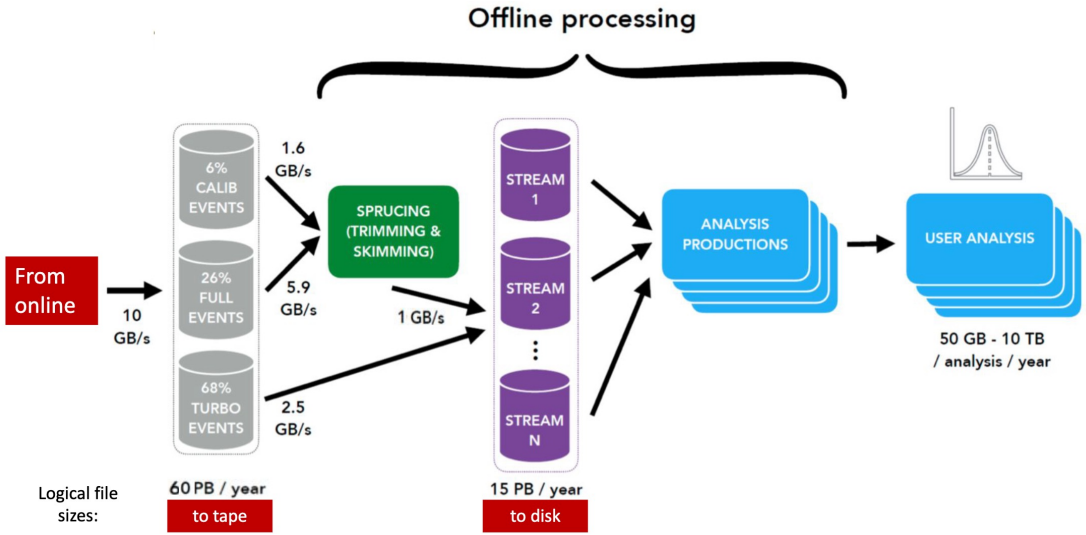

Data rates in practice

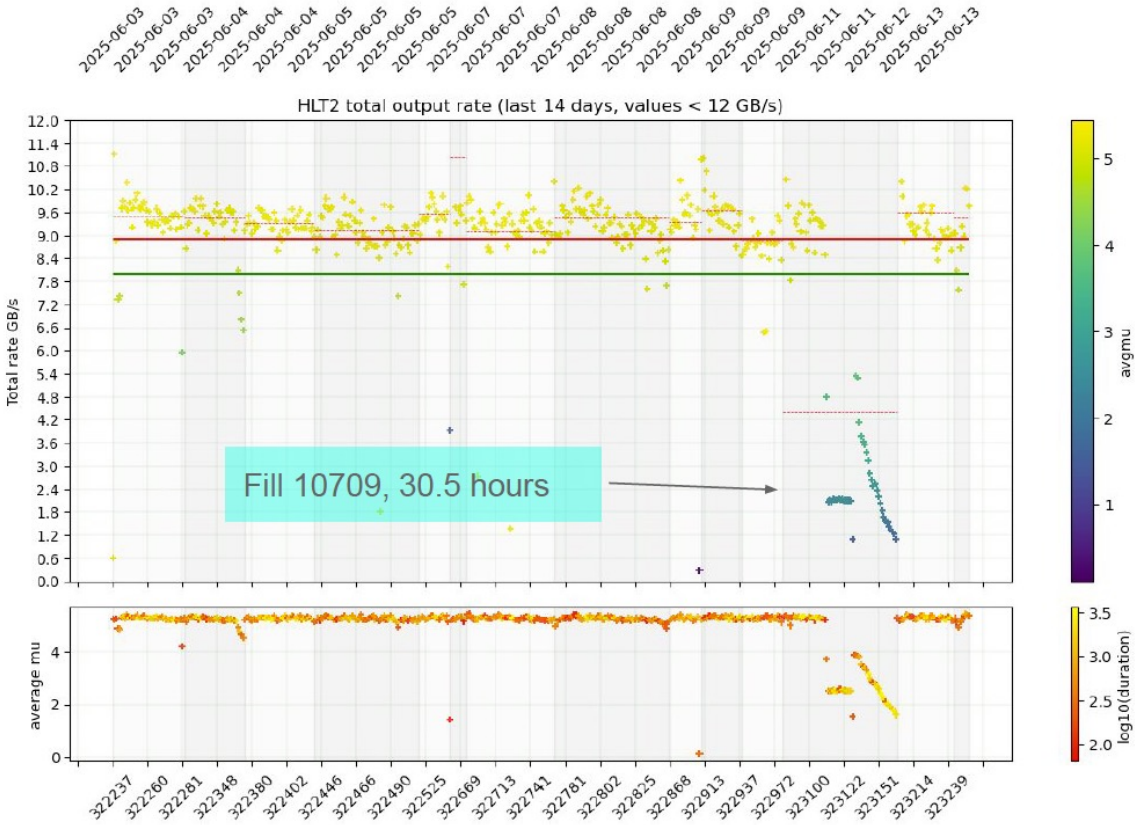

With 2025 data taking conditions we see 9.5 GB/s in runs with 2200 bunches and SMOG2

This will be reviewed during June Technical Stop

LHCb data processing

DIRAC DB update
+ DiracX first installation

- Making use of HLT farm as much as possible
  - until start of data taking
- Intervention in April to:
  - update the DBs, migrating their content to new charset
  - redeploy the infrastructure for the new version of DIRAC/DIRACX

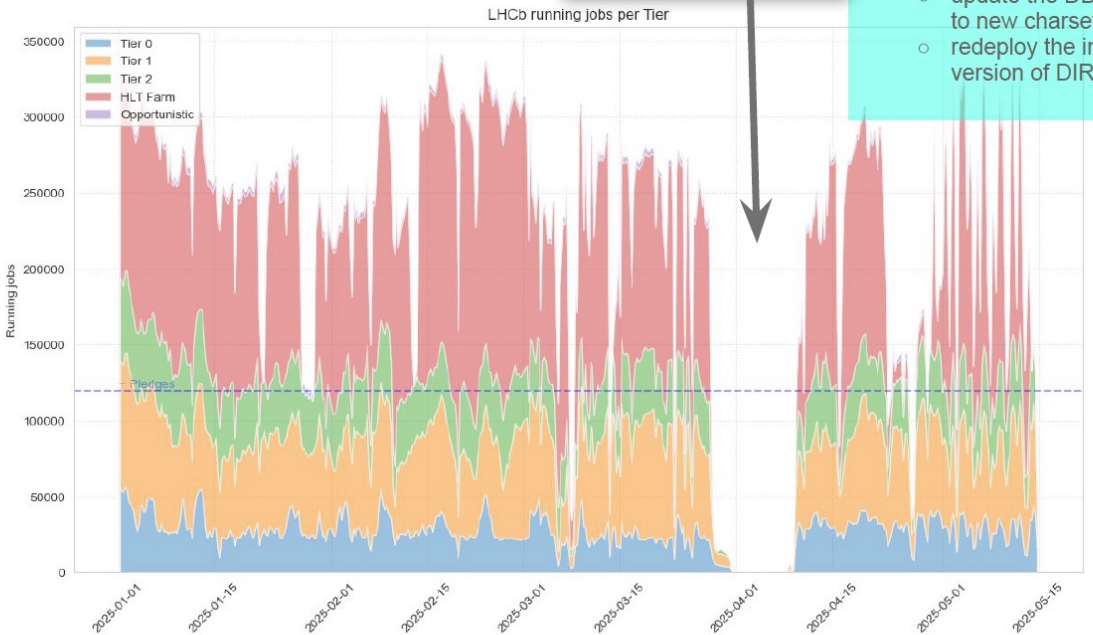

LHCb week #116 - Jun 16, 2025

Opportunistic resources

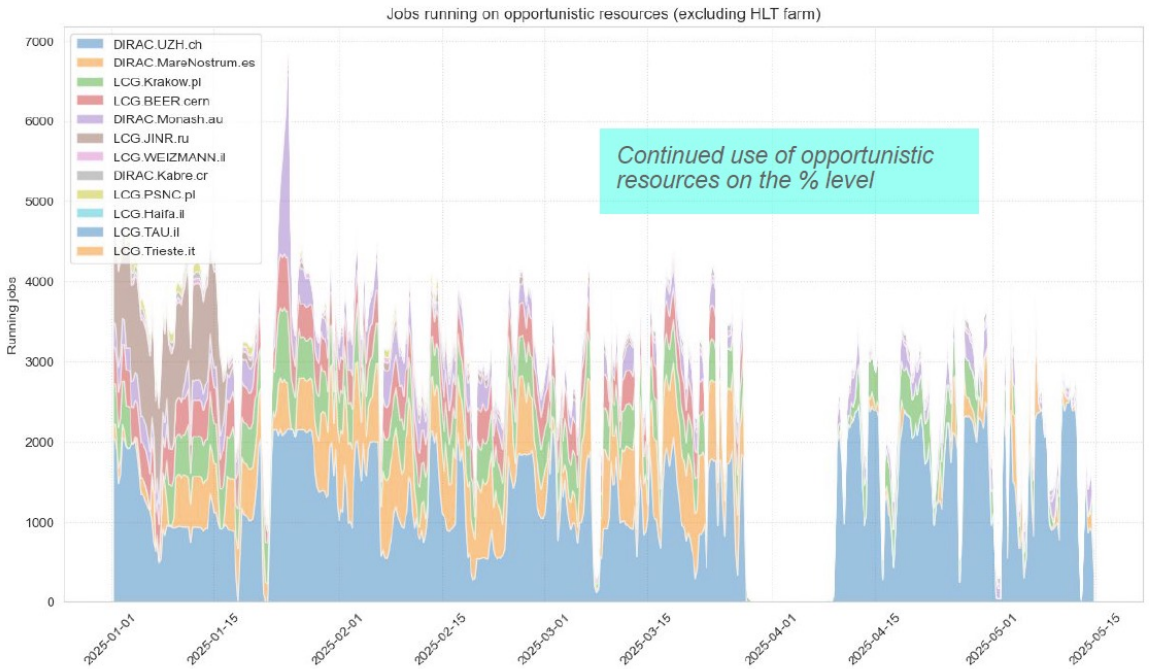

LHCb week #116 - Jun 16, 2025

CPU Resource usage

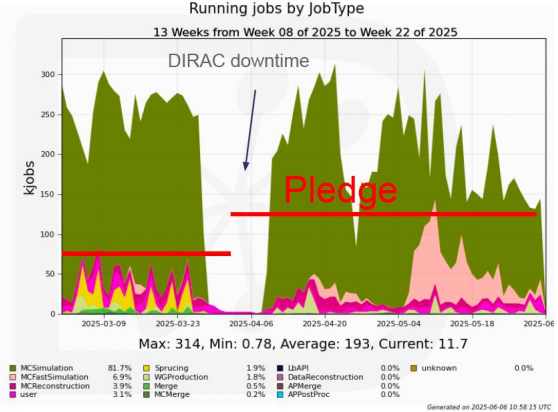

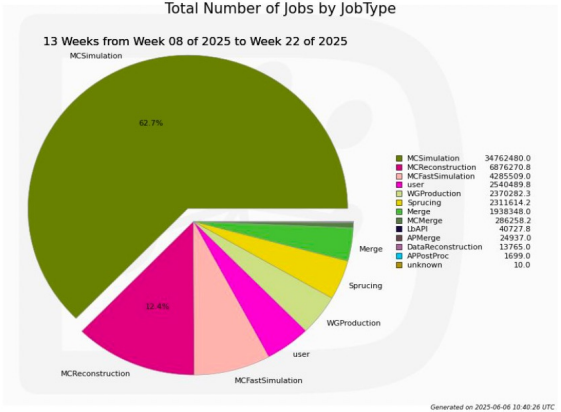

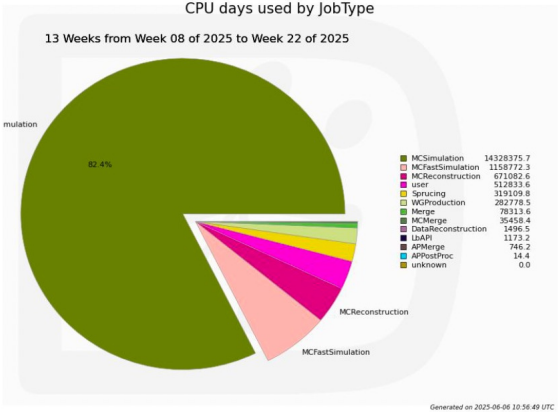

- ~55 million jobs executed last quarter
- 80% MC jobs, which constitutes ~90% of used CPU time
- The system was full most of the time, with an expected break in the beginning of April due to DIRAC upgrade
- The pledge was ~76k jobs before 01.04, ~123.5k since 01.04*

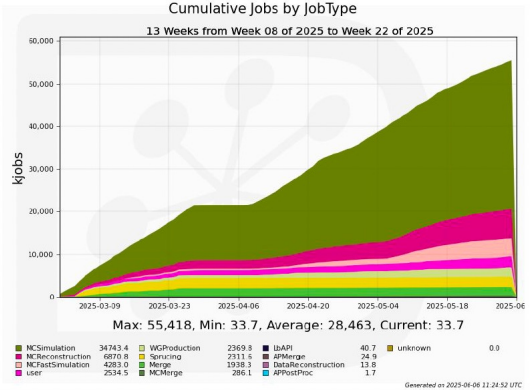

* average HS23 calculated as total HS23 hours over total CPU hours for the last year, based on accounting portal data

Science and Technology Facilities Council

NCB Meeting, 16.06.25

Simulated events

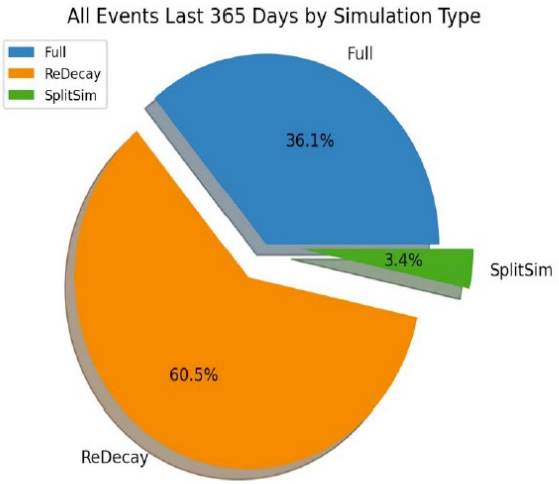

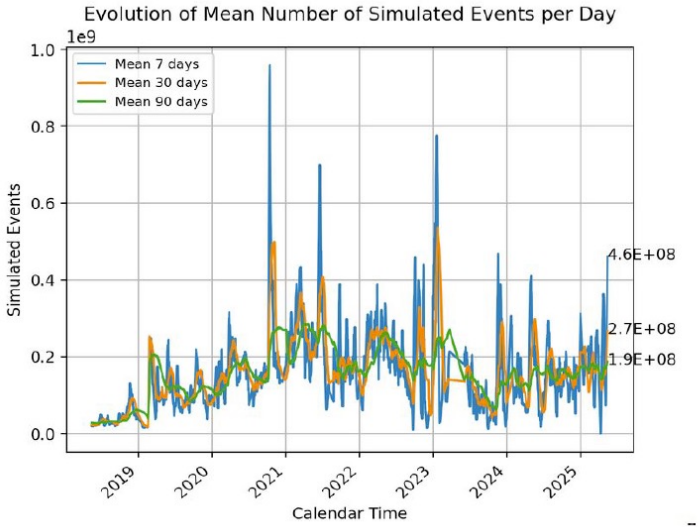

Increase in the use of FastSim in recent months (when it was 45%)

CPU Resource usage by site

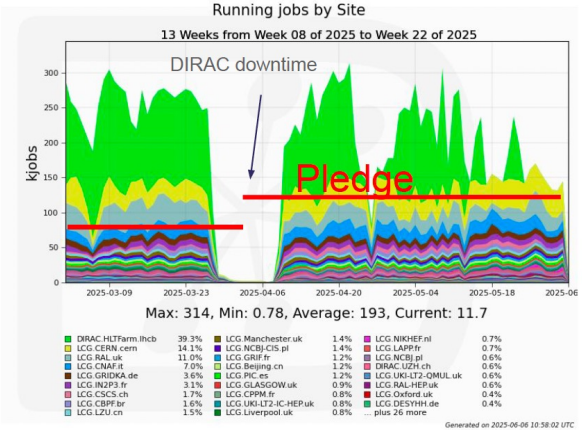

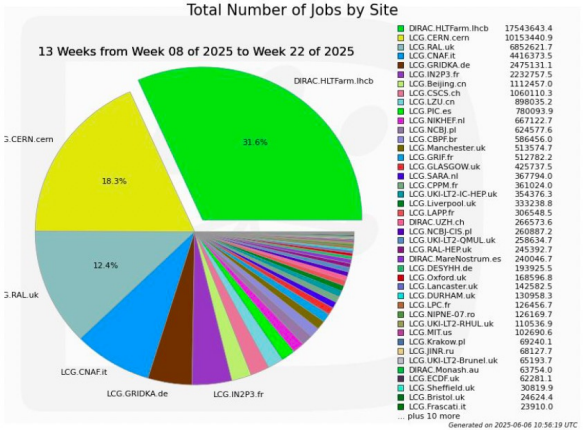

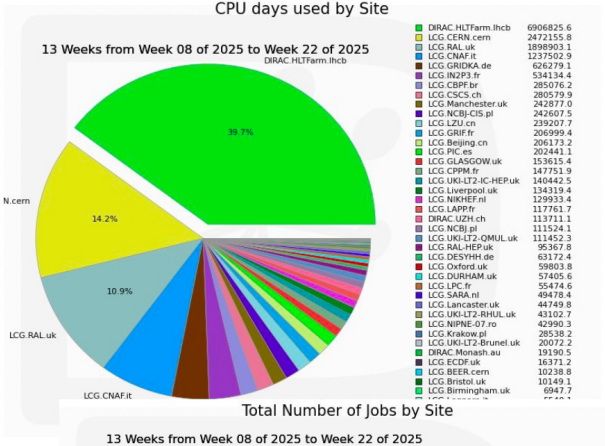

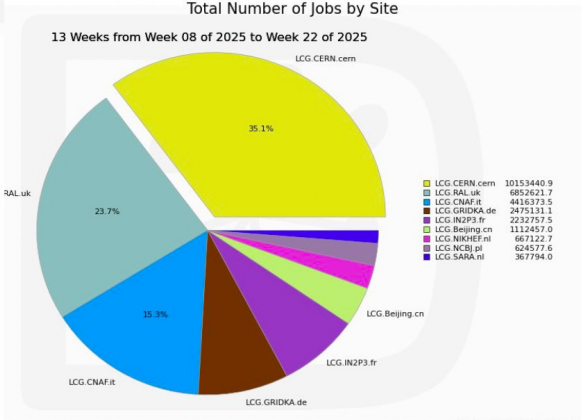

- HLTFarm was the main contributor, followed by CERN, RAL and CNAF
- Top 3 LCG sites are CERN (35% of LCG jobs), RAL (24%) and CNAF (15%)
  - SARA was underused
    - A bug in LHCbDIRAC prevented us from running MC jobs there (unable to use CentOS7)
- Tier-[01] sites executed ~50% of all jobs (40% cpu time)

Storage Resource usage

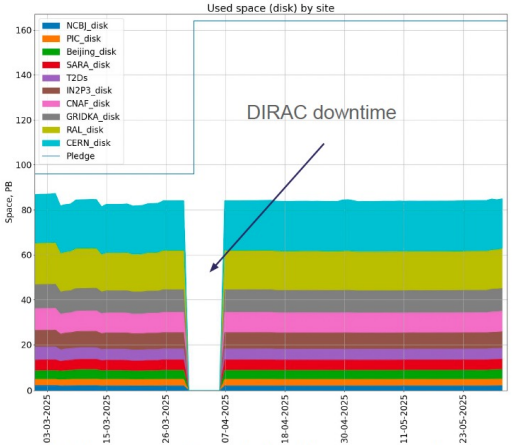

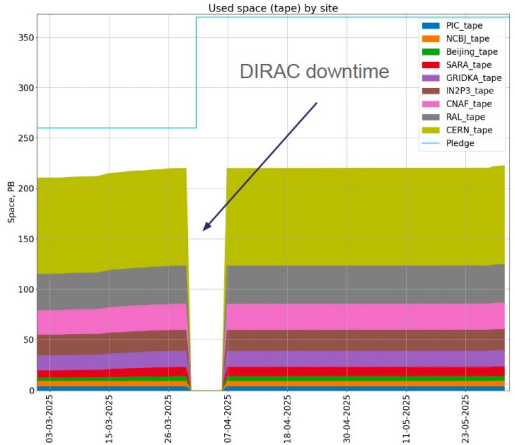

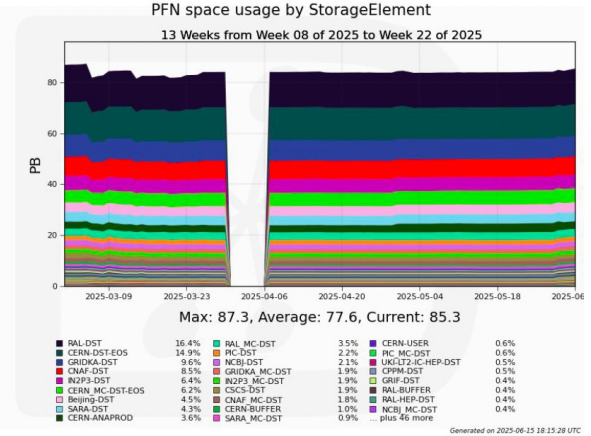

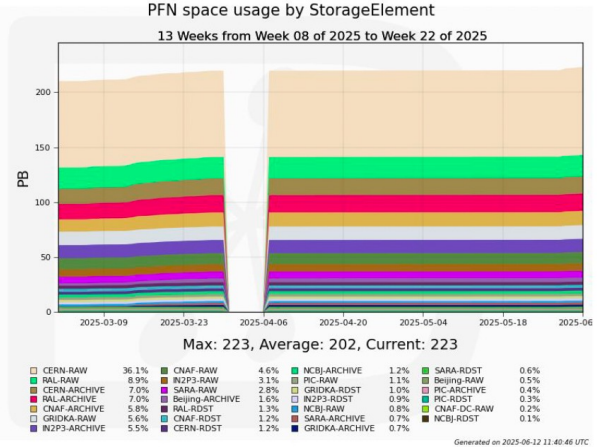

- Disk usage reduced after the reprocessing campaign
  - A few cleanups have been done
  - Total disk pledge for 2025 is **164PB**
    - Approximately **85PB** is used (40% at CERN, 62% at Tier-1 sites)
- Total tape pledge for 2025 is **370PB**
  - Approx. **223PB** (60%) is used (58% at CERN, 62% at Tier-1 sites)
- Dips are due do LHCbDIRAC upgrade (see below)

Science and Technology Facilities Council

Operational highlights

- Deep LHCbDIRAC downtime
  - Database upgrade and schema change, preparation for DiracX
  - Took a bit longer than expected, ~1.5 weeks, due to lengthy database upgrade operations
- Site operational issues
  - File loss at multiple sites/SEs (RAL, PIC, CNAF, GRIF, EOS-PILOT)
    - The GRIF ticket is no more there
  - Intermittent network issues at RAL
  - Storage issues at CNAF
- 2024 reprocessing campaign is over
- 2025 data distribution started

Using heterogeneous resources for MC

GPUs

- Relying on the Geant4-AdePT framework to support GPUs for particle transport
- Hybrid CPU-GPU Geant4 workflow integrated into Gaussino – the new LHCb MC framework
- Calorimeter Fast Simulation
  - Train AI model on the output on GPU
  - Produce hits by running inference on the generator part on the CPU or GPU

ARM

- Validation still ongoing, fixes incorporated into the Gaussino

Using heterogeneous resources for MC

HPC

- Using CSCS, Switzerland and BSC, Spain
- PushJobAgent improved considerably to scale better with the number of jobs
  - Running “pilot” at the central service, send user payload to the site to overcome the site-central service communications
- ToDo: still mostly using single-core jobs, poorly suitable for HPC
  - Need Gaussino with the support for multi-threaded processing
  - Allocation of whole nodes running several (up to 112) parallel jobs

HLT

- Used heavily for MC productions
- HLT SE is in the test - will allow to use it also for data processing

Tokens usage in LHCb

- LHCb IAM instance is in place
  - DIRAC registry is synchronized with it
  - Can deliver proxies
- Accessing CPU resources with tokens
  - Pilot submission
- Accessing storage resources
  - Still work in progress
  - Simple case – data distribution on the grid
    - DC24 used as a testbed
  - Difficult case – user and job interactions
    - Solution : DIRAC FC minting file access tokens (high rates)
- Access DiracX services with DiracX tokens

LHCb Tokens during DC24

- Efficiency of token-based transfers are much lower, compared to certificate-based

Mini DC’25

- No particular plans to participate in the mini DC using French sites
- If the decision will be take to participate
  - Existing DC’24 tools can be reused
  - A special dataset will be repared
  - Simple data injection script:
    - Placing an FTC request for a data transfer, file status monitoring, deletion upon successful transfer
    - Maintaining constant load on each channel
    - FTS requests are submitted with the DC tag, WLCG monitoring can be used
  - Exact mini DC parameters to be better defined before the start

LHCb RUN 3 data analysis model

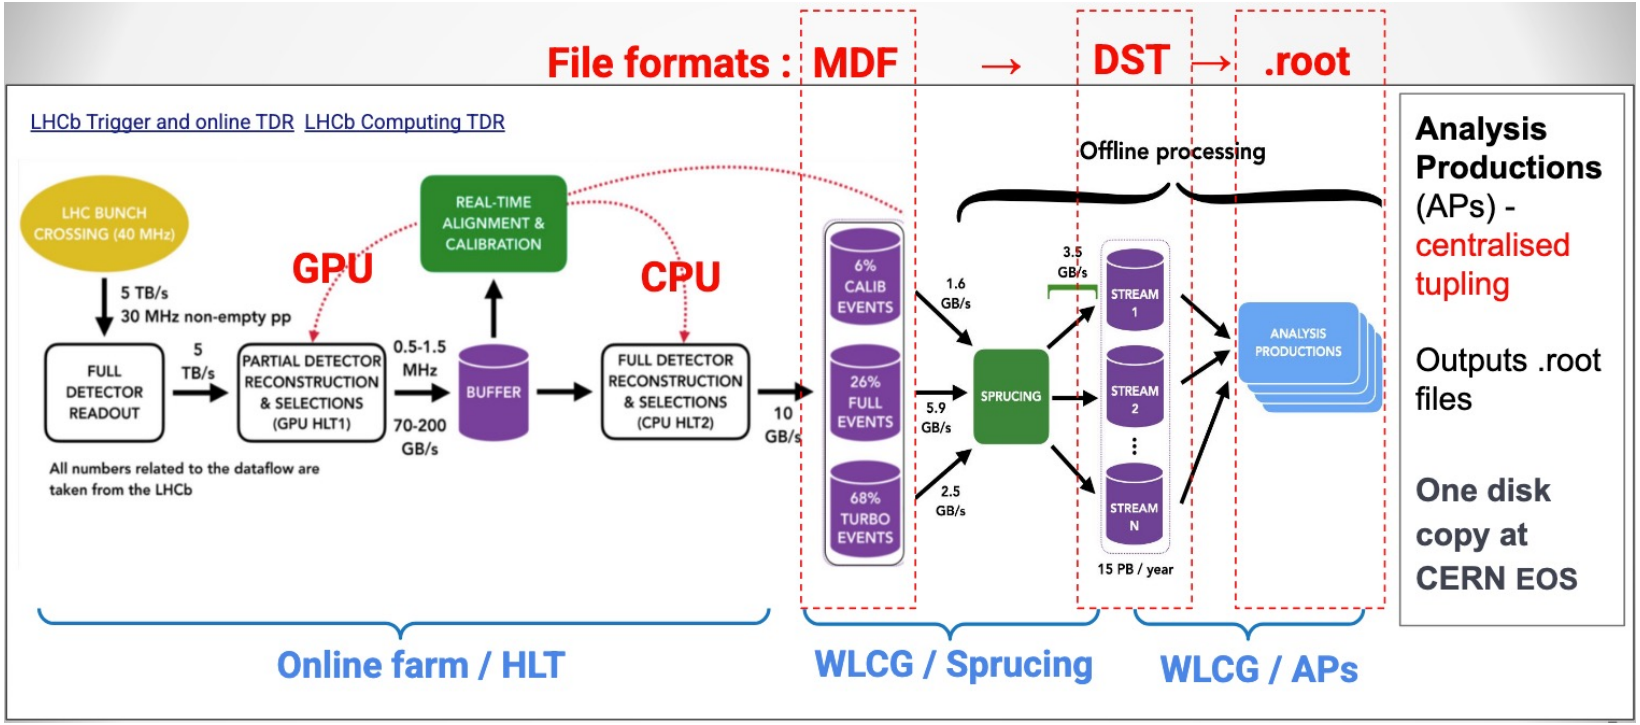

AF vs AP

- **Analysis productions** are an analysis facility in themselves
  - Declarative tupling via python options and yamls
  - Centralised productions run on WLCG
  - Exploit DIRAC transformation system
  - Full job testing on GitLab CI – avoid failures on the grid
- Positives
  - Efficient use of WLCG resources
  - Output referenced in bookkeeping - makes for easy cleanup
- Negatives
  - Very high I/O load on WLCG sites
  - High memory consumption causing jobs to be killed at some sites

AP Run 3 numbers

Storage

The storage used for run 3 analysis is part of the pledges, and includes:

- **Analysis Production output** (ntuples): ~2.6 PB (data+MC)
- User directories and WG areas on EOS at CERN (~1.3PB)
- User data on the WLCG (~1.2 PB)

Network

- Network is non-limiting for LHCb
- LHCb follows **job -> data model** (low WAN, high LAN)
  - WAN use is very small, only tuples transferred back

CPU

Analysis jobs form small part of LHCb WLCG load:

Analysis Productions ~6%

User jobs via Ganga: tupling, toy studies etc. ~4%

Sprucing 5%

MC productions 85%

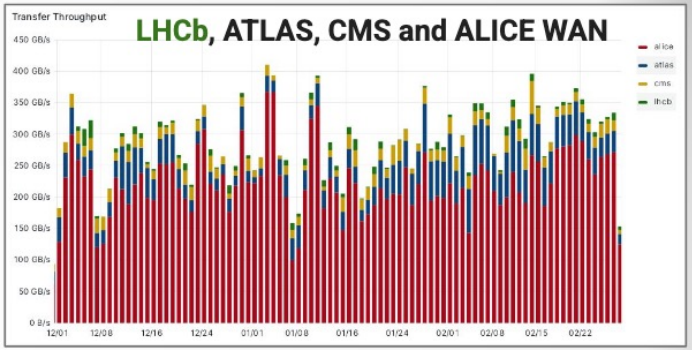

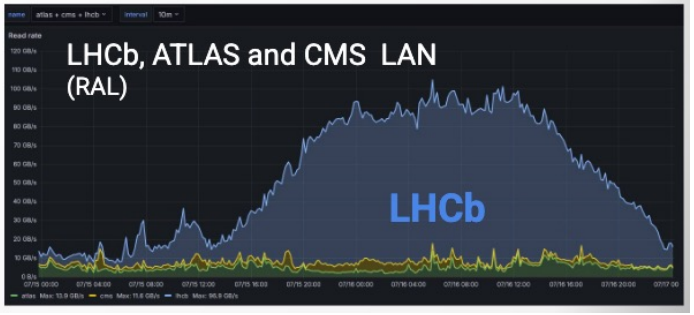

AF vs AP

- LHCb is opting for
  - centrally process/reduce datasets such that final analysis can be done by analyst on a “single machine”
  - with portable analysis environments can run any application in **AP**s and exploit as many resources as possible
  - Continuous R&D to cope with:
    - High I/O load, memory consumption, ability to run custom validated code…

- No particular request for AF
  - Using « standard » resources
  - Watching the current trends, keeping doors open …

DiracX migration

- DiracX
  - A complete rewrite of DIRAC
    - Keeping the successful architecture
    - Using new standards
    - Multi-VO, tokens, etc, naitively built in
  - Migration should be transparent to DIRAC users
  - LHCb performed the migration
    - April technical stop
    - DIRAC8.0 -> DIRAC9.0 + DiracX0.1

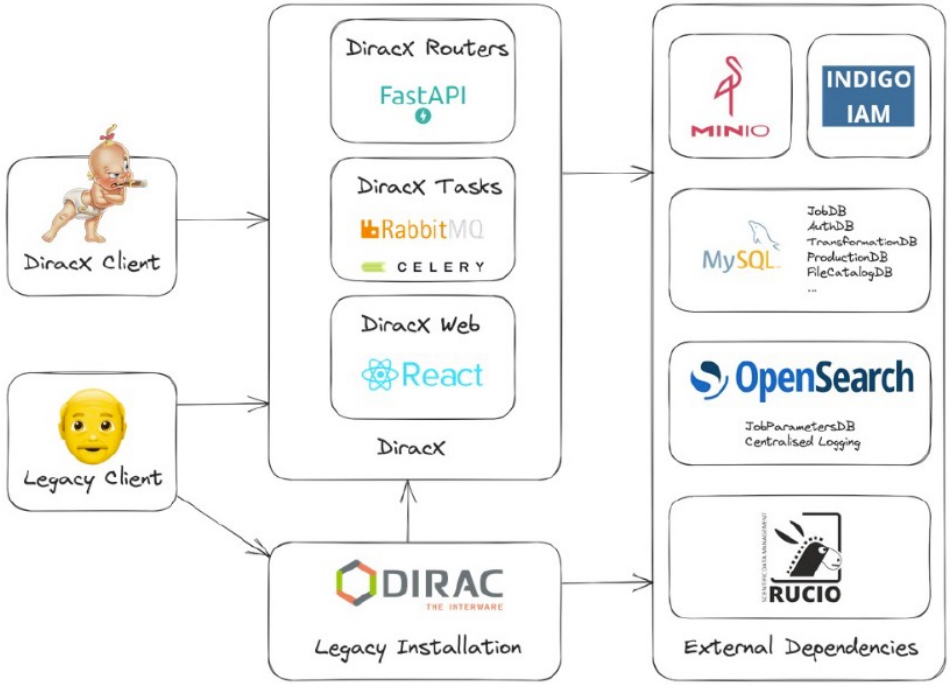

DiracX migration

- LHCb performed the migration during the April technical stop
  - DIRAC8.0 -> DIRAC9.0 + DiracX0.1
  - Took about 10 days of the service downtime
    - Multiple database optimizations
    - Several quick fixes
  - Only one DIRAC service (JobStateUpdate) is actually fully running with DiracX
    - One of the busiest DIRAC services
    - Old clients (pilots, jobs) are accessing it with a legacy adapter
  - The LHCb Production system is now running in full swing in the new setting
    - Paving the way for migration of other DIRAC installations

LHCbDiracX roadmap

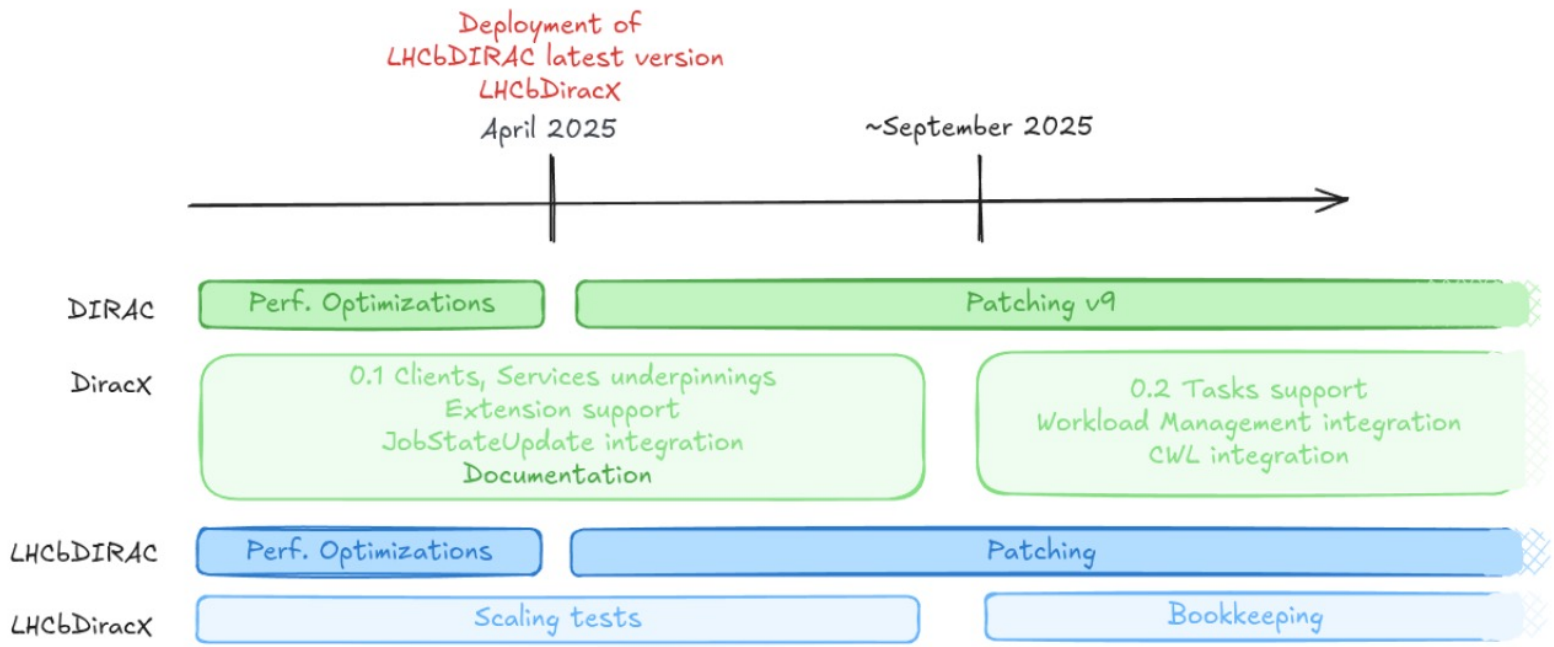

Resource requests for 2026 - CERN-RRB-2025-009

LHCb		2025		2026		
		RRB approved	Pledged	Request	2026 req. / 2025 RRB	C-RSG recomm.
CPU	Tier0	283	283	344	122%	344
	Tier1	928	856	1127	121%	1127
	Tier2	518	535	629	121%	629
	Total	**1729**	**1674**	**2100**	**121%**	**2100**
	HLT Others					
Disk	Tier0	54.9	54.9	70.9	129%	70.9
	Tier1	89.9	82.7	107.1	119%	107.1
	Tier2	17.4	15.2	20.7	119%	20.7
	Total	**162.2**	**152.8**	**198.7**	**123%**	**198.7**
Tape	Tier0	170.4	170.4	202.2	119%	202.2
	Tier1	194.8	164.2	233.7	120%	233.7
	Total	**365.2**	**334.6**	**435.9**	**119%**	**435.9**

The C-RSG considers that the LHCb Collaboration's resource requests for 2026 are necessary to achieve the experiment's physics programme.

Long-term (Run 4) resources evolution at « constant budget »

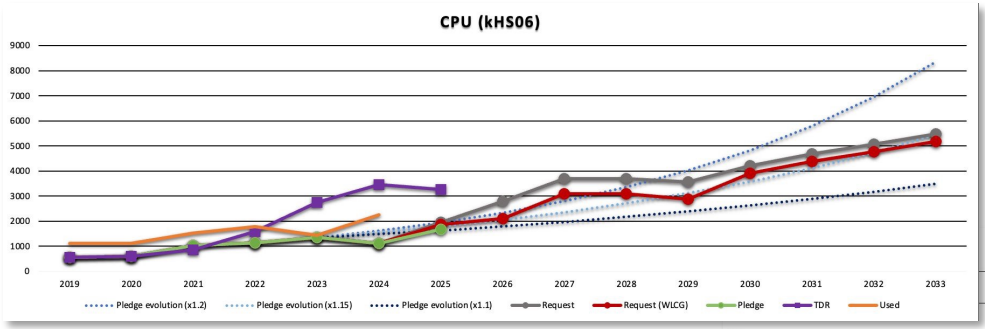

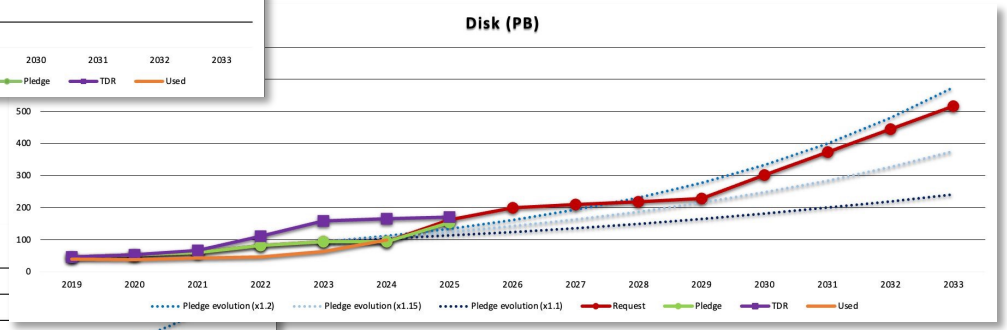

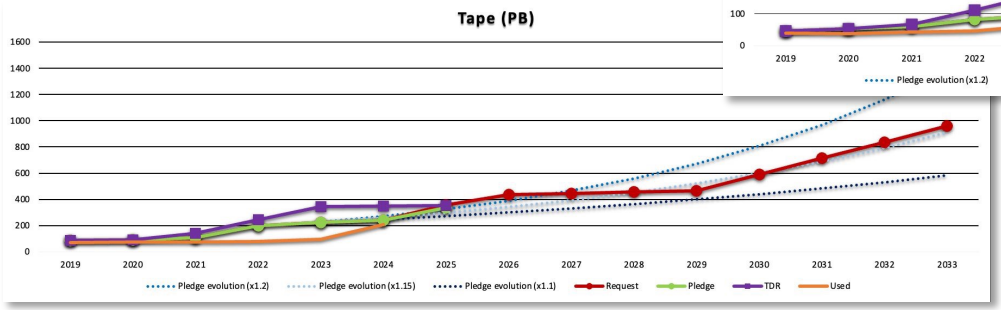

Outlook to Upgrade 2 – Run 5

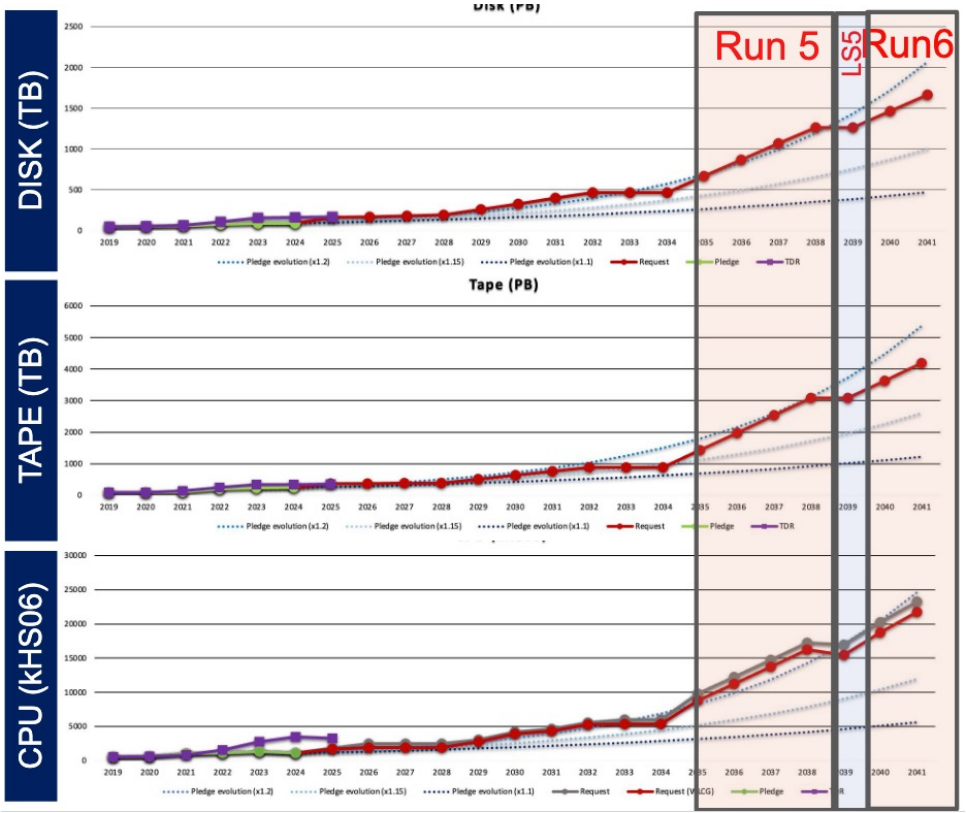

- Entering a different regime
  - Storage: a few exabytes
    - (x3-x5 wrt constant budget)
  - Compute: tens of MHS06
    - (x5 wrt constant budget)

Conclusions

- Busy YETS full of data re-processings
  - And DiracX released to production!
- LHCb Grid performing well
  - With some scope for improvements
- 2025 data taking has started
- Resource situation for 2025 still looks OK
  - 2026 resource request endorsed by C-RSG
- Tools available for the mini-DC’25
- No particular plans for AF, use AP intensively

Back-up slides

Distributed Computing: Future demand (Summary)

Expectations

- **Increased Luminosity:**
  - Allows the experiment to record a larger number of events.
  - Higher luminosity results in more complex event data.
- **Impacts on Storage:**
  - Larger datasets require significantly more storage capacity.
- **Monte Carlo Simulations:**
  - Increased data volume demands a greater number of MC simulations to ensure thorough analysis.
  - For some type of events, the generation can be more time-consuming.

Model assumptions		
	Upgrade I	Upgrade II
Peak L ($cm^{-2}s^{-1}$)	2×10^{33}	1.5×10^{34}
Yearly integrated luminosity (fb^{-1})	10	50
Logical bandwidth to tape (GB/s)	10	50
Logical bandwidth to disk (GB/s)	3.5	17.5
Running time (s)	5×10^6	
Trigger rate fraction (%)	26 / 68 / 6 Full / Turbo / TurCal	
Ratio Turbo/Full event size	16.7%	
Ratio full/fast/param. MC	40:40:20	
CPU work per event full/fast/param. MC (HS06.s)	1200 / 400 / 20	
Number of simulated events	$4.8 \times 10^9 / fb^{-1}$/year	
Data replicas on tape	2 (1 for derived data)	
Data replicas on disk	2 (Turbo); 3 (Full, TurCal)	
MC replicas on tape	1 (MDST)	
MC replicas on disk	0.3 (MDST, 30% of the total dataset)	

From the 7th Workshop on LHCb Upgrade II – Computing Model

LHCb data processing

As usual Simulation represents the majority of the jobs

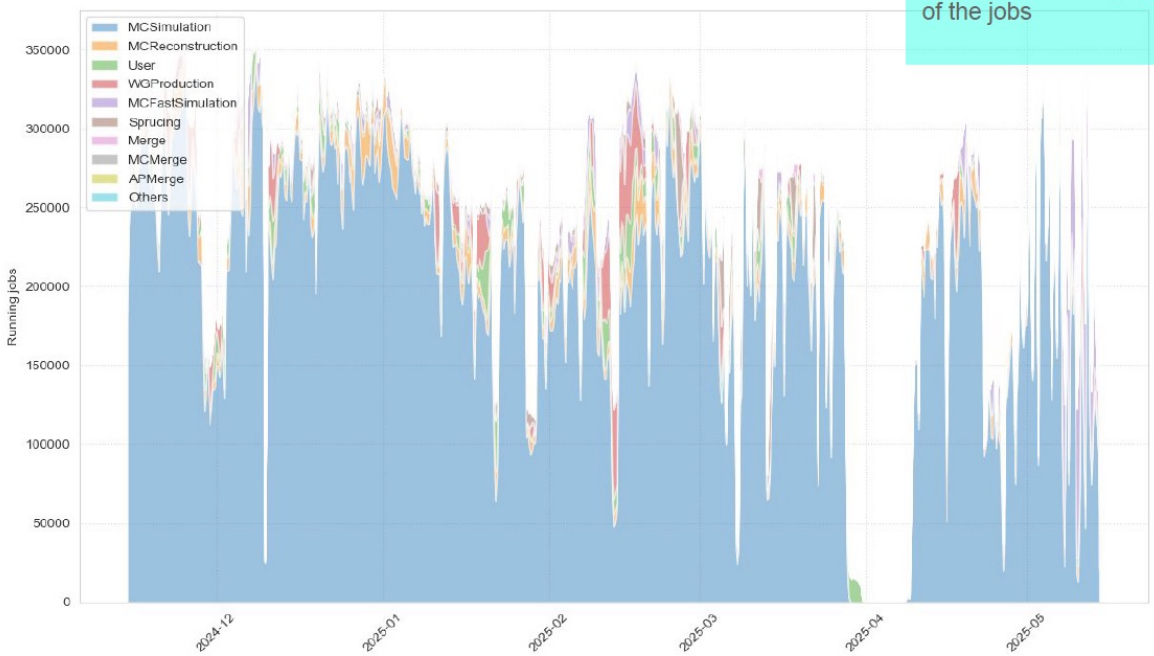

LHCb week #116 - Jun 16, 2025

IAM performance measurements

Get Access Tokens (client_credentials)
Roughly 1000 requests per second

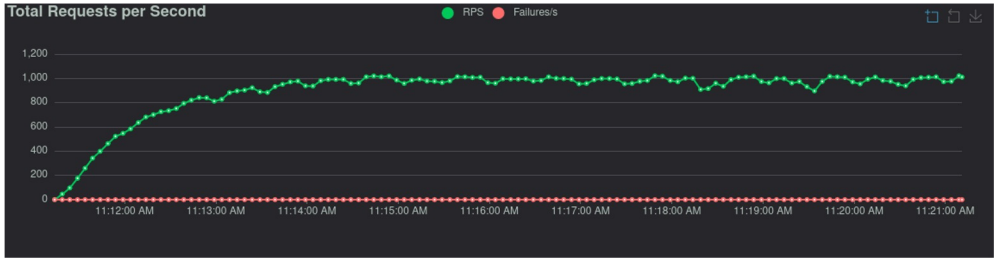

Token exchange to get a refresh token
Roughly 700 requests per second

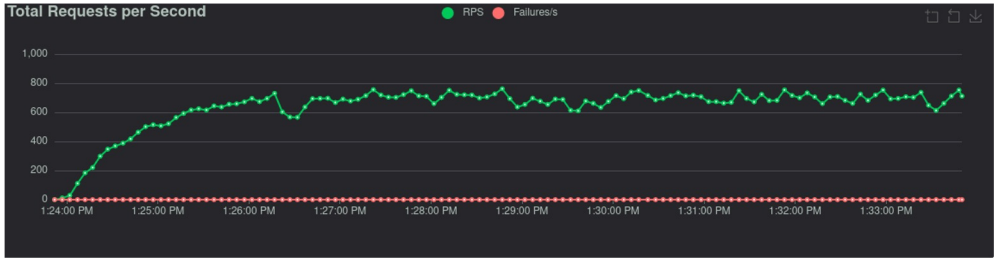

Get Access Tokens (using a refresh token)
Roughly 900 requests per second