# LHCb Computing Resources: 2025 requests

## LHCb Public Note

Issue: 0
Revision: 1

Reference: LHCb-PUB-2024-002
Created: 31st January 2024
Last modified: 29th February 2024

**Prepared By:** LHCb Computing Project
C. Bozzi/Editor

# Abstract

This document presents an estimate of the offline computing resources needed by LHCb in 2025. The computing requests are based on the Computing Model Technical Design Report for the LHCb Upgrade [LHCb-TDR-018], adjusted to the currently known LHC running schedule and the expected activities to be performed by the LHCb experiment.

# Table of Contents

# List of Tables

# 1. Introduction

This document presents an estimate of the LHCb experiment computing resources requirements for the 2025 WLCG year.
Section 2 recaps the major features of the LHCb computing model for Run3 and the main drivers of the offline computing resource requests. Section 3 shows the assumptions that have been made regarding the LHC running scenario and the LHCb plans for data taking. Section 4 presents the 2025 requests, with a summary given in Section 5. Concluding remarks are given in Section 6. An estimate for the long-term evolution of LHCb computing resources is given in Section 7. Replies to the C-RSG recommendations are shown in Section 8.

# 2. Computing model for LHCb in Run 3

The Computing Model for LHCb in Run 3 and its physics foundations are thoroughly discussed in a Technical Design Report [LHCb-TDR-018]. This section presents a recap of its basic features.

## 2.1. Basic features of the LHCb Computing Model

The concepts that were developed and implemented during the Run 2 data taking become predominant for Run3 data taking.

- The splitting of the High-Level Trigger in two parts, synchronous (HLT1) and asynchronous (HLT2) with data taking, enables the final detector alignment and calibration to be performed online in real time, thereby allowing for an offline-quality event reconstruction in HLT2 and avoiding almost completely the necessity of a costly offline reconstruction.
- The trigger system is entirely based on software. This increases the trigger efficiency for most of the physics programme by at least a factor 2. Furthermore, a five-fold increase of the instantaneous luminosity and the fact that the trigger selects signals with high purity, increases the event throughput to offline storage by at least an order of magnitude.
- From the processing flow point of view, the majority (70%) of triggered events are sent to the TURBO stream, where only high-level information (e.g., tracks, production and decay vertices, particle ID information) is saved to offline and the raw events are discarded. This is the case of events selected by exclusive trigger lines, as in the case of e.g., charm decays.
- A mechanism of “selective persistency” allows to tailor the quantities to be saved on storage on a per-trigger-line basis, ranging from e.g., two charged tracks to the entire event.
- More inclusive trigger lines as well as calibration lines (about 30% of the total) are saved in the “classic” FULL and TURCAL streams, where the entire event is persisted. The FULL stream is then further processed offline, where slimming and filtering criteria are applied, aimed respectively at saving only the interesting parts of the event (selective persistency) and increasing signal purities, thus reducing the footprint on disk storage.
- An additional offline event reconstruction is run only on part of the data corresponding to use cases such as detector commissioning, reconstruction studies, and to reconstruct streams that cannot be reconstructed online (e.g. due to timing constraints). These involve mainly events in the TURCAL stream.

From the previous points, it follows that the CPU needs are dominated by Monte Carlo simulation. As CPU work scales according to the integrated luminosity and pile-up, a detailed Geant4-based simulation of the detector would require at least a ten-fold increase in the resources. Faster simulation options are employed to mitigate the CPU requirements (see below).

The main data workflows are thus:

- The processing of the TURBO stream data to convert the LHCb-specific online format to the ROOT I/O-based offline format, and the subsequently streaming of these data. This workflow accounts for 0.01% of the CPU work on the Grid.
- The slimming and filtering of data in the FULL stream, and their subsequent streaming. Also in this case, the expected CPU work on the Grid is no more than a few percent of the total.
- The processing of the TURCAL stream, which is assumed to represent a small fraction of the CPU work as well.

## 2.2. Major drivers in offline resource requirements

As mentioned in Section 2.1, the production of simulated events dominates the offline CPU computing needs. LHCb has mitigated this by exploiting faster simulation options. In “ReDecay” the same underlying event is used several times (the default being 100) and only signals are generated and simulated each time. This simulation option, already in production, accounts for about 2/3 of the total simulated samples since several years. Another option, where only the response of the tracking detectors is simulated, has been in production as well and successfully used by analyses not requiring costly simulations of the calorimeters and the RICH detectors.
More fast simulation workflows are under preparation, such as the utilization of shower libraries and/or machine learning techniques to parametrize the response of the calorimeters. A full parametric simulation is also in development.
In all the above cases, the simulation workflow starts with events generation and the simulation of the detector response, where the latter accounts for the vast majority of computing work. The subsequent steps are the digitization of the detector signals and the emulation of the trigger.
The simulation is being adapted to run in a multi-threaded environment. This enables a significant reduction of the memory footprint, thereby opening the possibility to use resources, such as HPC farms and many-core architectures, where the memory per logical core is smaller than that of the usual grid computing nodes.

The storage needs are dominated by data and crucially depend on the HLT output bandwidth. A bandwidth of 10GB per live second of LHC is deemed sufficient to carry on the physics programme of LHCb. While the associated tape needs are incompressible, mitigations are possible for disk. As mentioned in Section 2.1, about 70% of triggered events are saved in the light TURBO format. However, the majority (7.5GB/s out of 10GB/s) of the bandwidth is taken by the remainder 30% of events in the FULL and TURCAL streams, where the entire event is saved. The events in these two latter streams are therefore slimmed and/or filtered offline, in a process dubbed *sprucing*[†], such that the total (logical) bandwidth to be saved on disk is only 3.5GB/s. Table 2-1 shows the extrapolated throughputs to tape and disk for the three data streams that are used.

---

[†] *sprucing* has replaced the Run1+Run2 *stripping*, *i.e.* a workflow by which events in the FULL and TURCAL streams are *skimmed* according to sets of selection criteria (*lines*) and the event content is *slimmed* to a size comparable to that of an event in the TURBO stream.

| 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% |

*Table 2-1: [taken from LHCb-TDR-018] Extrapolated throughput to TAPE and to disk (after offline processing), for the FULL, TURBO and CALIBRATION streams.*

The impact of simulated events on storage requests is small, as data produced during the intermediate steps are deleted, only the relevant information is persisted at the end, and analysis-dependent filtering criteria are generally applied.

## 2.3. Offline resource needs

The following basic assumptions enter in the calculation of the offline resource needs:

- Trigger output bandwidth, scaling with instantaneous luminosity and trigger rate, mitigated by processing online as much data as possible in the TURBO stream, and by an aggressive offline data reduction of the FULL and TURCAL streams.
- Simulation of Run 1 + Run 2 has negligible impact; the bulk of the simulation of a given year of Run 3 data taking starts slowly during that year, reaches the nominal level in the following year and, stays steady for the following 4 years, ramps down to 50% the year after and to zero afterwards.
- A mixture of full/fast/parametric simulations.
- Most of the simulation output is selectively persisted and aggressively filtered.
- The Run3 timeline is very different from the one that had been assumed in [LHCb-TDR-018]. For LHCb, 2022 and most of 2023 have been years of commissioning of the sub-detectors, many of which are new, and the software trigger system. An incident in the LHC vacuum system near the LHCb VELO detector damaged the RF foil that shields the VELO from the beam. This implied that the VELO was operated in the open position throughout 2023, with impact on the detector acceptance and resolution. The LHC operations in 2023 and 2024 have been shortened as well, following the global energy crisis. An LHC incident in mid-July 2023 brought the proton run to an anticipated end. The LHC resumed in Autumn 2023 and successfully provided heavy ion collisions. Nominal conditions are foreseen in 2024 (although with a shortened running time) and 2025. The third LHC long-shutdown (LS3) starts in 2026 and will last three years.

The basic parameters of the LHCb computing model are reported in Table 2-2.

| **Model assumptions for 2025** | |
|---|---|
| L ($cm^{-2}s^{-1}$) | $2 \times 10^{33}$ |
| Pileup | 6 |
| Running time *pp* collisions (s) | $6.3 \times 10^{6}$ |
| 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 |

*Table 2-2: Summary of the main assumptions of the LHCb computing model for 2025.*

# 3. LHC running scenario and LHCb data taking plans in 2025

The LHC schedule for 2025, provided by the LPC to the LHC experiments on June 26th, 2023, foresees a LHC running time of 6.3 $10^6$ seconds for proton collisions and <1.7 $10^6$ seconds of heavy-ion collisions in the 2025 calendar year, with an integrated luminosity for proton collisions at LHCb of less than 15 $fb^{-1}$.

The schedule assumes that the ion run in 2024 and/or 2025 could be extended to 5 weeks. Currently, 4 weeks of PbPb is foreseen for both years, but a short p-Pb run in one year and longer PbPb run in the other year is also a possibility. Finally, an additional 5 days of oxygen-oxygen (OO) collisions is also expected in 2025.

It is assumed that the throughput from the trigger farm to the offline system during pp collisions, the other parameter driving the offline storage requests in addition to the LHC live time, will be the nominal one (10GB per live second of the LHC) during the entire period foreseen for proton collisions in 2025.

LHCb plans to take heavy-ion collision data in 2025. In this document, an enhanced configuration (i.e. higher retention rate) with respect to the PbPb 2023 runs is foreseen, both for the oxygen and lead runs. Like in 2023, the future ion data will be processed mostly via the FULL stream, with selections made at the HLT1/HLT2/Sprucing levels. Fixed-target collisions using the SMOG2 system will also be recorded simultaneously for both lead and oxygen beams.

# 4. Resource requests for 2025

In this Section, the 2025 requests are presented. The 2024 pledges from the funding agencies, and a reassessment of the resource required for 2024, following the reality of 2023 data taking [LHCb-PUB-2024-003], are also shown for completeness.

## 4.1. CPU requests

1. For *sprucing* (both first pass and end-of-year re-sprucing), the CPU work to spruce one event in Run3 conditions is taken as the same as for an event during Run2.

2. no provision is made for the offline reconstruction of heavy-ion collision data, which is assumed to performed on the online farm.
3. simulation consists of two parts, the former dominating over the latter:
    a. The *simulation of Run3 pp collision data* follows the prescriptions made in the Computing Model TDR. In particular, the simulation of 2024 collisions will ramp up to the nominal level (4.8 $10^9$ events per *fb*$^{-1}$ per calendar year), while that of 2025 collisions will be at 50% of the nominal level. We assume also that the simulation of 2023 conditions will end in 2024, as the corresponding dataset is taken in non-standard detector conditions and will probably become obsolete in 2025. Following considerations that have already been reported earlier, we take the same event simulation time of Run2 Monte Carlo.
    b. *The simulation of Run3 heavy ion and fixed target collision data* is assumed to require 10% of the total work needed for the reconstruction of the real data counterpart.
    c. It is expected that the *simulation of Run1+Run2 pp collision data* in 2024 will require minimal additional requests.

A summary of the various parameters entering the CPU request corresponding to simulation is given in Table 4-1.

| | Run3 pp | Run3 HI |
|---|---|---|
| CPU work simulations 2025 (kHepScore23.y) | 1242 | 93 |
| Total number of events simulated in 2025 ($10^9$) | 80 | |
| Fraction full simulation | 0.36 | |
| Fraction fast simulation | 0.64 | |
| Fraction parametric simulation | 0.0 | |
| CPU work per event full simulation (kHepScore23.s) | 1.2 | |
| CPU work per event fast simulation (kHepScore23.s) | 0.12 | |
| CPU work per event parametric simulation (kHepScore23.s) | 0.02 | |

*Table 4-1: Summary of parameters entering the determination of the CPU work needed for simulation.*

4. The CPU work for user analysis in Run2 was found to scale with the CPU work for stripping. This is expected, as user jobs are principally processing data produced by the stripping. The same criterium is applied to analysis jobs in Run3, however with a 50% reduction factor. This considers (i) the fact that, according to the Computing Model TDR, most of the user analysis will be centrally managed with analysis productions and therefore with a much lower failure rate, and (ii) that the analysis framework has been completely reorganized, with emphasis given on CPU performance. Numerically:
    a. *Sprucing* work for 24+25 data: (68+82) = 150k HepScore23.y
    b. Required work: 150kHepScore23.y * 3.74 (Scaling factor analysis/stripping) / 2 (improvement over Run2) = 281 kHepScore23.y
    c. We then assume that there will be a residual tail of Run2 analysis, by taking half of the corresponding work measured during Run2: 75kHepScore23.y/2 = 38kHepScore23.y

   d. The sum of Run3+Run2 analysis work gives then 281+38 = 319 kHepScore23.y

5. LHCb uses O(100) virtual machines to support its offline computing infrastructure, for core services such as the build and nightly systems, software databases, messaging, and distributed computing services and agents. For 2025, this infrastructure requires 10kHepScore23.

A summary of the preliminary CPU requirements for 2025 is given in Table 4-2. With respect to the 2024 requests, the most important increase is due to simulation, namely that of Run3 collisions.

The CPU work that LHCb will get from the HLT farm in 2025 will be low, as the HLT farm will be used almost entirely for data taking activities during the LHC run, and for reconstructing heavy ions collision data during the (E)YETS.

| **CPU Work in WLCG year (kHepScore23.years)** | **2024**<br>**LHCb-PUB-2023-001** | **2024**<br>**THIS DOCUMENT** | **2025 prel.**<br>**LHCb-PUB-2023-003** | **2025**<br>**THIS DOCUMENT** |
|---|---|---|---|---|
| First pass sprucing | 70 | 68 | 82 | 82 |
| End-of-year sprucing | 70 | 68 | 82 | 82 |
| Simulation | 800 | 535 | 1336 | 1336 |
| Core and distributed computing infrastructure | 10 | 10 | 10 | 10 |
| User Analysis productions | 214 | 171 | 319 | 319 |
| **Total Work (kHepScore23.years)** | **1165** | **851** | **1829** | **1829** |
| **LHCb-TDR-018** | **3470** | **3470** | **3276** | **3276** |

*Table 4-2: Estimated CPU work needed for the different activities in 2025 (column "**2025 This document")**. The other columns show the 2024 requests, endorsed at the April 2023 RRB (column "**2024 LHCb-PUB-2023-001"), and** reassessed following the reality of 2023 data taking (column "**2024 This document"),** the preliminary 2025 requests (column "**2024 LHCb-PUB-2023-003");** the last row (**"LHCb-TDR-018"**) reports a comparison with the computing model TDR.*

## 4.2. Disk requests

Table 4-3 presents, for the different data classes, the forecast usage of disk space at the end of 2024. The various terms are due to:

1. Legacy Run1 and Run2 data, and their corresponding MC samples, in a single copy.
2. Data from Run3 pp collisions; the request is determined according to the Run3 Computing Model TDR; more specifically:
   a. the total throughput to disk is 3.5GB per "LHC live second", i.e., for each second LHC is giving stable beam collisions = 0.8 (FULL) + 2.5 (TURBO) + 0.2 (TURCAL), see Table 2-1.
   b. the LHC live time is assumed to be 6.3 million seconds.
   c. we save on disk 2 copies of TURBO stream, 2 copies of the latest (FULL+TURCAL) processing, 1 copy of the previous (FULL+TURCAL) processing.

2. data from ion-ion and fixed target collisions, and corresponding simulations; this disk provision is made by assuming:
   a. For ion-ion collisions: 6.2 billion triggered events (average size of 142kB/event) in the FULL stream.
   b. For fixed-target collisions: 187 billion triggered events (average size of 142kB/event) at the HLT1 level, with a retention rate of 10% passing HLT2/sprucing on the FULL stream.
   c. The above numbers apply also to the 2025 lead runs; in addition, the oxygen-oxygen runs are added as follows:

   i. For oxygen-oxygen collisions: 2.7 billion events (average size of 60kB/event) in the FULL stream.
   ii. For fixed-target collisions: 127 billion triggered events (average size of 60kB /event) at the HLT1 level, with a retention rate of 10% passing HLT2/sprucing on the FULL stream.

3. Run3 simulation of pp collisions, determined by following the Computing Model TDR with the same assumptions of point 3.a of Section 4.1 above.
4. User data and grid buffer data. The former (3.6PB) has been estimated by taking the sum of the annual increments observed in Run2 and LS2 for the analysis of existing data (0.1PB/year), and by assuming that the yearly increment of the space needed for a nominal year of Run3 data taking scales by a factor five. The latter has been estimated by assuming it is driven by the re-sprucing at the end of the year, and that the tape recall bandwidth (see below) can cope with the re-sprucing processing rate, allowing for a contingency of two weeks. We assume that re-sprucing, which involves a total of 47PB of data to be recalled from tape, will last two months. A contingency of two weeks would therefore correspond to a grid buffer space of 12PB.
5. Following the experience with the 2023 data taking and the associated shortage of disk space, further discussed in [LHCb-PUB-2024-003], a buffer of 10PB is requested at the Tier0. This serves two purposes:
   a. The temporary storage of data, coming from the online system, prior to storing them on the CTA tape system at CERN; a provision of 6.5PB guarantees a contingency of the order of one week;
   b. The storage of commissioning data, taken in 2022 and 2023, that are deemed to be important for studying detector performance, calibration, and alignment. This storage area of 3.5PB will be cleaned up and no longer required at the end of the 2025 data taking.

| **Disk storage usage forecast (PB)** | | **2024** | | **2024** | | **2025 prel.** | | **2025** | |
|---|---|---|---|---|---|---|---|---|---|
| | | **LHCb-PUB-2023-001** | | **This document** | | **LHCb-PUB-2023-003** | | **This document** | |
| Real data | Run1+Run2 pp data | 10.2 | 78.7 | 10.2 | 62.4 | 10.2 | 134.4 | 10.2 | 121.6 |
| | Run1+Run2 HI+SMOG | | | | | | | | |
| | Run3: FULL | 16.5 | | 12.5 | | 30.9 | | 26.9 | |
| | Run3: TURBO | 36.3 | | 27.5 | | 68.1 | | 59.3 | |
| | Run3: TURCAL | 4.5 | | 3.4 | | 8.4 | | 7.3 | |
| | Run3: Minimum bias | 0.0 | | 0.0 | | 0.0 | | 0.0 | |
| | Run3: HI+SMOG2 | 11.2 | | 8.8 | | 16.8 | | 17.9 | |
| Simulated data | Run1+Run2 Sim | 8.7 | 11.9 | 8.7 | 10.6 | 8.7 | 16.7 | 8.7 | 15.0 |
| | Run3 simulated data | 3.2 | | 1.9 | | 8.0 | | 6.3 | |
| Other | User data | 3.0 | 13.0 | 3.0 | 23.0 | 3.6 | 15.6 | 3.6 | 25.6 |
| | Grid Buffers | 10.0 | | 10.0 | | 12.0 | | 12.0 | |
| | Tier0 Buffer | 0.0 | | 10.0 | | 0.0 | | 10.0 | |
| **Total** | | | **103.6** | | **96.0** | | **166.7** | | **162.2** |
| | **LHCb-TDR-018** | | **165.0** | | **165.0** | | **171.0** | | **171.0** |

*Table 4-3: Disk Storage needed in 2025 for the different categories of LHCb data (column “**2025 This document”**). The other columns show the 2024 requests, endorsed at the April 2023 RRB (column “**2024 LHCb-PUB-2023-001”**), and reassessed following the reality of 2023 data taking (column “**2024 This document”**), the preliminary 2025 requests (column “**2024 LHCb-PUB-2023-003”**); the last row (**“LHCb-TDR-018”**)[‡] reports a comparison with the computing model TDR.*

---

[‡] Please note that in LHCb-TDR-2018 it is assumed that 2024 and 2025 would have been shutdown years for the LHC, hence only a small increase of disk storage was foreseen.

## 4.3. Tape requests

The forecast usage of tape space (Table 4-4) is the sum of:

1. The tape needed by the Run1+Run2 real (RAW+RDST+ARCHIVE) data at the end of 2022 and the simulated (ARCHIVE) data until the end of 2023, for a total of 81.9PB. This includes a contribution of 1.5PB to the ARCHIVE, due to the incremental stripping of Run2 data performed in 2023, which was unforeseen in previous requests.
2. The tape needed by the Run3 proton collision data, heavy-ion, and fixed target data, minimum bias / no-bias stream, and Run3 simulation. This request is dominated by pp data (FULL+TURBO+TURCAL), for which we assume an amount of data to be taken in 2025 of 10GB/s times 6.3 million seconds LHC live time = 63PB times 2 copies, for a total of 126PB.

| **Tape storage usage forecast (PB)** | | **2024**<br>**LHCb-PUB-2023-001** | | **2024**<br>**This document** | | **2025 prel.**<br>**LHCb-PUB-2023-003** | | **2025**<br>**This document** | |
|---|---|---|---|---|---|---|---|---|---|
| Run1 + Run2 | RAW data (pp+HI+fix target) | 36.9 | 79.4 | 36.9 | 80.9 | 36.9 | 80.4 | 36.9 | 81.9 |
| | RDST data (pp+HI+fixtarget) | 13.8 | | 13.8 | | 13.8 | | 13.8 | |
| | ARCHIVE | 28.7 | | 30.2 | | 29.7 | | 31.2 | |
| Run3 | pp data (FULL+TURBO+TURCAL) | 144.0 | 171.0 | 109.0 | 134.1 | 270.0 | 317.1 | 235.0 | 283.3 |
| | minimum bias / no-bias | 0.6 | | 0.6 | | 0.6 | | 0.6 | |
| | Heavy Ion + fixed target | 11.2 | | 13.3 | | 16.8 | | 22.4 | |
| | ARCHIVE (data+MC) | 15.1 | | 11.2 | | 3.7 | | 25.3 | |
| **Total** | | | **250.4** | | **215.0** | | **397.5** | | **365.2** |
| | **LHCb-TDR-018** | | **348.0** | | **348.0** | | **351.0** | | **351.0** |

*Table 4-4: Tape Storage needed in 2025 for the different categories of LHCb data (column “**2025 This document”**). The other columns show the 2024 requests, endorsed at the April 2023 RRB (column “**2024 LHCb-PUB-2023-001”**), and reassessed following the reality of 2023 data taking (column “**2024 This document”**), the preliminary 2025 requests (column “**2024 LHCb-PUB-2023-003”**); the last row (**“LHCb-TDR-018”**[§]) reports a comparison with the computing model TDR.*

---

[§] In LHCb-TDR-018, 2024 and 2025 were assumed to be shutdown years for the LHC (LS3)

# 5. Summary of 2025 requests

Table 5-1 shows the preliminary CPU (in kHepScore23.y), disk (in PB), and tape (in PB) requests for 2025, together with the endorsed 2024 requests, at the various tiers, as well as for the HLT farm and other opportunistic resources. The increase of the 2025 requests with respect to 2024 resources endorsed by the RRB in April 2023 are also shown. They are at the 60% level for CPU and disk, and at the 50% level for tape.

| LHCb | | 2024 | | | | | 2025 | |
|---|---|---|---|---|---|---|---|---|
| | | Request | Pledge | Pledge/req | 2024 req./ 2023 CRSG | 2024 req. / 2023 pledge | Request | 2025 req. / 2024 CRSG |
| **WLCG CPU** | **Tier-0** | 123 | 174 | 141% | 57% | 57% | 283 | 162% |
| | **Tier-1** | 404 | 542 | 134% | 57% | 68% | 928 | 162% |
| | **Tier-2** | 224 | 394 | 176% | 57% | 52% | 518 | 162% |
| | **HLT** | 50 | | | 100% | 100% | 50 | 100% |
| | **Sum** | 801 | 1110 | | 59% | 62% | 1779 | 160% |
| **Others** | | 50 | | | 100% | 100% | 50 | 100% |
| **Total** | | **851** | **1,110** | **130%** | **60%** | **63%** | **1,829** | **157%** |
| **Disk** | **Tier-0** | 35.4 | 30.6 | 86% | 117% | 117% | 54.9 | 180% |
| | **Tier-1** | 50.8 | 53.0 | 104% | 84% | 93% | 89.9 | 147% |
| | **Tier-2** | 9.8 | 9.4 | 96% | 84% | 124% | 17.4 | 147% |
| | **Total** | **96.0** | **93.0** | **97%** | **94%** | **103%** | **162.2** | **157%** |
| **Tape** | **Tier-0** | 97 | 117 | 121% | 107% | 107% | 170.4 | 146% |
| | **Tier-1** | 118 | 125 | 106% | 75% | 88% | 194.8 | 146% |
| | **Total** | **215.0** | **242.2** | **113%** | **87%** | **96%** | **365.2** | **146%** |

*Table 5-1: Evolution of offline computing requests in 2024-2025. Units are kHepScore23 for CPU, PB for disk and tape. The column "2024 request" corresponds to a reassessment of the 2024 requirements, following the reality of the 2023 data taking.*

# 6. Conclusion

This report summarizes a preliminary assessment of the offline computing requests needed by LHCb in 2025, utilising updated information on the LHC running conditions, and on the LHCb data taking plans. A summary of the requests is given in Table 6-1 for CPU, Table 6-2 for disk and Table 6-3 for tape, together with the 2024 resources endorsed at the April 2023 RRB.

For CPU, we assume that the HLT farm will be partly available during the winter shutdowns and not available during the LHC run, and that the opportunistic contributions will provide the same level of computing power as in the past, therefore we subtract the contributions from these two sites from our requests to WLCG. The required CPU resources are apportioned between the different Tiers considering the capacities that are already installed. The disk and tape estimates are broken down into fractions to be provided by the different Tiers using the distribution policies described in LHCb-PUB-2013-002.

We thank the C-RSG for their support and guidance.

| **CPU Power (kHepScore23)** | **2024** **pledges** | **2025** |
|---|---:|---:|
| Tier 0 | 174 | 283 |
| Tier 1 | 542 | 928 |
| Tier 2 | 394 | 518 |
| **Total WLCG** | **1110** | **1729** |
| | | |
| HLT farm | 50 | 50 |
| Opportunistic | 50 | 50 |
| **Total non-WLCG** | **100** | **100** |
| | | |
| **Grand total** | **1210** | **1829** |

*Table 6-1: CPU power requested at the different Tier levels in 2025. The 2024 pledges are also shown.*

| **Disk (PB)** | **2024** **pledges** | **2025** |
|---|---:|---:|
| Tier0 | 30.6 | 54.9 |
| Tier1 | 53.0 | 89.9 |
| Tier2 | 9.4 | 17.4 |
| **Total** | **93.0** | **162.2** |

*Table 6-2: LHCb Disk request for each Tier level in 2024 in 2025. The 2024 pledges are also shown. For countries hosting a Tier1, the Tier2 contribution could also be provided at the Tier1.*

| **Tape (PB)** | **2024** **pledges** | **2025** |
|---|---:|---:|
| Tier0 | 117 | 170 |
| Tier1 | 125 | 195 |
| **Total** | **242** | **365** |

*Table 6-3: LHCb Tape request for each Tier level in 2025. The 2024 pledges are also shown.*

# 7. Long-term evolution of LHCb computing resources

A long-term forecast of the LHCb computing requirements is shown in this section, to demonstrate that they will, in this long term, remain within canonical assumptions for increases in capacity. We take the expected increases in capacity to be between 10-20% per annum, driven by "Flat Cash" and referred to below as FC lines at yearly 10%, 15% and 20% increases.

The three figures below show, for CPU, Disk and Tape respectively:

- The FC-curves in different shades of blue.
- A hybrid line composed of:
    - For past years, including 2024: the actual pledged capacities in green.
    - For future years the projected total requirements in grey for CPU.
    - For future years the projected request to WLCG, allowing for the HLT farm, in red.
- The requests written in the LHCb Upgrade Computing Model TDR [LHCb-TDR-018] in purple.

A normalisation year of 2023 is used as this is the existing situation.

The CPU requirements are shown in Figure 1.

- The small drop between 2023 → 2024 reflects not only the updated LHC running schedule, but also the effects of the LHCb VELO event in 2023.
- The increase from 2024 → 2025 does indeed show a local year-on-year increase but, following the reassessment of 2023 and 2024 requests, this increase no longer exceeds the FC-curves.
- After 2026 our requests flatten and stay within the FC-curves until the end of Run4.

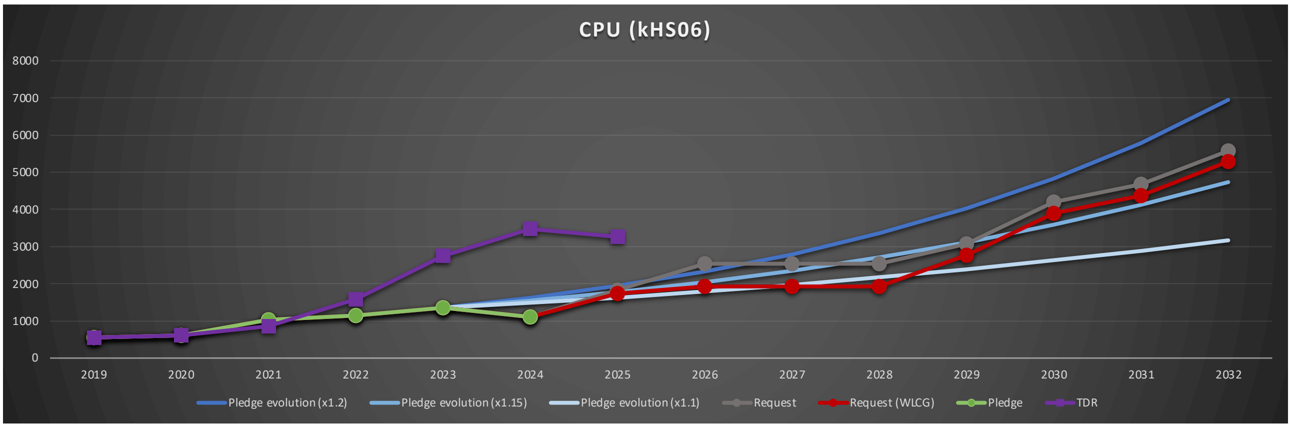

*Figure 1: LHCb CPU requirements. Blue lines show the FC-bands. Green shows actual pledges, Grey shows total projected requirements. Red shows WLCG requirements allowing for our HLT farm. Purple shows the requests made in the LHCb Upgrade Computing Model TDR.*

The Disk Storage request is show in Figure 2. Very similar comments pertain as made for the CPU request. There is a larger increase in 2025, exceeding the FC-curves, balanced by a long flat period which brings us back within the FC-curves by 2026 or 2027.

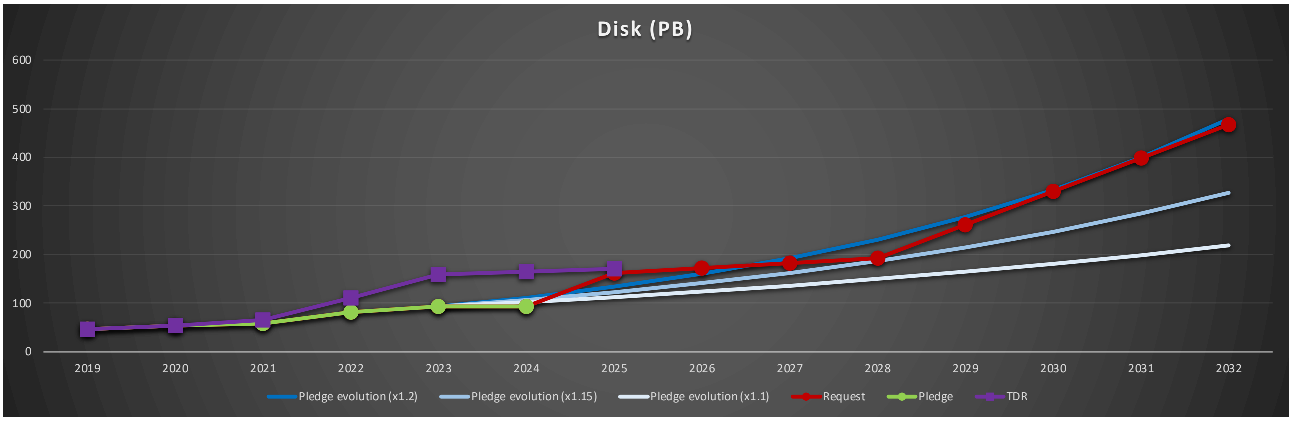

*Figure 2: LHCb Disk Storage requirements. Blue lines show the FC-bands. Green shows actual pledges, Red shows total projected requirements. Purple shows the requests made in the LHCb Upgrade Computing Model TDR.*

The Tape Storage requirements are shown in Figure 3. Similar comments as those made for disk apply.

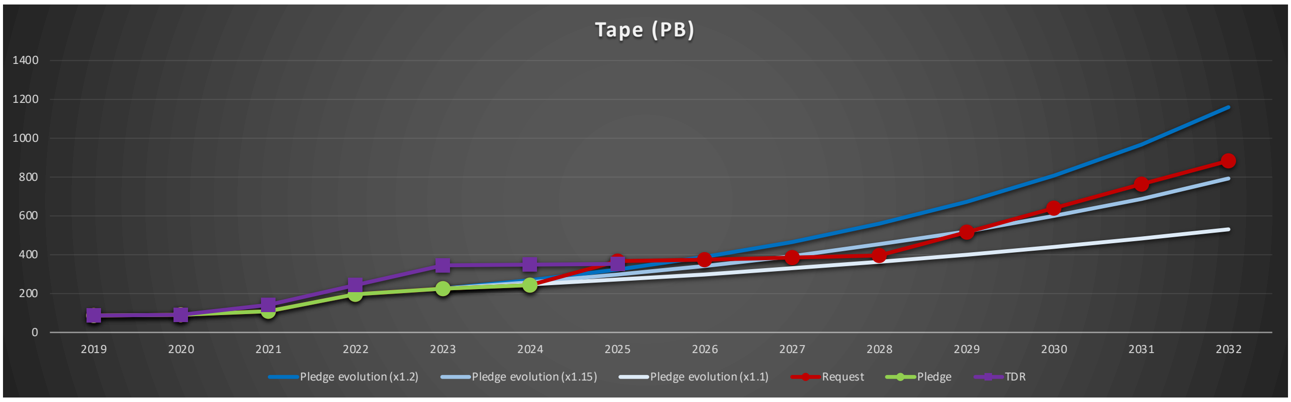

*Figure 3: LHCb Tape Storage requirements. Blue lines show the FC-bands. Green shows actual pledges, red shows total projected requirements. Purple shows the requests made in the LHCb Upgrade Computing Model TDR.*

LHCb understands that some countries are able and willing to provide larger steps in pledges in any given year, provided they will then make a commensurate smaller step in later years (i.e., buying ahead) and we welcome and thank them for this flexibility.
LHCb also understands that some countries are unable to do this and prefer to provide a smoother profile, and we recognise these constraints and thank them for what they can provide.

The information shown is qualitatively similar to what LHCb published in its TDR in 2018.

LHCb hopes that this information will allay some of the worries that been expressed by oversight bodies in respect of large year-on-year fractional increase requests submitted by LHCb. It is hoped that the "new detector pulse" effect set in this long-term context will allow oversight bodies and funding bodies to have confidence that LHCb requests remain approximately within "flat cash" limits.

# 8. Appendix: replies to the C-RSG recommendations

The C-RSG requested that *"the experiments provide a section that responds to the recommendations from the previous scrutiny. This response should address both the experiment specific recommendations and general recommendations relevant to the experiment."*

This appendix reports the actions that have been taken for each LHCb recommendation.

***LHCb-1*** *Considering that the CPU requirements are dominated by MC simulation for Run 3 data, the C-RSG encourages the LHCb Collaboration to explore strategies to reduce the simulation CPU footprint.*

LHCb is exploring several possibilities to reduce the simulation CPU footprint, complementing those already in use.

One goal is to replace detailed simulation in the most expensive parts of the detector (calorimeters and RICH systems) with faster counterparts, while producing similar types of output.

The LHCb simulation framework provides a dedicated custom simulation "hook" based on the Geant4 fast simulation interface to replace it with parametrised and machine learning models. Fast simulations can be enabled according to detector region, particle type and kinematic conditions. The interface is also being explored as a hook to enable the use of GPUs and exploit heterogeneous computing.

LHCb collaborates with the EP-SFT/Geant4 group in the context of the CaloChallenge initiative to find the best ML model for the generation of calorimeter showers, using a common benchmark geometry. Once a model is chosen, it can be retrained on the LHCb geometry, and implemented exploiting the interface mentioned above. The infrastructure to adopt new models for the LHCb calorimeter is essentially ready. A model based on a Variational Auto Encoder (VAE) has been tested up to the level of physics analysis. The differences between this fast simulation and the detailed simulation are already very small in some representative quantities and could be further reduced with additional training. Preliminary results indicate physics performance and trigger efficiencies to be in line with what is observed on the detailed simulation side for a few selected signal channels.

LHCb is also exploring the use of AdEPT, a GPU application for the simulation of EM showers. There are ongoing efforts with the CERN/EP-SFT group to integrate AdEPT into the LHCb simulation framework via the Geant4-based custom simulation hook, with the goal of offloading $e+/e-/\gamma$ (in specific regions) to GPUs and exploit the GPU's parallel computational capabilities. In this sense, "fast simulation" means a detailed simulation on a fast hardware, whereby the particles entering the calorimeter region fill the AdEPT pipeline, which then gives back the simulated information to the LHCb simulation framework for further manipulation.

Tests were also performed with the OPTICKS application, which provides an interface between Geant4 and the NVIDIA OptiX ray tracing engine to simulate photon propagation while maintaining the simulation of other particles on CPU. A simplified version of the LHCb RICH1 detector was used to validate the performance of OPTICKS and check its consistency with Geant4. The integration of OPTICKS into the LHCb simulation framework presents a considerable challenge which requires further effort. OPTICKS uses external packages; specific versions are required to avoid conflicts for use in distributed computing resources. Mitsuba, a graphical rendering software, compilable with LLVM or CUDA, is being investigated as a more portable and robust alternative to OPTICKS. Integration of either product within the LHCb simulation framework may leverage on the custom simulation hook with additional targeted development.

In parallel a novel framework implementing a flash-simulation paradigm via parametric functions and deep generative models is under development for integration within the general LHCb simulation framework. It provides analysis-level variables taking as input particles from physics generators, and parameterising the detector response and the reconstruction algorithms. It could provide samples when even higher simulated statistics may be required. This framework consists of a pipeline of machine-learning-based modules that allows, for selected sets of particles, to introduce reconstruction errors or infer high-level quantities via (non-)parametric functions. Although good agreement is observed by comparing key reconstructed quantities obtained against those from the existing detailed Geant4-based simulation, further validations are required as well as providing the analysis-level variables in the same format at that of the data.

In summary, promising work is currently ongoing, at different levels of maturity.

***LHCb-2*** *A significant contributor to the CPU budget is user analysis. The C-RSG encourages the LHCb Collaboration to increase the utilization of the centrally managed analysis production system that has been shown to improve CPU efficiency.*

Analysis productions (AP) are becoming more popular with respect to past years; Figure 3-8 of the 2023 resource usage report, attached here for completeness, shows that AP are no longer sporadic as previously reported, but more continuous throughout the year. AP are expected to ramp up as Run3 analyses will start. A physiological level of user jobs is nevertheless expected to stay, for prototyping new analyses and to execute workflows not related to data reduction such as toy Monte-Carlo studies and training of ML/AI models.

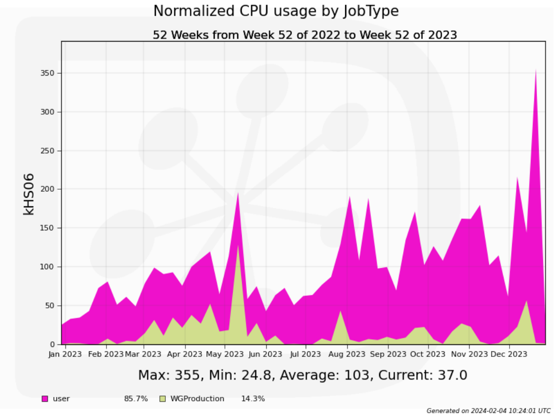

***LHCb-3*** *The physics data taken in 2023 during the pp and HI runs increase the storage requirements for 2025. These storage requests will need to be re-evaluated in the next scrutiny round in light of the actual physics data collected during 2023.*

The storage requests have been revaluated following the experience with data taking in 2023. Less storage is required for the 2023 datasets than previously anticipated, however there is a need for extra buffer space at the Tier0 to cope with operational issues, and to host temporary data that are used for commissioning, calibrating, and aligning the LHCb detector (and that will be deleted at the end of Run3).

## LHCb Risk Register

| Risk | Risk<br>Likelihood | Impact | Severity | Owner | Effect | Mitigation |
|---|---|---|---|---|---|---|
| **Funding** | | | | | | |
| Tape shortage at CERN | 2 | 4 | 8 | CERN | Cannot store data coming from HLT. Descoping of physics program | Decrease HLT throughput to offline by migrating more lines from FULL to TURBO and/or tightening |
| Tape shortage at Tier1 sites | 3 | 4 | 12 | Tier1 sites | | |
| Disk shortage | 3 | 3 | 9 | WLCG sites | Cannot analyse (part of) data taken. Impact on operations load and complexity, system resiliency. Delay in physics program | Park data on cold storage and activate if when feasible. Decide what to store on disk based on popularity. |
| Shortage of computing power | 3 | 2 | 6 | All sites | More time to get adequate simulation samples. Delay of physics analysis. Substantial changes in analysis model. | increase fraction of fast and parametric simulations. Increase pool of opportunistic resources |
| **Operations / Technology** | | | | | | |
| Availability of tape write bandwidth | 2 | 4 | 8 | CERN + Tier1 sites | Cannot store data coming from HLT. Descoping of physics program | Decrease HLT trhoughput. Prioritize physics program |
| Availability of tape read bandwidth | 2 | 4 | 8 | CERN + Tier1 sites | Cannot recall data for end-of-year re-sprucing. Delay of physics analysis | Delay re-sprucing to EYETS or LS3 |
| Underestimation of disk buffer | 3 | 3 | 9 | WLCG | Respruncing slowed-down and extending into data taking period | Delay re-sprucing or get space by temporarily removing other datasets |
| **Software** | | | | | | |
| Underestimation of sprucing work | 2 | 2 | 4 | LHCb | More computing power needed. YETS re-sprucing is delayed | recover computing power by delaying simulation production |
| Availability of sprucing application | 2 | 3 | 6 | LHCb | delay in (re-)sprucing campaigns. Buffer space is not sufficient | park unspruced data on tape. Delay YETS re-sprucing |
| Underestimation of analysis work | 2 | 2 | 4 | LHCb | more computing power needed | recover computing power by delaying simulation production |
| Underestimation of simulation work | 3 | 3 | 9 | LHCb/G4/HSF | Cannot afford sufficient simulation; physics analysis suffers | find optimization point by studying interplay between full/fast/parametric simulation and technological developments in each ofthem |

| **Likelihood:** | **Impact:** |
|---|---|
| *1: never expected to happen* | *1: we can deal with it, no problem* |
| *2: could happen but very unlikely* | *2: a bit of a hassle but not too bad* |
| *3: could well happen* | *3: can be managed, but with significant effort* |
| *4: will probably happen* | *4: crisis* |

# 9. Appendix: High-level Summary of Used and Requested Disk

The current and foreseen usage of disk space in LHCb is shown in the following table. Units are PetaBytes.

| Category | Period | Type | Current (03/02/24) | 2024 request | 2025 request |
|---|---|---|---|---|---|
| Persistent | Run1 + Run2 | pp data | 22.2 | 10.2 | 10.2 |
| | | HI + smog data | 2.7 | | |
| | | **TOTAL** | **24.9** | **10.2** | **10.2** |
| | Run3 | pp FULL | 0.58 | 12.5 | 26.9 |
| | | pp TURBO | 0.05 | 27.5 | 59.3 |
| | | pp TURCAL | 0.02 | 3.4 | 7.3 |
| | | HI + SMOG data | 1.25 | 8.8 | 17.9 |
| | | **TOTAL** | **1.9** | **52.2** | **111.4** |
| | **TOTAL DATA** | | **26.8** | **62.4** | **121.6** |
| | Run1 + Run2 | Monte-Carlo | 11.7 | 8.7 | 8.7 |
| | Run3 | Monte-Carlo | 2.3 | 1.9 | 6.3 |
| | **TOTAL Monte-Carlo** | | **14.0** | **10.6** | **15.0** |
| **TOTAL DATA + MC** | | | **40.8** | **73.0** | **136.6** |
| User | | | 2.8 | 3.0 | 3.6 |
| Cache | | | 0 | 0 | 0 |
| Buffer (Tier0) | | | 11.9 | 10 | 10 |
| Buffer (Grid) | | | 8.2 | 10 | 12 |
| **GRAND TOTAL** | | | **63.7** | **96.0** | **162.2** |