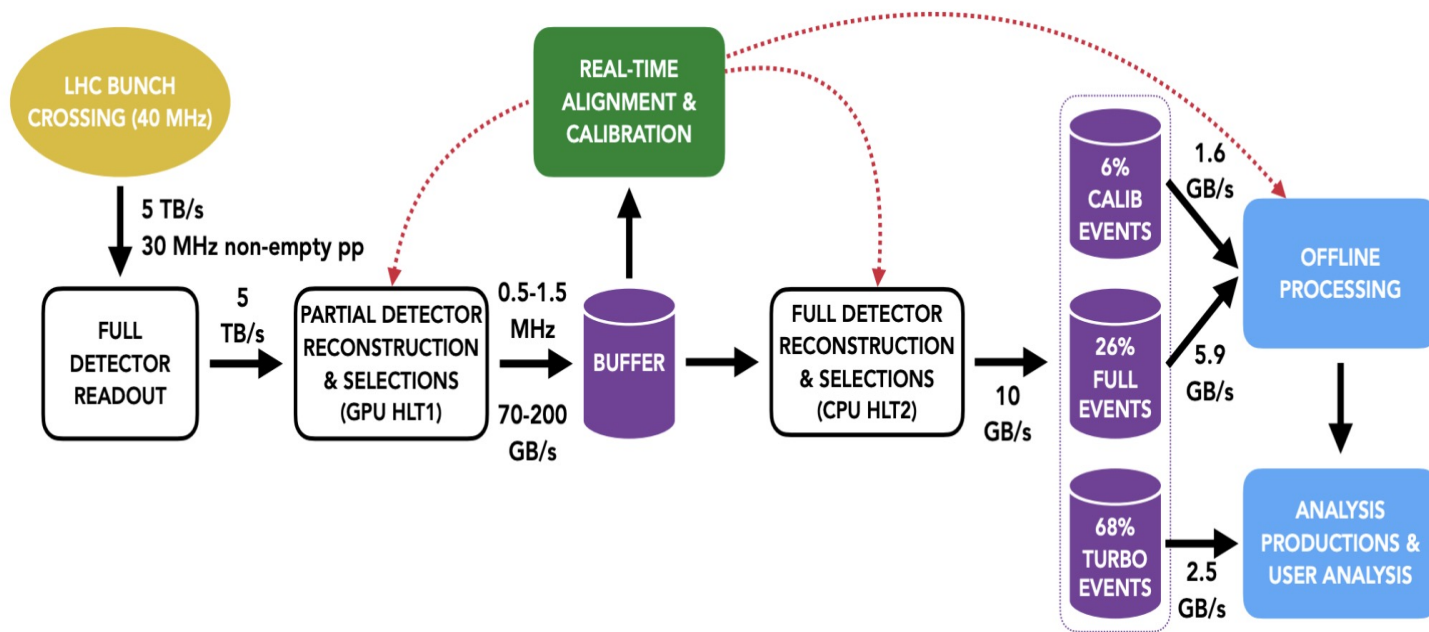




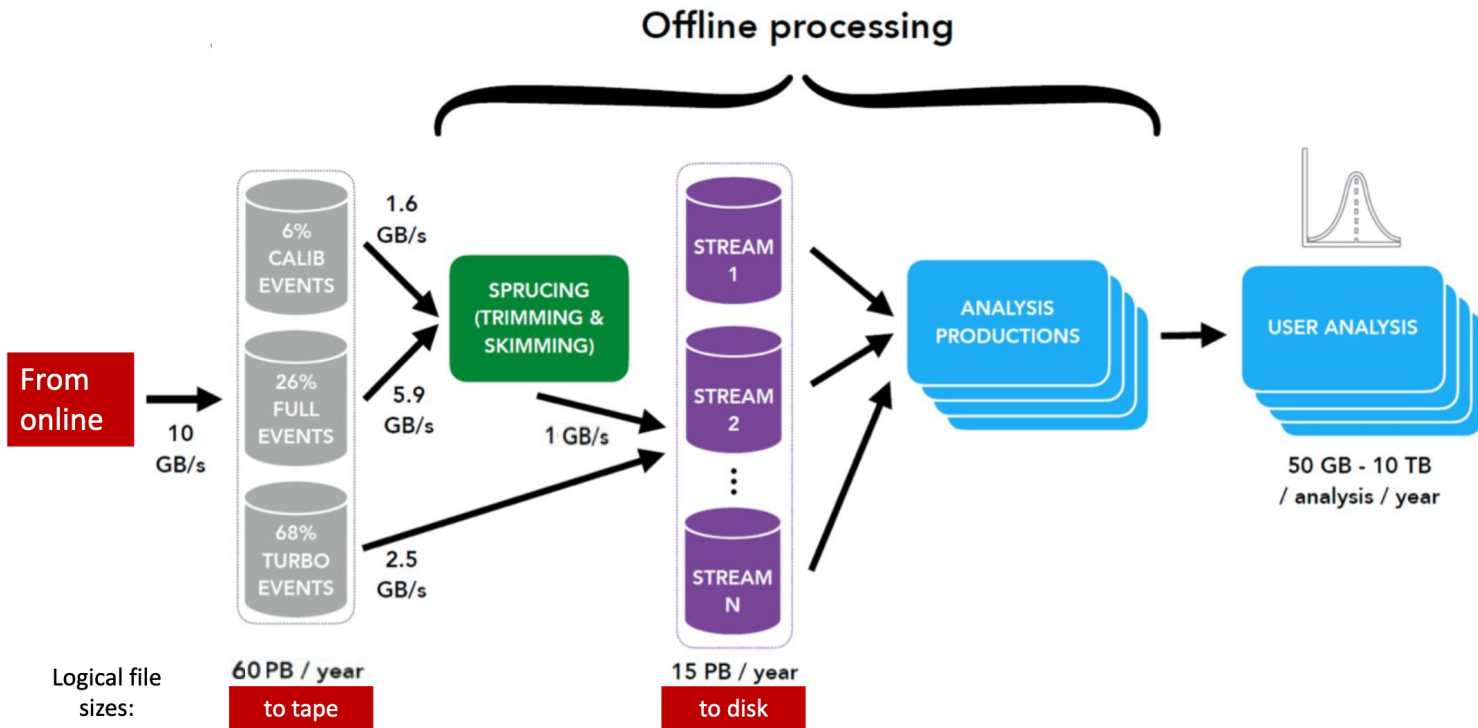
LHCb RUN3 activity

*Andrei Tsaregorodtsev on behalf
of the LHCb Computing Project,
Journées LCG-France
1st December, 2023, Lyon*

LHCb Run3 Computing model reminder: online





LHCb Run3 Computing model reminder: offline



Distributed computing operations

- Computing work: **MC** production (89%), physics analysis (9%), data reconstruction and selection (stripping/sprucing, 2%)
- Usual mix of sites T0:T1:T2 ~ 40:40:20




Similar pattern in 2024-2025

Distributed computing operations

- Simulating about **100 million** events per day in the last three months
- **1/3 of events** produced with **detailed simulation** in the last 365 days
- long tail of Run1+Run2 simulation, confirming computing model assumption
- Small Run3 simulation samples produced for specific purposes (early measurements, HLT bandwidth division, detector and software performance studies, etc.)


All Events Last 365 Days by Simulation Type



2023 – turned out to be a commissioning year:



- Velo acceptance limited
- UT is in commissioning mode
- Underuse of storage resources

Evolution of Mean Number of Simulated Events per Day



Lots of real data productions

- Offline re-run of **HLT2** on **2022** proton collision data
 - Small but **precious dataset**
 - VELO closed to nominal position
 - Validation completed, production imminent
 - Sprucing will follow afterwards
- **Sprucing of 2023** proton collision data completed
 - Commissioning and collision data
 - Before and after June technical stop
- **Stripping of Run2** datasets
 - Key to LHCb physics output in the next couple of years
 - 2016 dataset running, 50% complete
 - 2017 and 2018 will follow
 - Staging from tape massive amounts of data
 - Stripped data need additional 2PB logical (4PB physical) disk space




Heavy ion data

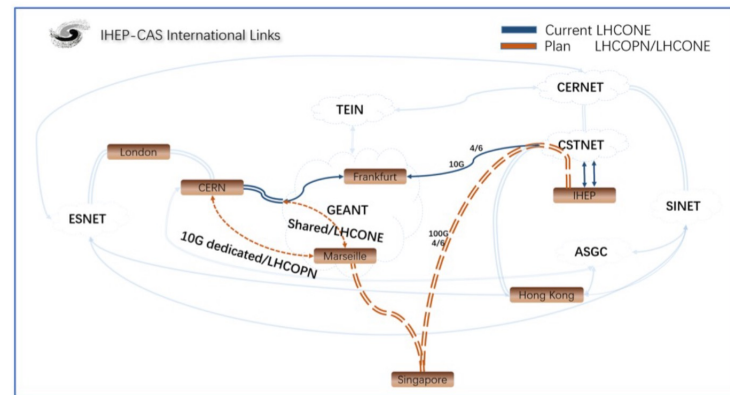
- **Heavy ion collision data processing**
 - Run HLT2 on the online farm at the pit
 - HLT farm will be not used for simulation during the shutdown
 - Run sprucing offline
 - Expected 4PB (2PB) logical file size on tape (disk)
 - Physical file size will be twice (two copies)
- **HI+fixed target** is representing about 10% of required resources
 - This is the assumption for both CPU and storage for 2024-2025 request
 - Otherwise goes through the same workflow as pp-collisions data
 - HI run is to be confirmed for 2024

New LHCb Tier1 sites

- The NCBJ (Warsaw) and IHEP (Beijing) Tier2 sites have progress towards becoming Tier1 sites for LHCb
 - requirements in terms of network, storage (most notably: tape), services, service level agreement
- NCBJ: network OK, computing hardware and configuration ready
 - Tape software has been upgraded in September, data challenge has been performed successfully
 - Final commissioning in December'23
- IHEP: computing HW and SW OK, network ~OK
 - 100Gbps available in shared mode, functional tests OK, data challenge imminent
 - Dedicated 10Gbps link to LHCOPN is in progress
 - Final commissioning 1H'24
- IHEP and NCBJ will add ~5% each to the LHCb resources




Institute of High Energy Physics
Chinese Academy of Sciences




Opportunistic resources

- Fraction of jobs executed on totally opportunistic resources stable at a few percent level
- Zurich (CH), Beer (CERN), Monash (AU), Sibir (RU), Barcelona (ES)
- Barcelona Supercomputing Center (MareNostrum), SDumont.br and CSCS are production HPCs
 - CINECA – GPU access
 - small center in Costa Rica (Kabre) recently added
- WLCG sites not pledging to LHCb in France, Israel, Italy are also utilised opportunistically



Jobs at the French sites



GPU usage in LHCb

LHCb (main messages) – J.Flix, Nov'23 GDB

- The Allen framework is used in Run3 to execute on the event building farm GPUs the first HLT stage (HLT1). Offline GPUs could be used to emulate the HLT1 in simulation through Allen
- A few LHCb analysts use GPUs, but they use resources on institute clusters which are outside WLCG. Actual usage is unknown. Fast simulations use machine-learning techniques currently implemented on CPUs, but they might be ported on GPUs in future. LHCb might be using opportunistically the GPUs on the online event builder farm when not in use for data taking
- No established plans on GPU resource demands or for future utilization at sites. Some discussion have started in the context of GPU usage for physics analysis, in the context of analysis facilities. November 2023 workshop to understand how to use GPUs/accelerators also in later processing stages, in the context of future Run4/Run5 LHCb Upgrades
- LHCb is totally dominated by simulation (using above 90% of compute resources). R&D towards porting (parts of) simulation on GPUs, provided they can be advantageous in terms of performance and required resource

Tokens usage in LHCb

- LHCb IAM instance is in place
- Accessing CPU resources with tokens
 - Most of the HTCondorCE sites are configured for token access
 - Tests with ARC6 sites were successful, ARC7 tests are ongoing
- Accessing storage resources
 - “Following the WLCG schedule” – not very aggressive
 - Unlikely to be used in the DC’24 in February

2024 pledges

- Snapshot from CRIC
- CPU OK
- Tape ~OK
- Disk -15%

| Tier | Pledge Type | Year | LHCb Required | LHCb Pledged | LHCb Balance |
|------|-------------|------|---------------|--------------|--------------|
| 0 | Disk | 2024 | 30600 | 30600 | 0 % |
| 0 | Tape | 2024 | 117100 | 117100 | 0 % |
| 0 | CPU | 2024 | 174000 | 174000 | 0 % |
| 1 | Tape | 2024 | 133300 | 125155 | -6 % |
| 1 | Disk | 2024 | 61200 | 52986 | -13 % |
| 1 | CPU | 2024 | 572000 | 542293 | -5 % |
| 2 | Disk | 2024 | 11800 | 9381 | -20 % |
| 2 | CPU | 2024 | 319000 | 393541 | 23 % |
| Tier | Pledge Type | Year | LHCb Required | LHCb Pledged | LHCb Balance |

Preliminary 2025 requests

LHCb-PUB-2023-003

| Model assumptions for 2025 | |
|---------------------------------------|-----------------------------|
| L ($\text{cm}^{-2}\text{s}^{-1}$) | 2×10^{33} |
| Pileup | 6 |
| Running time <i>pp</i> collisions (s) | 6.3×10^6 |
| Output bandwidth (GB/s) | 10 |
| Fraction of Turbo events | 73% |
| Ratio Turbo/FULL event size | 16.7% |
| Ratio full/fast/param. simulations | 36:64:0 |
| Data replicas on tape | 2; 1 for derived data |
| Data replicas on disk | 2 (Turbo); 3 (FULL, TurCal) |
| Simulation replicas (disk and tape) | 1 |

| stream | rate fraction | TAPE throughput (GB/s) | TAPE bandwidth fraction | DISK throughput (GB/s) | DISK bandwidth fraction |
|-------------|---------------|---------------------------|-------------------------------|------------------------------|----------------------------|
| FULL | 26% | 5.9 | 59% | 0.8 | 22% |
| Turbo | 68% | 2.5 | 25% | 2.5 | 72% |
| Calibration | 6% | 1.6 | 16% | 0.2 | 6% |
| Total | 100% | 10.0 | 100% | 3.5 | 100% |

Preliminary 2025 requests

LHCb-PUB-2023-003

| CPU Work in WLCG year (kHepScore23.years) | 2024 LHCb-PUB- 2023-001 | 2025 prel. THIS DOCUMENT |
|---|-------------------------------|--------------------------------|
| First pass sprucing | 70 | 82 |
| End-of-year sprucing | 70 | 82 |
| Simulation | 800 | 1336 |
| Core and distributed computing infrastructure | 10 | 10 |
| User Analysis productions | 214 | 319 |
| Total Work (kHepScore23.years) | 1165 | 1829 |
| LHCb-TDR-018 | 3470 | 3276 |

| Disk storage usage forecast (PB) | | | 2024 | 2025 prel. | |
|----------------------------------|---------------------|------|-------------------|---------------|-------|
| | | | LHCb-PUB-2023-001 | This document | |
| Real data | Run1+Run2 pp data | 10.2 | 78.7 | 10.2 | 134.4 |
| | Run1+Run2 HI+SMOG | | | | |
| | Run3: FULL | 16.5 | | 30.9 | |
| | Run3: TURBO | 36.3 | | 68.1 | |
| | Run3: TURCAL | 4.5 | | 8.4 | |
| | Run3: Minimum bias | 0.0 | | 0.0 | |
| | Run3: HI+SMOG2 | 11.2 | | 16.8 | |
| Simulated data | Run1+Run2 Sim | 8.7 | 11.9 | 8.7 | 16.7 |
| | Run3 simulated data | 3.2 | | 8.0 | |
| Other | User data | 3.0 | 13.0 | 3.6 | 15.6 |
| | Buffers | 10.0 | | 12.0 | |
| Total | | | 103.6 | | 166.7 |
| LHCb-TDR-018 | | | 165.0 | | 171.0 |

| Tape storage usage forecast (PB) | | 2024 LHCb-PUB-2023-001 | | 2025 prel. This document | |
|----------------------------------|-----------------------------|---------------------------|-------|-----------------------------|-------|
| Run1 + Run2 | RAW data (pp+HI+fix target) | 36.9 | 79.4 | 36.9 | 80.4 |
| | RDST data (pp+HI+fixtarget) | 13.8 | | 13.8 | |
| | ARCHIVE | 28.7 | | 29.7 | |
| Run3 | pp data (FULL+TURBO+TURCAL) | 144.0 | 171.0 | 270.0 | 317.1 |
| | minimum bias / no-bias | 0.6 | | 0.6 | |
| | Heavy Ion + fixed target | 11.2 | | 16.8 | |
| | ARCHIVE (data+MC) | 15.1 | | 3.7 | |
| Total | | 250.4 | | 397.5 | |
| LHCb-TDR-018 | | 348.0 | | 351 | |

Preliminary 2025 requests

- 60% increase with respect to 2024 request
- Long tail of the R1+R2 simulation
- **HL+fixed target** is representing about 10% of required resources

| LHCb | | 2024 | | 2025 (prelim.) | |
|-------------|--------|---------|-----------------------|----------------|-----------------------|
| | | Request | 2024 req./2023 pledge | Request | 2025 req. / 2024 CRSG |
| WLCG CPU | Tier-0 | 174 | 81% | 283 | 162% |
| | Tier-1 | 572 | 96% | 928 | 162% |
| | Tier-2 | 319 | 74% | 518 | 162% |
| | HLT | 50 | 100% | 50 | 100% |
| | Sum | 1115 | 89% | 1779 | 160% |
| Others | | 50 | 100% | 50 | 100% |
| Total | | 1,165 | 93% | 1,829 | 157% |
| Disk | Tier-0 | 30.6 | 101% | 49.2 | 161% |
| | Tier-1 | 61.2 | 112% | 98.5 | 161% |
| | Tier-2 | 11.8 | 150% | 19.0 | 161% |
| | Total | 103.6 | 111% | 166.7 | 161% |
| Tape | Tier-0 | 117.1 | 129% | 189.3 | 162% |
| | Tier-1 | 133.3 | 99% | 208.1 | 156% |
| | Total | 250.4 | 111% | 397.5 | 159% |

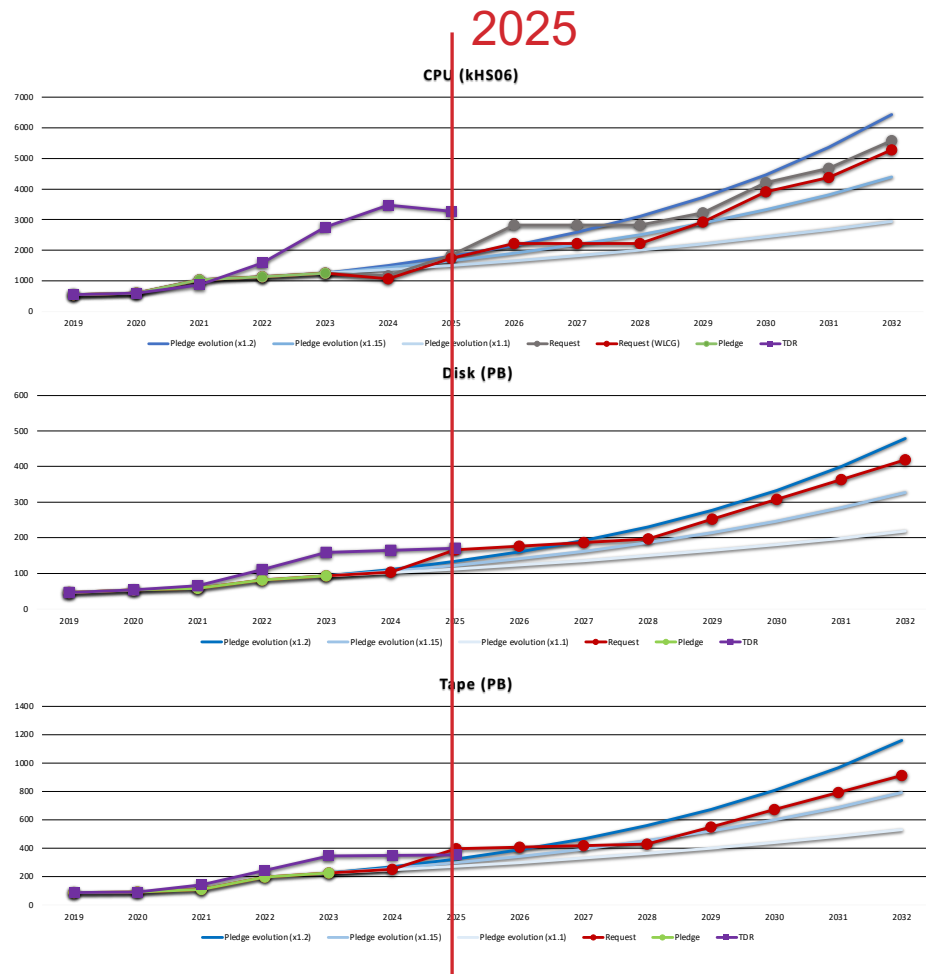
Disk space at T1/T2 sites

- The division of the disk space request between T1 and T2-D sites is not strict
 - Prefer to concentrate disk storage on fewer sites to help data manager's operations
 - Countries without T1 are welcome to provide T2-D sites
- T2-D sites are mostly running user analysis jobs
 - Data stored can depend on analysis production procedures
 - Often results in a low occupancy of T2-D storage

Outlook for Run4

- 2024 requests ~ 2023 requests
- resource profile flattens in LS3 and returns to the “constant budget” line in Run4

C. Bozzi. LHCb



Conclusions

- Smooth operations: 90% MC + 10% data (re)processing and user analysis
- New T1 centers (NCBJ,IHEP) will be soon going into production contributing ~5% of resources each
- Stable use of opportunistic resources on few percent level (non-pledged, HPC)
- Tokens usage is following the WLCG schedule
- 2025 request is showing 60% increase
- Flattening request for LS3 will result in returning to the flat budget curve for Run4

Back-up slides

Storage usage

Overall, LHCb are using ~135PB, an increase of ~5PB over the last quarter

This is divided between Disk (~62PB) and Tape (~73PB). Note that the disk value is still ~18PB over the requested amount for 2023

Transfer rates peaked at 50 GB/s with an average between 5-10 GB/s

