LHCb Computing Resources: 2023 requests

LHCb Public Note

Issue:	0
Revision:	1
Reference:	LHCb-PUB-2022-010
Created:	1st February 2022
Last modified:	26th April 2022

Prepared By: LHCb Computing Project
C. Bozzi/Editor

Abstract

This document presents the offline computing resources needed by LHCb in the 2023 WLCG year. The computing requests are based on the Computing Model Technical Design Report for the LHCb Upgrade, 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 the computing resources requirements that will be required by LHCb for the 2023 WLCG year.
The 2023 requests have been updated since the preliminary assessment, shown in LHCb-PUB-2021-008. The main motivation for revising them is due to a change in the LHCb data taking plans in 2022, which impacts the 2023 requests in terms of storage and, to a lesser extent, also in compute power needed. Section 2 gives an overview of the expected scenario for data taking in 2022. Section 3 shows the assumptions that have been made regarding the LHC running scenario, and the LHCb plans for data taking in 2023. Section 4 presents the 2023 requests, with a summary given in Section 5. Concluding remarks are given in Section 6, while Section 7 contains replies to the C-RSG recommendations.

2. Revised LHCb data taking plans in 2022

The pandemic situation further delayed an already tight schedule for the complete installation of the LHCb Upgrade before the closure of the experimental cavern. The end of LS2, previously foreseen at the end of January 2022, was reported to end of March 2022.

LHCb foresees to commission the Upgrade detector in parallel with the LHC set up and during the intensity ramp. Data taking will thus be possible when the LHC will give collision with full machine (>1200 bunches) around mid-July. These data will be precious to perform calibration and alignment tasks, and to early "rediscover" signals and check the performance of the new (albeit incomplete) detector in terms of efficiency and resolution for charged particle tracking and identification, and calorimeter reconstruction.

The above plan has an impact on the computing resources needed in 2022 and subsequent years. Storage requirements will be lower. The driving parameters for storage requests is the LHC useful "live time" and the instantaneous luminosity. These have been assumed to correspond to 2.1 million seconds (75 days of data taking with a stable beams efficiency of 33%), and at most an instantaneous levelled luminosity of 1×10^{33} $cm^{-2}s^{-1}$, for an integrated luminosity of $2.1 fb^{-1}$, a factor 5.7 lower than previous assumptions. The resulting storage requirements[1] are roughly a third lower than those previously anticipated (and endorsed in previous scrutiny rounds). Compute resources related to Run3 will also decrease, as they are largely driven by simulation, and the number of events to simulate for a given data-taking year scales with the number of events collected in real data.

3. LHC running scenario and LHCb data taking plans in 2023

The LHC schedule for 2023 was provided by the LPC to WLCG and the LHC experiments.

[1] Storage requirements include also Run1+Run2 data.

This schedule foresees a LHC running time of 6 10^6 seconds for proton collisions and 1.2 10^6 seconds of ion collisions[2] in the 2023 calendar year. The design instantaneous luminosity for LHCb is 2.0 10^{33} cm^{-2} s^{-1}. Therefore, an integrated luminosity for proton collisions of about 12 fb^{-1} is foreseen in 2023.

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 2023.

For ion-ion collisions, it is assumed that the associated requests are the same that will be needed for the corresponding running period in 2022.

The Computing Model for LHCb has been described in detail elsewhere. It has not been changed with respect to the previous report to the C-RSG.

4. Resources requests for 2023

In this Section, the preliminary 2023 requests are presented. The main assumptions regarding the LHCb computing model are briefly reminded below.

4.1. CPU requests

1. For *sprucing*[3] (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
2. no provision is made for the offline reconstruction of proton-ion collision data.
3. simulation consists of three parts, the second of which dominates the other two:
   a. The *simulation of Run1+Run2 pp collision data* has been estimated by assuming that the ratio of events generated with fast/full simulation will be the same as those measured in 2019-2021. The time needed to simulate one event is reduced by 40% with respect to 2019, as measured during the validation campaign of the most recent simulation version. It is expected that the requests for Run1 and Run2 simulation in 2023 will decrease with respect to the current levels. The number of events generated per fb^{-1} in 2023 is assumed to be half of that measured in 2019-2021.
   b. The *simulation of Run3 pp collision data* follows the prescriptions made in the Computing Model TDR. In particular, the simulation of 2023 collisions will be at 50% of the nominal level; it will ramp up to the nominal level (4.8 10^9 events per fb^{-1} per calendar year) in 2024. For this request, we have changed the computing model parameters. We are now assuming a split of 36:64:0 for the events produced with detailed:fast:parametric simulations, instead of the baseline of 40:40:20. This split reproduces the relative amount of events produced with detailed and fast simulations observed in the last three years. We assume that parametric simulation will not be used in 2023, as (i) the associated application has not been validated nor deployed in production yet, and (ii) we advocate that these simulations

[2] In the preliminary 2023 requests, proton-ion instead of ion-ion collisions were envisaged in 2023. It was subsequently clarified that this would not be the case. The requests have thus been updated accordingly. The difference is anyway very small.

[3] *sprucing* replaces the current definition of *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.

will be used in physics analysis only after the understanding of the performance of the LHCb Upgrade, an essentially new detector, will have reached an adequate level.

c. *The simulation of Run3 heavy ion and fixed target collision data* is assumed to take 10% of the work needed for their reconstruction

As requested at the last scrutiny round, a summary of the various parameters entering the CPU request corresponding to simulation is given in the following table:

	Run1+Run2	Run3 pp	Run3 HI
CPU work simulations 2023 (kHS06.y)	283	634	93
Total number of events simulated in 2023 (10^9)	17	38	
Fraction full simulation	0.36	0.36	
Fraction fast simulation	0.64	0.64	
Fraction parametric simulation	0	0.0	
CPU work per event full simulation (kHS06.s)	1.3	1.3	
CPU work per event fast simulation (kHS06.s)	0.1	0.1	
CPU work per event parametric simulation (kHS06.s)	0	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 is 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 that: (i) according to the Computing Model TDR, most of the user analysis will be centrally managed with Working Group productions and therefore with a much lower failure rate, and (ii) the analysis framework is being completely reorganized, with emphasis given on the CPU performance.
5. LHCb is currently using 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 2023, this infrastructure will require 10kHS06.

A summary of CPU requirements for 2023 is given in Table 4-2, together with the 2022 requests presented in April 2021 in LHCb-PUB-2021-002.

With respect to the 2022 requests, the most important increase is due to simulation, namely that of Run3 collisions.

As remarked in previous report, the CPU work that LHCb will get from the HLT farm in 2023 will be much lower than in 2019/2020, as the HLT farm will be used almost entirely for data taking activities during the LHC run, and for possibly reconstructing PbPb collision data during the (E)YETS.

CPU Work in WLCG year (kHS06.years)	2022 LHCB-PUB-2021-002	2023 LHCB-PUB-2021-008	2023 THIS DOCUMENT
First pass sprucing	80	80	80
End-of-year sprucing	80	80	80
Simulation	870	1800	1010
Core and distributed computing infrastructure	10	10	10
User Analysis and working group productions	220	330	240
Total Work (kHS06.years)	**1260**	**2300**	**1410**
LHCb-TDR-018	**1580**	**2750**	

Table 4-2: Estimated CPU work needed for the different activities in 2023, with comparison with the Computing Model TDR (LHCb-TDR-018), the 2022 requests, and the preliminary 2023 requests.

4.2. Disk requests

Table 4-3 presents, for the different data classes, the forecast usage of disk space at the end of 2023. The various terms are due to:

1. legacy Run1 and Run2 data, and their corresponding MC samples; the disk request (10.2 data + 8.7 MC = 18.9PB in total) corresponds to a single copy of the most recent data processing (8.2PB), plus two copies for the three incremental stripping cycles that have been performed in 2021 (2 copies* 3 cycles* 0.33PB/cycle = 2PB), a single copy of simulation (6.6PB) and of the data samples that are going to be generated in 2022 (2.1PB).
2. data from Run3 pp collisions; the disk request is determined according to the Run3 Computing Model TDR.
3. data from heavy-ion and fixed target collisions, and corresponding simulations; this disk provision is made by assuming, as in the 2022 requests, that the number of events to be collected from heavy ion, ion-gas and proton-gas collisions are 1, 1, and 3.3 billion respectively, with event sizes of 980, 530 and 60kB.
4. Run3 simulation of pp collisions, determined by following the Computing Model TDR with the same assumption of point 3.b in Section 4.1 above.
5. User data and buffer data. The former has been estimated by taking the current occupancy of the user data space (1.1PB) and its annual increments (0.1PB/year during Run2), and by assuming that the yearly increment of the space needed for 2023 data will scale by a factor five: 1.1PB + 0.1PB (Run1+2 2021) + 2*0.1PB (Run1+2 2022 and 2023) + 2*0.5PB (Run3 2022 and 2023) = 2.4PB. 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 44PB of data to be recalled from tape, will last two months. A contingency of two weeks would therefore correspond to a buffer space of 11PB, to be compared with a current usage of 6PB over a period of five months.

The increase of the disk request in 2023 with respect to 2022 is driven by the volume of data that are going to be taken. An aggressive policy in keeping legacy Run1+Run2 data on disk has been assumed, by reducing the number of copies and datasets. Being heavily filtered, simulation does not add significantly.

Disk storage usage forecast (PB)		2022 LHCb-PUB-2021-002		2023 LHCb-2021-008		2023 This document	
Real data	Run1+Run2 pp data	10.2	65.9	10.2	114.6	10.2	78.8
	Run1+Run2 PbPb + SMOG						
	Run3: FULL	13.7		27.4		16.2	
	Run3: TURBO	30.3		60.6		35.7	
	Run3: TURCAL	3.7		7.4		4.4	
	Run3: Minimum bias	2.4		0.0		0.0	
	Run3: PbPb + SMOG2	5.6		8.9		11.2	
Simulated data	Run1+Run2 Sim. Data	8.7	11.0	8.7	14.9	8.7	11.5
	Run3 simulated data	2.2		6.2		2.8	
Other	User data	1.8	12.8	2.4	13.4	2.4	13.4
	Buffers	11.0		11.0		11.0	
Total			**89.6**		**142.9**		**102.5**
	LHCb-TDR-018		**111.0**		**159.0**		

Table 4-3: Disk Storage needed for the different categories of LHCb data, with comparison with the Computing Model TDR (LHCb-TDR-018), the 2022 requests, and the preliminary 2023 requests.

4.3. Tape requests

The forecast usage of tape space is the sum of:

1. The tape needed by the Run1+Run2 real (RAW+RDST+ARCHIVE) data at the end of 2021 and the simulated (ARCHIVE) data until the end of 2023, for a total of 78.4PB
2. The tape needed by the Run3 proton collision data (141.4PB FULL/TURBO/TURCAL + 15.6PB ARCHIVE), heavy ion and fixed target data (11.2PB in total), minimum bias / no-bias stream (0.6PB), and Run2+Run3 MC (2PB)

Table 4-4 shows, for the different data classes, the forecast total tape usage at the end of 2023. The increase over 2022 is entirely driven by the Run3 data.

Tape storage usage forecast (PB)		2022 LHCb-PUB-2021-002		2023 LHCb-2021-008		2023 This document	
Run1 + Run2	RAW data (pp+HI+fixed target)	38.4	82.1	38.4	83.1	36.9	78.4
	RDST data (pp+HI+fixed target)	13.7		13.7		13.8	
	ARCHIVE	30.0		31.0		27.7	
Run3	pp data (FULL+TURBO+TURCAL)	120.1	137.8	240.1	276.2	141.4	169.9
	minimum bias / no-bias	0.6		0.6		0.6	
	Heavy Ion Data + fixed target	5.6		10.6		11.2	
	ARCHIVE (data+MC)	11.5		25.4		16.7	
Total			**219.9**		**359.3**		**248.3**
	LHCb-TDR-018		**243.0**		**345.0**		

Table 4-4: Tape Storage needed for the different categories of LHCb data, with comparison with the Computing Model TDR (LHCb-TDR-018), the 2022 requests, and the preliminary 2023 requests.

5. Summary of 2023 requests

Table 5-1 shows the CPU, disk, and tape requests for 2022 and 2023 at the various tiers, as well as for the HLT farm and other opportunistic resources. For 2023, the preliminary requests presented in Autumn 2021 (LHCb-PUB-2021-008) are reported. The final 2023 requests, presented in this document, are in the column labeled "2023, this document". The increase of the 2023 requests with respect to 2022 are also shown.

		FINAL 2022 requests				LHCb-PUB-2021-008		THIS DOCUMENT		
LHCb		2022				2023 (prelim.)		2023		
		Request	Pledge	Pledge/req	2022 req. / 2021 CRSG	Request	2023 req. / 2022 CRSG	Request	2023 req./ 2022 CRSG	2023 req. / 2022 pledge
WLCG CPU	Tier-0	189	189	100%	108%	361	190%	215	114%	114%
	Tier-1	622	515	83%	108%	1185	191%	707	114%	137%
	Tier-2	345	333	96%	107%	657	190%	391	113%	118%
	HLT	50	50	100%	100%	50	100%	50	100%	100%
	Sum	1206	1086	90%	108%	2252	187%	1364	113%	126%
Others		50	50	100%	100%	50	100%	50	100%	n/a
Total		**1,256**	**1,136**	**90%**	**107%**	**2,302**	**183%**	**1,414**	**113%**	**124%**
Disk	Tier-0	26.5	26.5	100%	141%	42.2	159%	30.3	114%	114%
	Tier-1	52.9	47.8	90%	141%	84.4	159%	60.5	114%	127%
	Tier-2	10.2	6.9	68%	141%	16.2	159%	11.6	114%	168%
	Total	**89.6**	**81.2**	**91%**	**141%**	**142.9**	**159%**	**102.5**	**114%**	**126%**
Tape	Tier-0	81	81	101%	184%	131.6	163%	91	113%	112%
	Tier-1	139	116	83%	184%	227.7	163%	157	113%	135%
	Total	**219.9**	**197.3**	**90%**	**184%**	**359.3**	**163%**	**248.3**	**113%**	**126%**

Table 5-1: Evolution of offline computing requests in 2022-2023.

6. Conclusion

This report summarizes the offline computing requests needed by LHCb in 2023, utilising updated information on the LHC running conditions, and on the data taking plans of the upgraded LHCb detector, including heavy ion collisions. 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 2022 resources pledged to LHCb. The resources needed by LHCb in 2022 are lower than those previously requested and pledged; their deployment can therefore be staged.

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 (kHS06)	2022	2023
Tier 0	189	215
Tier 1	622	705
Tier 2	345	395
Total WLCG	**1156**	**1315**
HLT farm	50	50
Opportunistic	50	50
Total non-WLCG	**100**	**100**
Grand total	**1256**	**1415**

Table 6-1: CPU power requested at the different Tier levels in 2022 and 2023

Disk (PB)	2022	2023
Tier0	26.5	30.3
Tier1	52.9	60.5
Tier2	10.2	11.7
Total	**89.6**	**102.5**

Table 6-2: LHCb Disk request for each Tier level in 2022 and 2023. ***For countries hosting a Tier1, the Tier2 contribution could also be provided at the Tier1.***

Tape (PB)	2022	2023
Tier0	81	91
Tier1	139	157
Total	**220**	**248**

Table 6-3: LHCb Tape request for each Tier level in 2022 and 2023.

7. 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 MC simulation for Run 3 data dominates the CPU requirements of the experiment. CPU resource calculations for Run 3 simulation have been performed using preliminary estimates taken from the upgrade LHCb computing proposal released in 2018. The C-RSG recommends refining these estimates using more realistic event simulation times from actual Run 3 conditions. Likewise, the C-RSG encourages LHCb to explore increasing the fraction of events produced with fast simulation.*

We thank the C-RSG for spotting this. Using preliminary estimates to determine the CPU work due to fast MC simulation for Run3 events was indeed an oversight. We are now requiring that the time needed to simulate one event using fast techniques scales with the time needed to perform a detailed simulation in the same manner as measured on Run2 simulations (see Table 4-1 above). The fraction of events produced with fast simulation has been increased to that measured on Run2 simulation productions. We are also assuming that no events will be simulated with parametrized simulation, as (i) the associated application has not been deployed in production yet and (ii) we advocate that these simulations will be used in physics analysis only after the understanding of the performance of the LHCb Upgrade, an essentially new detector, will have reached an adequate level.

LHCb-2 *The CPU cycles obtained from opportunistic HPC resources continue to be marginal. The C-RSG restates its recommendation of devoting appropriate effort to access and exploit HPC facilities. The availability of HPC resources for HEP computing is growing and its contribution to LHCb data analysis would mitigate the pressures the new computing model places on WLCG CPU resources*

We acknowledge that other experiments have made significant use of HPC resources and that these could be helpful to us. We continue to work on configuring the Barcelona Supercomputing Center (BSC) for simulation productions. With the help of colleagues from ATLAS and PIC, we have finally a computing element at PIC that we can use to submit jobs on BSC. Tests are still ongoing at the time of this writing. The opportunistic usage of the SantosDumont supercomputing cluster in Brasil has slowed down, as the center is more and more used by other communities that are granted priority access.

The LHCb effort devoted to accessing and exploiting HPC facilities has not increased in the last six months. However, we would like to stress that human effort is not the main blocking factor in exploiting HPC centers, as there are other limiting factors such as access policies, hardware availability and configurations that must be performed by site managers rather than LHCb people.

LHCb-3 *LHCb faces tremendous computing challenges in Run 3. While recognizing the very significant reductions the collaboration has made to reduce the CPU requirements, the C-RSG is concerned about an apparent shortage of personnel available for computing activities. Some of the C-RSG's past recommendations could not be addressed apparently due to personnel constraints. In particular, software improvements require to further contain the growth of computing resources requirements involve a substantial expert workforce. The C-RSG encourages the experiment and the funding agencies to increase the support for LHCb computing development.*

We believe there has been some misunderstanding which led to this remark. In the specific case of the past recommendations, we believe you are referring to finding an automatic way to determine popular and less popular data so that disk space could be recovered. As a matter of fact, a popularity agent exists in production since several years. It determines the usage of all our datasets on a weekly basis, it runs completely automatically, and it is used to prepare the popularity reports in every scrutiny round. It gives a list of datasets that have not been used, with an adjustable date threshold. This list of datasets can then be used to delete datasets in a single DIRAC command, that is manually executed, and which takes a few seconds of the LHCb data manager's time. If an analyst needs a deleted dataset to be recovered, he/she opens a JIRA task with the relevant bookkeeping path. The data manager then issues a DIRAC command that stages this dataset from tape. Therefore, a quasi-automatic mechanism to deal with unused datasets has been already in place for some time. When we wrote in previous reports that no effort was available to automatize the handling of unused datasets, we had something much more sophisticated in mind, involving e.g. automatic data cleaning, copying, requesting, staging, expiration dates, etc. We understand these mechanisms have been implemented in other VOs. They have led to a more intense use of the tape systems and increased usage of network. Rather than implementing a full infrastructure to deal with this, we believe that a simpler solution based on the quasi-automatic mechanism described above would serve the purpose without impacting on computing operations significantly. Certainly, one needs to fine-tune the deletion cycles towards optimizing the amount of disk space that can be recovered and the users requests to stage datasets back on disk. We believe that cleaning up datasets every few months would serve the purpose.

That said, it is nevertheless true that we have a shortage of personnel available for computing activities. We try address this in the following ways:

- We have introduced a computing shifter role. whereby an LHCb collaborator, even without specific computing proficiency, monitors the distributed computing infrastructure in terms of workload management, data management, site availability and reliability. Computing experts are therefore relieved from performing such tasks and can commit more effort to tasks requiring their expert knowledge. The shifter role will eventually be expanded to also cover data quality and simulation aspects.

- We have split the production management, previously a one-person role, in three separate roles: simulation, analysis and real data productions. We identified a new production manager for real data. We are setting up a pool of people that will manage the analysis productions.
- We have prepared a list of tasks and roles, including the required efforts in terms of FTEs and time span, in a "collaboration agreement" to be submitted to the LHCb collaborating institutes so that they can record their efforts and better communicate their commitments in grant applications. This should hopefully help in identifying more contribution to the computing activities.

ALL-1 *The C-RSG thanks all four experiments for the responses to the Spring 2021 recommendations. It appreciates the constructive discussions the group had with the computing coordinators of each experiment.*

Thanks for this remark and thank you for the constructive discussions**.**

ALL-2 *The C-RSG notes that computing resource allocations are usually made at the beginning of the year, providing a well-defined and committed resource for the subsequent 12 months. The group is concerned that reductions in the WLCG MOU commitment mid-year compromise the experiments' ability to complete their physics programs.*

We are very sympathetic with this remark.

ALL-3 *The use of GPUs as a source of T1 and T2 computing is increasing, but the benchmarks necessary to account for the GPU computing resources have not been established. The C-RSG recommends that an interim approach to GPU accounting be established to allow a consistent assessment of the total resources available to each collaboration.*

We believe that, rather than a remark to experiments, this is a suggestion towards WLCG to work in that direction. We notice that a WLCG benchmarking working group has been established for quite some time, that this group is working towards the redefinition of a new CPU benchmark, based on the VO workloads. We understand this group has also GPU benchmarking in their agenda, although with a lower priority. LHCb will cooperate in providing any relevant workloads for that. At the moment, the HLT1 application (i.e. the first stage of the software trigger) is capable of running on GPUs.

ALL-4 *Given that the C-RSG sees a future where processor architectures may evolve, the group suggests the experiments continue to assess the long-term outlook for architectures alternative to the current X86 platform (e.g., ARM, Power, RISC-V) and anticipate porting implications.*

Assessing alternative architectures is indeed an important topic. This is one of the tasks that we explicitly mentioned in the collaboration agreement discussed above.

ALL-5 *The C-RSG notes the increasing importance of networking in view of the changing computing models for the LHC experiments.. The group requests that experiments report on any risks to their computing plans from possible network constraints.*

The sheer data volume originating from the LHCb Upgrade detector indeed dictates for an increase of throughput in geographical network (mainly LHCOPN, to ship data from the CERN Tier0 to the Tier1 sites, but also LHCONE and general-purpose networks to transfer data to Tier2-D sites and simulation from Tier2 and other sites to storage elements on T1 and T2D sites). Also, the throughput to and from tape system is increasing. A network challenge and a tape challenge have been performed last autumn, where all LHCb VOs were involved, and where it was shown that the network is adequate to support the requirements. Another tape test will be performed in March, where LHCb will also assess its capability to read back from tape, in conjunction with the other LHC VOs, at the aggregated speed required for the end-of-year resprucing campaign.

The impact from possible network constraints to the LHCb computing activities is high: in absence of adequate networking from CERN to Tier1 sites, we would not be able to transfer and safely stow custodial data and process them at the required rates. Constraints from LHCONE and general-purpose network would have a moderate impact on the availability of simulated datasets. LHCb does not make use of disk caches and dynamical replicas; therefore, the impact of network constraints on physics analysis is small. We consider the likelihood of network shortages to be low. We believe that the overall risk to the LHCb computing plans from possible network constraints is small.