

LHCb Computing Resources: preliminary 2025 requests

LHCb Public Note

Issue: 0
Revision: 1

Reference: LHCb-PUB-2023-003
Created: 1st August 2023
Last modified: 29th August 2023

Prepared By: LHCb Computing Project
C. Bozzi/Editor

Abstract

This document presents a preliminary 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

1.	INTRODUCTION.....	5
2.	COMPUTING MODEL FOR LHCb IN RUN 3	5
2.1.	BASIC FEATURES OF THE LHCb COMPUTING MODEL.....	5
2.2.	MAJOR DRIVERS IN OFFLINE RESOURCE REQUIREMENTS.....	6
2.3.	OFFLINE RESOURCE NEEDS	7
3.	LHC RUNNING SCENARIO AND LHCb DATA TAKING PLANS IN 2025.....	8
4.	PRELIMINARY RESOURCE REQUESTS FOR 2025	9
4.1.	CPU REQUESTS	9
4.2.	DISK REQUESTS	11
4.3.	TAPE REQUESTS.....	12
5.	SUMMARY OF PRELIMINARY 2025 REQUESTS	14
6.	CONCLUSION.....	15
7.	LONG-TERM EVOLUTION OF LHCb COMPUTING RESOURCES (COPIED FROM LAST REPORT)	16
8.	APPENDIX: REPLIES TO THE C-RSG RECOMMENDATIONS	18
9.	APPENDIX: HIGH-LEVEL SUMMARY OF USED AND REQUESTED DISK	21

List of Tables

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.....	7
Table 2-2: Summary of the main assumptions of the LHCb computing model for 2025. ...	8
Table 4-1: Summary of parameters entering the determination of the CPU work needed for simulation.	10
Table 4-2: Estimated CPU work needed for the different activities in 2025 (column “ 2025 This document ”). The 2024 requests, endorsed at the April 2023 RRB (column “ 2024 LHCb-PUB-2023-001 ”), and a comparison with the computing model TDR (row “ LHCb-TDR-018 ”) are also shown.	11
Table 4-3: Disk Storage needed in 2025 for the different categories of LHCb data (column “ 2025 This document ”). The 2024 requests, endorsed at the April 2023 RRB (column “ 2024 LHCb-PUB-2023-001 ”), and a comparison with the computing model TDR (row “ LHCb-TDR-018 ”) are also shown.	12
Table 4-4: Tape Storage needed in 2025 for the different categories of LHCb data (column “ 2025 This document ”). The 2024 requests, endorsed at the April 2023 RRB (column “ 2024 LHCb-PUB-2023-001 ”), and a comparison with the computing model TDR (row “ LHCb-TDR-018 ”) are also shown.	13
Table 5-1: Evolution of offline computing requests in 2024-2025. Units are kHepScore23 for CPU, PB for disk and tape.....	14
Table 6-1: CPU power requested at the different Tier levels in 2025. The 2024 requests, endorsed at the April 2023 RRB, are also shown	15
Table 6-2: LHCb Disk request for each Tier level in 2024 in 2025. The 2024 requests, endorsed at the April 2023 RRB, are also shown. For countries hosting a Tier1, the Tier2 contribution could also be provided at the Tier1.....	15
Table 6-3: LHCb Tape request for each Tier level in 2025. The 2024 requests, endorsed at the April 2023 RRB, are also shown.....	15

1. Introduction

This document presents a preliminary 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 preliminary 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* replaces 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 basic assumptions that enter in the calculation of the offline resource needs are

- 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 (see Figure 2-1).
- A mixture of full/fast/parametric simulations.
- Most of the simulation output is selectively persisted and aggressively filtered.
- The Run3 timeline is now very different from the one that had been assumed in [LHCb-TDR-018]: due to the COVID-19 pandemics, LS2 has ended in 2022. 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. The LHC operations in 2023 and 2024 have been shortened, following the energy crisis. Nominal conditions are expected in 2025. The third LHC long-shutdown (LS3) starts in 2026 and will last three years.

Data taking year	Simulation year										
	X	X+1	X+2	X+3	X+4	X+5	X+6	X+7	X+8	X+9	X+10
X	50%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
X+1		50%	100%	100%	100%	100%	100%	100%	100%	100%	100%
X+2			50%	100%	100%	100%	100%	100%	100%	100%	100%
X+3				50%	100%	100%	100%	100%	100%	100%	100%

Figure 2-1: Model for Monte-Carlo simulation production; the simulation of data taken on year X starts at 50% of the nominal value in that year, ramps to the nominal value in year X+1, stay constant until year X+4, ramps down to 50% in year X+5 and to zero afterwards

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

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

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 was provided by the LPC to the LHC experiments on June 26th, 2023, see Figure 2-2:

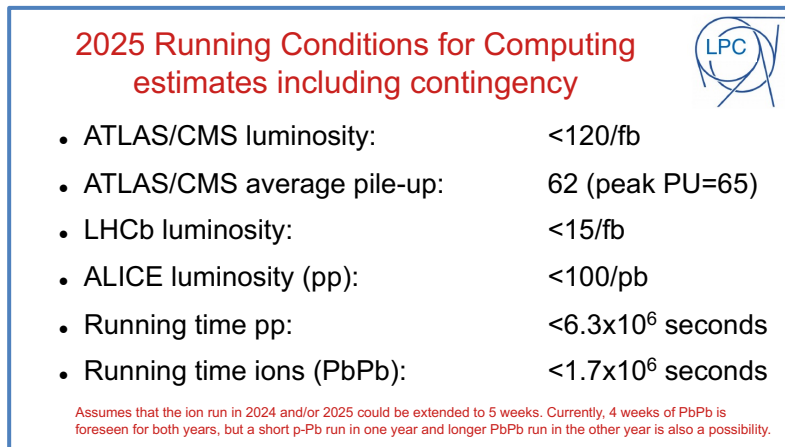


Figure 2-2: Running conditions for the 2025 LHC data taking, as communicated by the LPC to WLCG.

This schedule foresees a LHC running time of $6.3 \cdot 10^6$ seconds for proton collisions and $<1.7 \cdot 10^6$ 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.

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, the same configuration as that used for proton-ion collisions in 2016, which consisted of 5.4 billion triggers in the (FULL+TURBO) streams and 2.2 billion triggers in a NOBIAS stream, is assumed. This might be revised, based on the experience that is going to be gained in the upcoming 2023 heavy-ions collision data taking.

4. Preliminary resource requests for 2025

In this Section, the preliminary 2025 requests are presented. The resource requests for 2024 are also shown for completeness [[LHCb-PUB-2023-001](#)]. The 2024 requests were scrutinized by the C-RSG and endorsed at the April 2023 RRB.

The CPU requests are given by using the new HepScore23 benchmark, which uses HEP-specific workloads and therefore offers a more accurate representation of the real-world performance of computer hardware than the previously used HS06 benchmark. The ratio between the two benchmarks has been set to one for a reference server. The dispersion of the HepScore23 benchmark for different hardware is well within 10%. A normalization factor of one has been used for the requests shown in this document.

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 Run1. This assumption is based on the following considerations:
 - a. The LHCb stripping application (DaVinci) was found to be a factor 5 slower on Run3 simulations of minimum bias events with respect to Run1+Run2 simulations.
 - b. The Run3 sprucing application (Moore) is the same one used in HLT2[‡], it is therefore optimised for performance. When run on the same selection lines, still not the complete set but a representative one, Moore runs a factor 5 faster than DaVinci on Run3 simulations.
2. no provision is made for the offline reconstruction of heavy-ion collision data.
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 (4.8 10⁹ events per fb⁻¹ per calendar year), while that of 2025 collisions will be at 50% of the nominal level.

We take the same event simulation time of Run2 Monte Carlo produced in the *Sim09* simulation cycle. This is motivated by considerations that were already shown earlier, and that we report in the following:

1. When using the more recent *Sim10-v1*[§] simulation cycle, simulation time was improved by a factor two with respect to *Sim09*, for both Run2 and Run3 detector/conditions.
2. The increased instantaneous luminosity (corresponding to $v=7.6^{**}$) results in a 70% increase in the *Sim10-v1* simulation time of the Run3 detector and conditions with respect to *Sim09* Run2 ($v=1.6$). There is a mild dependency on instantaneous luminosity, e.g. *Sim10-v1* Run3 simulation at $v=3.8$ is 15% slower than *Sim09* Run2 simulation.

[‡] when used offline in Sprucing, the Moore application does not re-run reconstruction, but it only processes different (and slower) selection lines than those processed online in HLT2.

[§] “*Sim09*” and “*Sim10*” are major simulation cycles. The latter uses a more recent versions of Geant4, (10.6 instead of 9.6), code optimization, new version of compilers, and optimization of RICH LHCb-custom background processes.

^{**} v is the average number of pp interactions per bunch crossing; this includes elastic and diffractive processes. The average number of visible interactions per bunch crossing μ is lower due to the LHCb acceptance. It scales with luminosity and the inverse of the number of bunches. Roughly speaking, $v=1.6$, 3.8 and 7.6 correspond to luminosities of 4·10³², 1·10³³ and 2·10³³ cm⁻² s⁻¹.

3. A new *Sim10-v2* cycle has just been released; it contains a Geant4 patch^{††} that is about two times faster than *Sim10-v1* for the Run3 detector and $v=1.6$. This scaling factor is not expected to change when using the nominal Run3 luminosity ($v=7.6$)
- 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)	710	93
Total number of events simulated in 2025 (10^9)	99	
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 thorough investigation showed that a significant fraction of time for Run3 simulation was spent in the RICH detectors. This was tracked down to the implementation of surface boundaries for the propagation of optical photons exacerbated by the high number of reflective surfaces in the geometry description of the new RICH photodetectors. By backporting the optimized implementation of the Geant4 code available in a more recent version, 10.7, with respect to what used in *Sim10-v1*, 10.6, the speed of the overall simulation for the Run3 detectors has been halved. The change has been verified to have no effect on physics distributions.

- a. Sprucing work for 23+24+25 data: $(26+68+82) = 176\text{kHepScore23.y}$
 - b. Required work: $176\text{kHepScore23.y} * 3.74$ (Scaling factor analysis/stripping) / 2 (improvement over Run2) = 329kHepScore23.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: $75\text{kHepScore23.y}/2 = 38\text{kHepScore23.y}$
 - d. The sum of Run3+Run2 analysis work gives then $329+38 = 368\text{kHepScore23.y}$
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, together with the requirements for 2024, discussed in the April 2023 RRB [[LHCb-PUB-2023-001](#)].

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 LHCb-PUB- 2023-001	2025 prel. THIS DOCUMENT
First pass sprucing	70	82
End-of-year sprucing	70	82
Simulation	800	1633
Core and distributed computing infrastructure	10	10
User Analysis productions	214	368
Total Work (kHepScore23.years)	1165	2176
LHCb-TDR-018	3470	3276

Table 4-2: Estimated CPU work needed for the different activities in 2025 (column “2025 This document”). The 2024 requests, endorsed at the April 2023 RRB (column “2024 LHCb-PUB-2023-001”), and a comparison with the computing model TDR (row “LHCb-TDR-018”) are also shown.

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 proton-ion and fixed target collisions, and corresponding simulations; this disk provision is made by assuming:
 - a. for proton-ion collisions: 5.4 billion triggers in the FULL and TURBO streams (of events sizes of 180 and 75kB, and average event size of 120kB), and about 2.2 billion triggers in a NOBIAS stream (50kB/event). The events in the RECO stream are subsequently stripped with an average retention rate of 20%
 - b. for fixed target collisions: 1 and 3.3 billion triggers for ion-gas and proton-gas collisions, with corresponding event sizes of 530 and 60kB (same as in the 2024 requests)
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 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 buffer space of 12PB.

Disk storage usage forecast (PB)		2024 LHCb-PUB-2023-001		2025 prel. This document	
Real data	Run1+Run2 pp data	10.2	78.7	10.2	134.4
	Run1+Run2 HI+SLOG				
	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+SLOG2	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	

Table 4-3: Disk Storage needed in 2025 for the different categories of LHCb data (column “2025 This document”). The 2024 requests, endorsed at the April 2023 RRB (column “2024 LHCb-PUB-2023-001”), and a comparison with the computing model TDR (row “LHCb-TDR-018”)^{‡‡} are also shown.

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 80.4PB

^{‡‡} Please note that in LHCb-TDR-018 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.

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 LHCb-PUB-2023-001		2025 prel. This document	
Run1 +	RAW data (pp+HI+fix target)	36.9	79.4	36.9	80.4
	RDST data (pp+HI+fixtarget)	13.8		13.8	
Run2	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	

Table 4-4: Tape Storage needed in 2025 for the different categories of LHCb data (column “2025 This document”). The 2024 requests, endorsed at the April 2023 RRB (column “2024 LHCb-PUB-2023-001”), and a comparison with the computing model TDR (row “LHCb-TDR-018”^{§§}) are also shown.

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

5. Summary of preliminary 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 90% level for CPU, and at the 60% level for storage, evenly distributed over the Tier levels.

LHCb		2024		2025 (prelim.)	
		Request	2024 req./2023 pledge	Request	2025 req. / 2024 CRSG
WLCG CPU	Tier-0	174	81%	340	195%
	Tier-1	572	96%	1114	195%
	Tier-2	319	74%	622	195%
	HLT	50	100%	50	100%
	Sum	1115	89%	2126	191%
Others		50	100%	50	100%
Total		1,165	93%	2,176	187%
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%

Table 5-1: Evolution of offline computing requests in 2024-2025. Units are kHepScore23 for CPU, PB for disk and tape.

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 endorsed at April 2023 RRB	2025 preliminary
Tier 0	174	340
Tier 1	572	1114
Tier 2	319	622
Total WLCG	1065	2076
HLT farm	50	50
Opportunistic	50	50
Total non-WLCG	100	100
Grand total	1165	2176

Table 6-1: CPU power requested at the different Tier levels in 2025. The 2024 requests, endorsed at the April 2023 RRB, are also shown

Disk (PB)	2024 endorsed at April 2023 RRB	2025 preliminary
Tier0	30.6	49.2
Tier1	61.2	98.5
Tier2	11.8	19.0
Total	103.6	166.7

Table 6-2: LHCb Disk request for each Tier level in 2024 in 2025. The 2024 requests, endorsed at the April 2023 RRB, are also shown. For countries hosting a Tier1, the Tier2 contribution could also be provided at the Tier1.

Tape (PB)	2024 endorsed at April 2023 RRB	2025 preliminary
Tier0	117	189
Tier1	133	209
Total	250	398

Table 6-3: LHCb Tape request for each Tier level in 2025. The 2024 requests, endorsed at the April 2023 RRB, 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 2023: 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 exceeding the FC-curves. Firstly, this is exacerbated by the fact that we have reduced the 2024 request. Secondly, we argue this is natural as 2025 will be the first year the LHCb Run 3 detector will run at full data taking capacity, and such steps may be expected for any upgraded detector.
- However, after 2025 our request flattens and comes back within the FC-curves by 2026 or 2027.




Figure 3: 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 shown in Figure 2. Very similar comments pertain as made for the CPU request. There is a larger increase in 2025 balanced by a long flat period which brings us back within the FC-curves by 2026 or 2027.




Figure 4: 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 5




Figure 5: 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 that which LHCb published in its TDR in 2018.

LHCb hopes that this information will allay some of the worries that been expressed by oversight bodies recently 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 and general recommendation.

LHCb-1 *The LHCb Collaboration requests no increase in resources for 2024, compared to the resources already approved by the C-RRB for 2023. This decision is based on the fact that the resources allotted for 2023 will not be fully utilised due to changes in the physics programme following the incident affecting the VELO detector operations. However, less than 90% of the disk resources at the T1 and T2 levels approved by the C-RRB for 2023 have been pledged. The C-RSG encourages the funding agencies to provide the LHCb collaboration with the 2023 approved resources in time for the 2024 data-taking.*

Thanks for this comment, which we completely share.

LHCb-2 *The C-RSG notes that the zero growth of LHCb computing resources in 2024 will be followed by a significant increase (of almost 100%) in 2025, as per the long-term projections shared by the LHCb Collaboration. The experiment considers it acceptable to stagger the growth over 2024 if it makes the procurement process for the sites easier to manage.*

We agree with this comment. We continuously engage with funding agencies on how they are going to procure resources. This report shows that there is an almost 100% increase in CPU requirements, which is certainly a significant step for many funding agencies. The storage increase (60%) is lower, although significant. As stated in the previous section, 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.

LHCb-3 *The C-RSG recommends that the LHCb Collaboration allocates the necessary effort to carry out the activities identified to decrease CPU consumption in simulations (such as optimizing pile-up handling) and reduce storage footprint (such as implementing more aggressive data compression).*

The LHCb collaboration is actively planning to procure structured effort to carry out these activities.

Some work has already started. As an example, it has been recently shown that the introduction of the “FastMath” compilation option gives a 5% speed improvement on the HLT application. When combined with Profile Guided and Link Time Optimisations, a gain of up to 10% can be reached. The corresponding impact on the simulation application Gauss has yet to be studied at the time of this writing.

Investigations are ongoing on how to further reduce the size of objects coming from HLT2 without losing physics information.

Some simulation productions were launched by increasing the ROOT basket size, which in turn may result in more efficient compression of the resulting files. However, this implied a significant increase in memory consumption, which went beyond the maximum allowed limits at several sites. Therefore, we had to roll-back to the default solution.

ALL-1 *The C-RSG requests that the collaborations report in subsequent scrutinies high-level summaries of the manner in which their disk space is utilised, how they have optimised the allocations and how they propose to allocate the requested space. This summary should identify the space used for: i) persistent datasets (and differentiating between primary copies and additional repli- cas), ii) cached storage that is used to hold datasets for short periods of time, and iii) buffering space used, for example, to transfer data from one storage media to another. This information will assist the C-RSG in better understanding the pressures on disk utilization and help identify best practice.*

The utilization of disk space is regularly reported in the resource usage document that LHCb submits every year for the spring scrutiny round. The proposal to allocate the requested space is given in every resource request report (for this document, please see Table 4-3). Regarding the identification of the space used for the various activities:

- i. Persistent datasets: the information is given separately for real data and Monte-Carlo, and Run1+Run2 and Run3 conditions.
- ii. We do not use caches at all in LHCb.
- iii. Buffering space information is reported separately both in the usage and in the request documents.

Appendix 9 gives the requested information in tabular form. A couple of remarks:

- In the absence of sizeable Run3 data, persistent datasets (39.5PB) are dominated by Run1+Run2 data and simulations, which are kept in two copies, instead of the single copy foreseen in the computing model.
- The usage of buffer space is considerably larger than anticipated. This is due to space allocated on CERN EOS for transient commissioning data (5PB) and operational reasons (10.5PB), both of which are expected to be reduced significantly.

***ALL-2** The C-RSG recognises the significant efforts of the collaborations and WLCG to identify additional T1 sites. As noted in earlier scrutiny rounds, the addition of such sites is essential to mitigate the effects on the physics programmes of the loss of any existing T1 facility.*

As already reported earlier, the potential loss of a Tier-1 site is a concern for LHCb.

At the moment, we are using the Russian Tier1 center (RRCKI) as an opportunistic facility. This site provides less than 5% of the total resources required by LHCb, so its impact is moderately small.

Two additional Tier1 sites were identified, in NCBJ Swierk (Poland) and IHEP Beijing (China). The WLCG Overview Board has endorsed in December 2022 their plan to become Tier1 sites for LHCb, setting an 18-months transition period to address all the needed steps. Both NCBJ and IHEP are already Tier2-D centers for LHCb, offering significant compute work for simulation and physics analysis, as well as disk storage. We defined milestones on data management, storage, and network transfers, and the sites are actively collaborating with the LHCb distributed computing team to reach these milestones already by the end of 2023.

Having two additional Tier1 sites mitigates the pressure on storage resources, especially tape. The overall contribution of these two sites is of the order of 10% of the resources required by LHCb.

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

Likelihood:

- 1: never expected to happen
- 2: could happen but very unlikely
- 3: could well happen
- 4: will probably happen

Impact:

- 1: we can deal with it, no problem
- 2: a bit of a hassle but not too bad
- 3: can be managed, but with significant effort
- 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 (21/08/23)	2024 request	2025 request	
Persistent	Run1 + Run2	pp data	19.0	10.2	10.2	
		HI + smog data	3.6			
		TOTAL	22.6			10.2
	Run3	pp FULL	0.02	16.5	30.9	
		pp TURBO		36.3	68.1	
		pp TURCAL		4.5	8.4	
		HI + SMOG data		11.2	16.8	
		TOTAL		0.02	68.5	124.2
	TOTAL DATA			22.6	78.7	134.4
	Run1 + Run2	Monte-Carlo	14.9	8.7	8.7	
	Run3	Monte-Carlo	2.0	3.2	8.0	
	TOTAL Monte-Carlo			16.9	11.9	16.7
	TOTAL DATA + MC			39.5	90.6	151.1
User			2.8	3.0	3.6	
Cache			0	0	0	
Buffer			22.7	10	12	
GRAND TOTAL			65.0	103.6	166.7	