LHCb Computing Resources: preliminary 2026 requests

LHCb Public Note

Issue:	0
Revision:	1
Reference:	LHCb-PUB-2024-005
Created:	1st August 2024
Last modified:	2nd September 2024
Prepared By:	LHCb Computing Project C. Bozzi/Editor

Abstract

This document presents a preliminary estimate of the offline computing resources needed by LHCb in 2026. 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. Two scenarios are presented, corresponding to whether 2026 is a data taking year or not.

Table of Contents

List of Tables

1. Introduction

This document presents a preliminary estimate of the LHCb experiment computing resources requirements for the 2026 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 2026 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 (copied from previous reports)

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. Nominal conditions were reached in 2024, which also included an extension of 5 weeks that were advanced from 2025.

  The third LHC long-shutdown (LS3) starts in 2026 and will last three years. An alternative scenario is being considered, where Run3 is extended in 2026 and LS3 starts in 2027. Therefore, the computing requirements for 2026 are evaluated in the two scenarios.

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

Model assumptions for 2026		
	as LS3 year	As data-taking year
L ($\mathrm{cm}^{-2}\mathrm{s}^{-1}$)	0	2×10^{33}
Pileup	n/a	6
Running time *pp* collisions (s)	0	6.3×10^{6}
Output bandwidth (GB/s)	0	10
Fraction of Turbo events	0	73%
Ratio Turbo/FULL event size	0	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 2026.

3. LHC running scenario and LHCb data taking plans in 2026

In case of data taking in 2026, LHC running times of $6.3\ 10^{6}$ seconds for proton collisions and $<1.7\ 10^{6}$ seconds of heavy-ion collisions are foreseen, with an integrated luminosity for proton collisions at LHCb of less than 15 fb^{-1}.

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 2026.

For heavy-ion collisions, the enhanced configuration (i.e. higher retention rate) with respect to the PbPb 2023 runs, introduced in the previous report, is assumed. 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.

4. Resource requests for 2026

In this Section, the preliminary 2026 requests are presented, according to the two scenarios previously described.

4.1. CPU requests

In case 2026 is not a data taking year:

1. 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 2025 collisions will ramp up to the nominal level ($4.8\ 10^{9}$ events per fb^{-1} per calendar year). 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 2026 will require minimal additional requests.

2. 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
3. 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 2026, this infrastructure requires 10kHepScore23.

In case of data taking in 2026:

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 2025 collisions will ramp up to the nominal level (4.8 10^9 events per fb^{-1} per calendar year), while that of 2026 collisions will be at 50% of the nominal level. 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 2026 will require minimal additional requests.

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

2026	As LS3 year		As data-taking year	
	Run3 pp	Run3 HI	Run3 pp	Run3 HI
CPU work simulations 2026 (kHepScore23.y)	1711	93	2180	93
Total number of events simulated in 2026 (10^9)	110		140	
Fraction full:fast:parametric simulation (%)	36:64:0		36:64:0	
CPU work per event full:fast:parametric simulation (kHepScore23.s)	1.2:0.12:0.02		1.2:0.12:0.02	

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+26 data: (68+82+82) = 233k HepScore23.y
   b. Required work: 233kHepScore23.y * 3.74 (Scaling factor analysis/stripping) / 2 (improvement over Run2) = 435 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 435+38 = 473 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 2026, this infrastructure requires 10kHepScore23.

A summary of the preliminary CPU requirements for 2026 is given in Table 4-2. With respect to the 2025 requests, the most important increase is due to simulation, namely that of Run3 collisions.

In case of data taking, the CPU work that LHCb will get from the HLT farm in 2026 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. Otherwise, it is planned to get about 500kHepScore23.year from the HLT farm.

CPU Work in WLCG year (kHepScore23.years)	**2025** **LHCb-PUB-2024-002**	**2026** **As LS3 (preliminary)**	**2026** **As data-taking (preliminary)**
First pass sprucing	82	0	82
End-of-year sprucing	82	0	82
Simulation	1336	1804	2273
Core and distributed computing infrastructure	10	10	10
User Analysis productions	319	319	473
Total Work (kHepScore23.years)	**1829**	**2133**	**2921**
LHCb-TDR-018	**3276**		

*Table 4-2: Estimated CPU work needed for the different activities in 2026 as LS3 (column "**2026 as LS3**") or data-taking year (column "**2026 as data-taking**"). The other column shows the 2025 requests, endorsed at the April 2024 RRB (column "**2024 LHCb-PUB-2024-002**"); 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 2026. The various terms are due to:

1. Legacy Run1 and Run2 data, and their corresponding MC samples, in a single instance.
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. In case of data-taking, the LHC live time is assumed to be 6.3 million seconds.
   c. we save on disk 2 instances of TURBO stream, 2 instances of the latest (FULL+TURCAL) processing, 1 instance 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. in case of ion-ion collisions in 2026: 6.2 billion triggered events (average size of 142kB/event) in the FULL stream.
   b. In case of fixed-target collisions in 2026: 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.
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 (4.1PB) 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 if not taking data in 2026, or at the end of the 2026 data taking otherwise.

Disk storage usage forecast (PB)		**2025** **LHCb-PUB-2024-002**		**2026** **As LS3 (preliminary)**		**2026** **As data-taking (preliminary)**	
Real data	Run1+Run2 pp data	10.2	121.6	10.2	123.3	10.2	180.5
	Run1+Run2 HI+SMOG						
	Run3: FULL	26.9		26.9		41.3	
	Run3: TURBO	59.3		59.3		91.1	
	Run3: TURCAL	7.3		7.3		11.2	
	Run3: Minimum bias	0.0		0.0		0.0	
	Run3: HI+SMOG2	17.9		19.7		26.8	
Simulated data	Run1+Run2 Sim	8.7	15.0	8.7	19.6	8.7	21.7
	Run3 simulated data	6.3		10.9		13.0	
Other	User data	3.6	25.6	4.2	19.2	4.2	26.2
	Grid Buffers	12.0		5.0		12.0	
	Tier0 Buffer	10.0		10.0		10.0	
Total			**162.2**		**162.1**		**228.5**
	LHCb-TDR-018		**171.0**				

*Table 4-3: Disk Storage needed for the different categories of LHCb data in 2026 as LS3 (column “**2026 as LS3”**) or data-taking year (column “**2026 as data-taking”**). The other column shows the 2025 requests, endorsed at the April 2024 RRB (column “**2024 LHCb-PUB-2024-002”**); the last row (**“LHCb-TDR-018”**) reports a comparison with the computing model TDR.*

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.
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, in case of data taking in 2026, an amount of data of 10GB/s times 6.3 million seconds LHC live time =63PB times 2 instances, for a total of 126PB.

Tape storage usage forecast (PB)		**2025** **LHCb-PUB-2024-002**		**2026** **As LS3 (preliminary)**		**2026** **As data-taking (preliminary)**	
Run1 + Run2	RAW data (pp+HI+fix target)	36.9	81.9	36.9	82.9	36.9	82.9
	RDST data (pp+HI+fixtarget)	13.8		13.8		13.8	
	ARCHIVE	31.2		32.2		32.2	
Run3	pp data (FULL+TURBO+TURCAL)	235.0	283.3	235.0	287.8	361.0	433.7
	minimum bias / no-bias	0.6		0.6		0.6	
	Heavy Ion + fixed target	22.4		24.2		31.3	
	ARCHIVE (data+MC)	25.3		28.0		40.8	
Total			**365.2**		**370.7**		**516.6**
	LHCb-TDR-018		**351.0**				

*Table 4-4: Tape Storage needed for the different categories of LHCb data in 2026 as LS3 (column “**2026 as LS3**”) or data-taking year (column “**2026 as data-taking**”). The other column shows the 2025 requests, endorsed at the April 2024 RRB (column “**2024 LHCb-PUB-2024-002**”); the last row (**“LHCb-TDR-018”**) reports a comparison with the computing model TDR.*

5. Summary of preliminary 2026 requests

Table 5-1 and Table 5-2 show the preliminary CPU (in kHepScore23.y), disk (in PB), and tape (in PB) requests at the various tiers, as well as for the HLT farm and other opportunistic resources, by assuming that 2026 is a LS3 or a data-taking year, respectively. For comparison, the 2025 requests endorsed at the April 2024 RRB are shown, together with the expected increase in 2026. They are at the 60% level for CPU and disk, and at the 50% level for tape.

LHCb		**2025**		**2026 (prelim.)**	
		Request	**2025 req. / 2024 CRSG**	**Request**	**2026 req. / 2025 CRSG**
WLCG CPU	**Tier-0**	283	162%	251	89%
	Tier-1	928	162%	823	89%
	Tier-2	518	162%	459	89%
	HLT	50	100%	500	1000%
	Sum	1779	160%	2033	114%
Others		50	100%	100	200%
Total		**1,829**	**157%**	**2,133**	**117%**
Disk	**Tier-0**	54.9	180%	54.8	100%
	Tier-1	89.9	147%	89.9	100%
	Tier-2	17.4	147%	17.4	100%
	Total	**162.2**	**157%**	**162.1**	**100%**
Tape	**Tier-0**	170.4	146%	172.7	101%
	Tier-1	194.8	146%	197.9	102%
	Total	**365.2**	**146%**	**370.7**	**102%**

Table 5-1: Evolution of offline computing requests in 2025-2026, ***in the assumption that 2026 is an LS3 year****. Units are kHepScore23 for CPU, PB for disk and tape.*

LHCb		**2025**		**2026 (prelim.)**	
		Request	**2025 req. / 2024 CRSG**	**Request**	**2026 req. / 2025 CRSG**
WLCG CPU	**Tier-0**	283	162%	454	160%
	Tier-1	928	162%	1487	160%
	Tier-2	518	162%	830	160%
	HLT	50	100%	50	100%
	Sum	1779	160%	2821	159%
Others		50	100%	100	200%
Total		**1,829**	**157%**	**2,921**	**160%**
Disk	**Tier-0**	54.9	180%	77.3	141%
	Tier-1	89.9	147%	126.7	141%
	Tier-2	17.4	147%	24.5	141%
	Total	**162.2**	**157%**	**228.5**	**141%**
Tape	**Tier-0**	170.4	146%	244.2	143%
	Tier-1	194.8	146%	272.3	140%
	Total	**365.2**	**146%**	**516.6**	**141%**

*Table 5-2: Evolution of offline computing requests in 2025-2026****, in the assumption that 2026 is a data-taking year****. 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 2026, 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 2025 resources endorsed at the April 2024 RRB.

For CPU, we assume that the HLT farm will be partly available during the winter shutdowns and not available if the LHC runs, 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)	**2025**	**2026 as LS3**	**2026 as data-taking**
Tier 0	283	251	454
Tier 1	928	823	1487
Tier 2	518	459	830
Total WLCG	**1729**	**1533**	**2771**
HLT farm	50	500	50
Opportunistic	50	100	100
Total non-WLCG	**100**	**600**	**150**
Grand total	**1829**	**2133**	**2921**

Table 6-1: CPU power requested at the different Tier levels in 2026. The 2025 requests (endorsed at the April 2024 RRB) are also shown.

Disk (PB)	**2025**	**2026 as LS3**	**2026 as data-taking**
Tier0	54.9	54.8	77.3
Tier1	89.9	89.9	126.7
Tier2	17.4	17.4	24.5
Total	**162.2**	**162.1**	**228.5**

Table 6-2: LHCb Disk request for each Tier level 2026. The 2025 requests (endorsed at the April 2024 RRB) are also shown. For countries hosting a Tier1, the Tier2 contribution could also be provided at the Tier1.

Tape (PB)	**2025**	**2026 as LS3**	**2026 as data-taking**
Tier0	170	173	244
Tier1	195	198	272
Total	**365**	**371**	**516**

Table 6-3: LHCb Tape request for each Tier level in 2026. The 2025 requests (endorsed at the April 2024 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 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.

The CPU requirements obtained **by assuming no data taking in 2026** are shown in **Figure 1, top**.

- The increase from 2025 → 2026 does indeed show a local year-on-year increase that does not exceed the FC-curves; the request to WLCG is actually decreasing, as it is assumed that the HLT is fully available in 2026
- After another small increase in 2027, the requests flatten and stay within the FC-curves until the end of Run4.

The CPU requirements obtained **by assuming nominal data taking in 2026** are shown in **Figure 1, bottom**.

- The increase from 2025 → 2026 does indeed show an increase that exceeds the FC-curves. As the HLT farm will be busy with data taking, this increase needs to be sustained by WLCG.
- In 2027 and 2028, the requests flatten but they still exceed the FC-curves; the WLCG request is however mitigated by the availability of the HLT farm in LS3, and stays within the FC-curves from 2028 onwards
- From 2029 onwards, the requests stay within the FC-curves until the end of Run4.

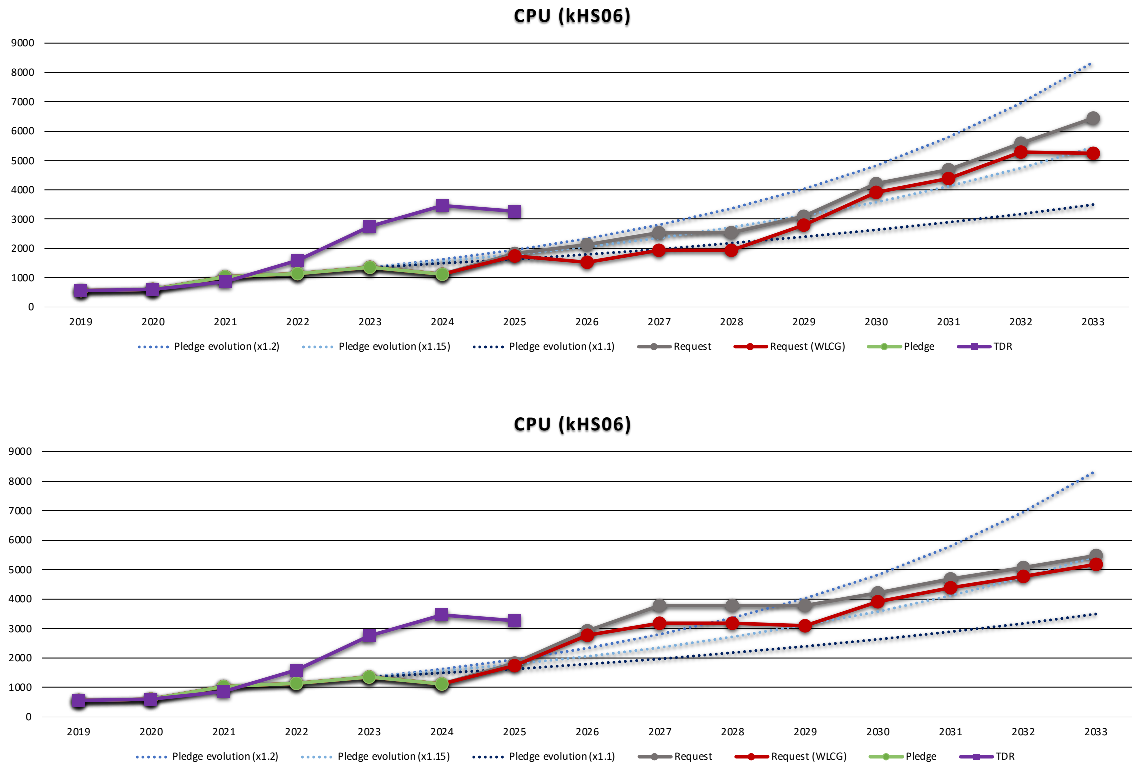

Figure 1: LHCb CPU requirements, ***assuming 2026 is a LS3 (top) or a data-taking (bottom) year****. Blue lines show the FC-bands. Green shows actual pledges, Grey shows total projected requirements. Red shows WLCG requirements allowing for our HLT farm. Purple shows the requests made in the LHCb Upgrade Computing Model TDR.*

The Disk Storage request obtained by **assuming no data taking in 2026** is shown in **Figure 2, top**. The Disk Storage request obtained by **assuming data taking in 2026** is shown in **Figure 2, bottom**. Very similar comments pertain as made for the CPU request. There is a larger increase in 2025-2026, exceeding the FC-curves, balanced by a long flat period which brings us back within the FC-curves by 2028 or 2029.

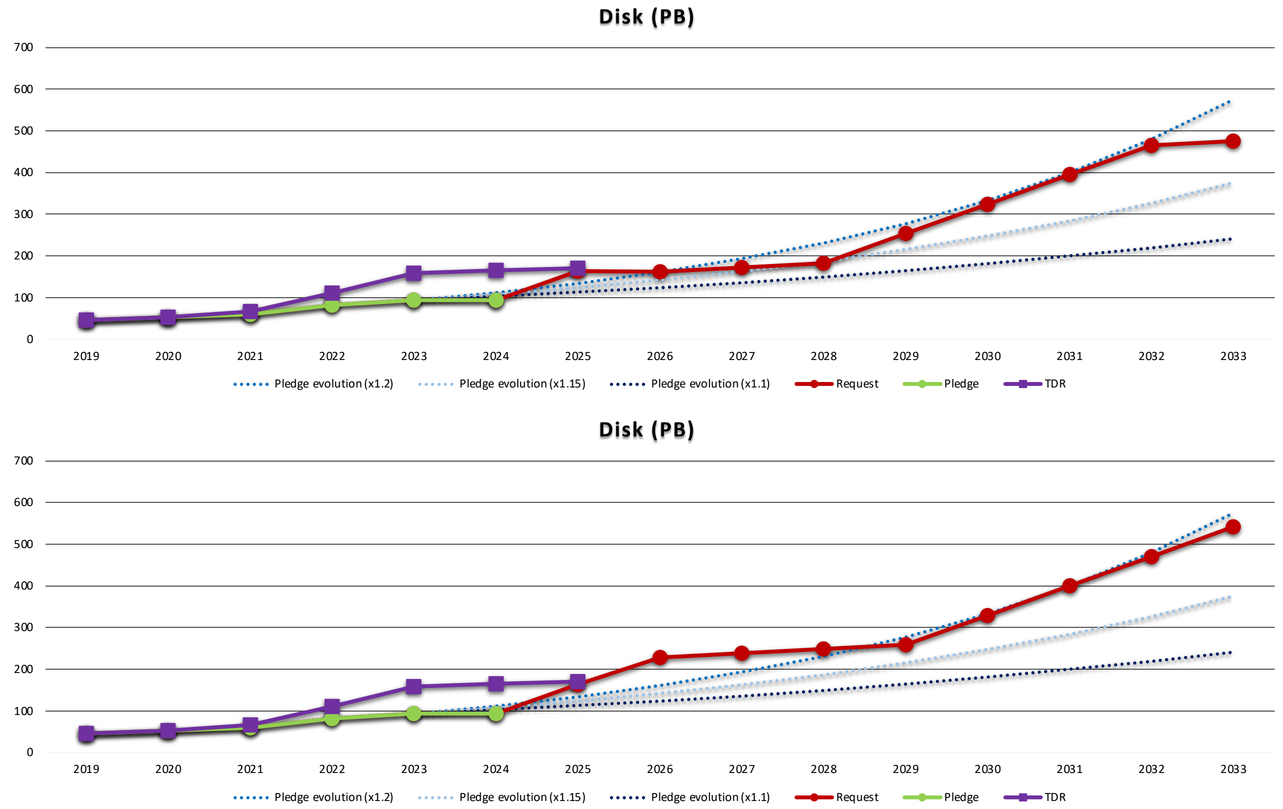

*Figure 2: LHCb Disk Storage requirements, **assuming 2026 is a LS3 (top) or a data-taking (bottom) year**. Blue lines show the FC-bands. Green shows actual pledges, Red shows total projected requirements. Purple shows the requests made in the LHCb Upgrade Computing Model TDR.*

The Tape Storage requirements obtained by **assuming no data taking in 2026** is shown in **Figure 3, top**. The Tape Storage request obtained by **assuming data taking in 2026** is shown in **Figure 3, bottom**. Similar comments as those made for disk apply.

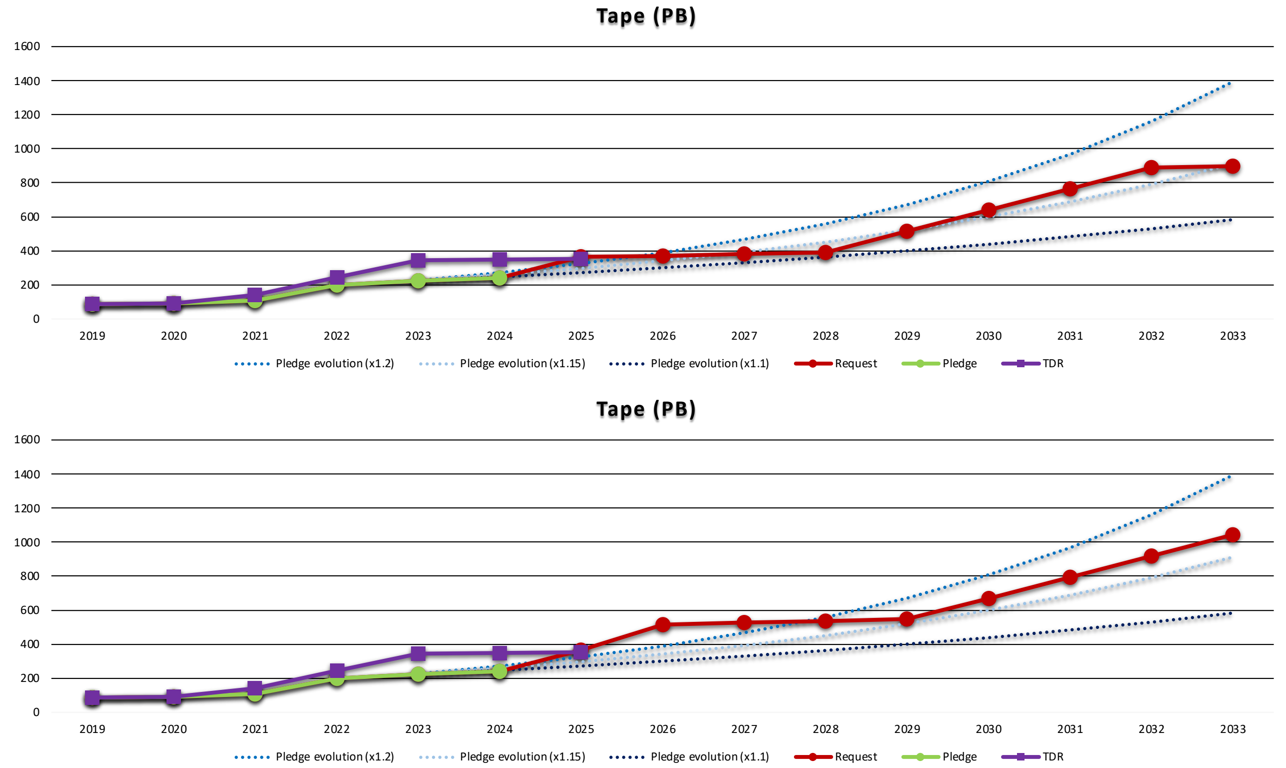

*Figure 3: LHCb Tape Storage requirements, **assuming 2026 is a LS3 (top) or a data-taking (bottom) year**. Blue lines show the FC-bands. Green shows actual pledges, red shows total projected requirements. Purple shows the requests made in the LHCb Upgrade Computing Model TDR.*

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 *Estimates for the required CPU for data processing, simulation and analysis are based on extrapolations rather than on measurements on real data or simulations using the Run 3 nominal running conditions. The C-RSG recommends reassessing those parameters using the data collected in 2024 and the corresponding simulations.*

A reassessment of the parameters driving the CPU requests using the data collected in July and August 2024 is currently in progress. For data processing and analysis, a preliminary estimate shows that the CPU work needed to spruce one event is significantly less than previous extrapolation, whereas the CPU work needed for user analysis and analysis production is significantly more. However, the sum of the two requests is in the same ballpark of those obtained by extrapolating. While this is encouraging news, we prefer for the moment to keep the requests based on extrapolations, as we are still tuning the relevant software applications; in particular, we noticed significant inefficiencies in the physics analysis framework, due to in-time compilation of code snippets which also increases the memory consumption of the applications. We estimate that we can have a much firmer picture by the end of the current data-taking, therefore in time for finalizing the 2026 requests.

Concerning simulation, we have indeed produced Run3 samples, however the bulk of simulation is still on Run1+Run2 conditions (we remind that the bulk of simulation for a given data-taking year starts in the following WLCG year, so the 2024 simulation samples are still not very significant. Moreover, the simulation application has not been frozen yet for production, therefore basing our requests on the currently measured numbers seems not prudent. We therefore prefer to use the extrapolation-based requests, while the simulation application is being finalized and deployed in production.

LHCb-2 *The coordination of the software and computing areas has been reorganized within LHCb to optimize resources, increase efficiency, and improve communication. The coordinator will become a member of the experiment's management board, raising the project's visibility within the Collaboration. The C-RSG requests that the Collaboration provides a report in the next scrutiny round on the effectiveness of this new organization in addressing the shortage of personnel for computing and software support.*

The LHCb Software and Coordination Board started meeting in April 2024, discussing immediate issues such as computing resources for data recorded in 2024, but also longer-term issues related to development practices, future frameworks and associated development. The Software and computing Board coordinator reports to the collaboration board where the shortage of personnel is brought to the attention of the LHCb team leaders so that they can take actions and identify local resources to help. The coordinator reports also during the plenary sessions of LHCb collaboration weeks to inform a large fraction of the collaboration members on the Software and Computing issues.

2024 is a key year for the experiment, with all sub-detectors operational, gathering more than 6 fb-1 so far. The focus has been to ensure stability of the data taking software and therefore to avoid disruptive changes in the current organization of the projects and to support the existing teams by identifying cross project issues and aligning priorities accordingly.

One way to address the shortage of personnel is to increase efficiency by limiting effort duplication between the projects. This has long been a concern for LHCb, which has tried to either use existing common solutions whenever possible or participate in collaborations to develop new ones. The Software and Computing Board naturally helps with communication and strengthens this point.

Attracting new effort on software and computing requires a clear strategy for the future, in line with the interests of the collaborating institutes. This will be discussed during the Computing Workshop organized by the board after this year's data taking.

In conclusion, it is too early to quantify the effectiveness of the new organization yet, as this year has been focused on the success of the data taking and since providing new effort inherently has latency in being put in place. As the year ends, the focus of the Software and Computing board will move to establishing a software and computing strategy where collaborations can be found, in order to alleviate the resource problems. The new organization allowed the software and computing activities to significantly gain in visibility within the collaboration.

Actions have been taken in response to the C-RSG's comments to all LHC experiments:

ALL-1 *The use of compact data formats, first introduced by CMS and more recently adopted by ATLAS, has reduced some pressure on disk resources while increasing throughput of physics analyses to published results. The C-RSG encourages all the Collaborations to continue to focus development efforts in this area both for the analysis of the Run 3 data and for the HL-LHC computing frameworks.*

LHCb analyses are based on particles (with tracks, vertices, particle ID information) and their combinations, and not only on global event properties. This already led the collaboration to find ways to limit resource use in LHC Run 2, with the introduction of the Turbo model that only keeps the signal candidates, subsequently improved with "selective persistency" which allows saving the part of the event relevant to the analysis. Turbo events are much smaller (an order of magnitude) than raw data, but this limits the physics analysis as part of the events is discarded. The possibility to record all reconstructed objects in the event is possible in the "full" stream, for a limited number of analyses. An extra stream recording information needed for offline calibration and performance measurement is also available.

This system is very flexible, the size of the events is very variable, depending only on what is kept by the selective persistency configuration. This can go from a few kB to a full event size. This flexibility requires a mechanism in place to limit resource use and ensure a fair split between the working groups. Each working group submits a number of configurations (trigger lines, or sprucing lines) which are tested on reference input files, to validate their output and check the overall output bandwidth of the system.

In order to further reduce the event size, LHCb started compressing data files with the zstandard algorithm in 2024, providing good compression ratios and fast decompression speed. This has proven more effective than the previously used event level compression.

Furthermore, LHCb continues to monitor development in this area and to improve the software to limit the resource use whenever possible.

ALL-2 *Increasingly the LHC Collaborations are relying on the use of HPC systems. These systems progressively provide a large fraction of their capabilities in form of accelerators, such as GPUs, rather than CPUs. Their effective and efficient utilisation by the LHC community is non-trivial and requires significant changes to the software architecture. Furthermore the connectivity to external networks, the authentication and authorisation services, and batch systems differ. This mandates careful adaptation of workloads and workflow management systems. For this, additional sustained investment in expert developers is essential. The C-RSG recommends identifying adequate mechanisms to fund these activities.*

DIRAC, the workflow management system used by LHCb, has been improved over the years to support the HPC systems made available to the experiment. These developments are shared with all members of the DIRAC collaboration, but it is likely that a sustained effort will be needed, as all HPC systems have their particularities. This however requires a limited amount of person power in complement to the current distributed management system team.

Adapting the workloads to accelerators is a much larger problem as it may require a complete redesign of the applications to be effective. For example, in order to implement its first level trigger on GPUs for LHC Run 3, the experiment had to rethink and implement its filtering code from scratch, in order to take advantage of the parallelism available for the Allen application. This showed the importance of attracting and retaining expert developers. LHCb has had some success collaborating with external computing groups. This approach is crucial to benefit from GPU experts' help, but the funding needed for long term activities remains an issue.

ALL-3 *To reduce confusion in terminology, the C-RSG requests that when experiments refer to a set of data it uses the term “instance” rather than “copy” or “replica”. This would mean that a data set stored twice (e.g., to ensure data integrity) would be described as two instances.*

This document has been adapted accordingly.

LHCb Risk Register

Risk	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	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	Resprucing 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:
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