

LHCb Computing Resources: 2011 re- assessment, 2012 request and 2013 forecast

LHCb Public Note

Issue: V3
Revision: 4

Reference: LHCb-PUB-2011-009
Created: 1st Mar 2011
Last modified: 4nd Mar 2011

Prepared By: LHCb Computing Project
R. Graciani/Editor

Abstract

This note covers the following aspects: re-assessment of computing resource usage estimates for 2011 data taking period, request of computing resource needs for 2012 data taking period and a first forecast of the 2013 needs, when no data taking is foreseen.

Estimates are based on 2010 experience and latest updates to the LHC schedule, as well as on a new implementation of the computing model simulation tool. Differences in the model and deviations in the estimates from previous presented results are stressed.

Document Status Sheet

Table 1-1 Document Status Sheet

1. Document Title: [Project Name Qualification] User Requirements Document			
2. Document Reference Number: [Document Reference Number]			
3. Issue	4. Revision	5. Date	6. Reason for change
Draft	1	1 Mar 2011	First version
V1	2	2 nd March 2011	Full version
V2	3	4 th March 2011	Final version for review
V3	4	4 th March 2011	Final version

1.	INTRODUCTION.....	1
2.	SUMMARY OF 2010 RESOURCE USAGE.....	2
3.	CHANGES IN THE LHCb COMPUTING MODEL	5
4.	RE-ASSESSMENT OF 2011 NEEDS	7
5.	ESTIMATES OF 2012 NEEDS	9
6.	FORECAST OF 2013 NEEDS	10
7.	SUMMARY	11

List of Tables

Table 1-1 Document Status Sheet.....	ii
Table 2-1: RAW data collected by LHCb during 2010 and its distribution between the different SEs. Only physics data (155 TB in total) is replicated to Tier1s.....	2
Table 2-2: Snapshot of the Disk Storage usage at the different LHCb Tier0/1s. The view of the site, provided by SLS, and the view of LHCb taken from the LFC are shown and compared to the pledge.....	2
Table 2-3: Snapshot of Tape Storage usage at the different LHCb Tier0/1s. SLS only provides usage of the disk cache so it is not included. The LHCb view is taken from the LFC and compared to the pledges.	3
Table 2-4: Average CPU power used and efficiency as reported by WLCG Accounting portal.....	3
Table 2-5: Raw CPU work (in days), Normalized CPU work (in kHS06·day), and average CPU power in HS06 for all LHCb Jobs. Normalization is calculated based on the known requirements of a Reference Monte Carlo simulation producing 24 million events requiring 660 kHS06·day.	4
Table 3-1: New and old per event CPU and Storage needs for the different LHCb applications and data types relevant for this document.	5
Table 7-1: LHCb 2011-13 CPU work estimated need grouped by Tier level and by Activity.	11
Table 7-2: LHCb 2011-13 CPU peak power estimated need grouped by Tier level and by Activity.	11
Table 7-3: LHCb 2011-2013 Disk Storage estimated need grouped by Tier level and type of Data.	13
Table 7-4: LHCb 2011-2013 Tape Storage estimated need grouped by Tier level and type of Data.	13
Table 7-5: Re-assessment of 2011 needs without an extra 1 kHz of trigger rate for charm physics.....	14

1. Introduction

This document summarizes the last updates of LHCb computing resource usage estimates. It covers the period 2011-2013¹ and is based on the experience from 2010 and the latest updates to the LHC schedule. For the 2011, it must be considered as an update, in the light of 2010 experience, of the request already presented and approved last year. For 2012, it must be considered the LHCb request to LHCC based on our current best knowledge of the computing and physics plans of the collaboration. For 2013, when no data taking is foreseen, it must be considered as a first attempt that will likely be re-evaluated in the light of 2011 experience and any possible updates in the machine schedule.

The LHCb data taking conditions during 2010, and those expected for 2011, deviate notably from those used in the past to produce the LHCb computing model and previous resource usage estimates. Due to these changes and based on the 2010 experience several changes have been introduced in the LHCb computing model to best adapt to the new conditions. These changes have been studied with one main target in mind: minimize the impact on the expected physics performance of the detector. This is the leading argument in this document.

To better incorporate the changes a new simulation tool has been implemented. Using the old model, the results from the new tools have been validated against the old one. All estimates presented in this document are based on the new tool².

This document is organized as follows: the usage of resources in 2010 is presented in Section 2, the changes introduced in the computing model are given in Section 3, estimates for 2011-13 are described in Sections 4, 5 and 6, and finally a summary is presented in Section 7.

¹ For the purpose of this document a given year always refers to the period between April 1st of that year and March 31st of the year after.

² The new simulation tool is a set of python modules steer by some configuration files that describe the different activities included in the simulation. The code and configuration files used can be obtained at <https://coma.ecm.ub.es/svn/lhcbcompsim/tags/CRSG-2011-03-04>, the simulation can be executed by issuing the command:

```
# python runsim.py -o results/CRSG LHCb_Sites.cfg LHCb_Reconstruction.cfg  
LHCb_Reconstruction_2011_Clean.cfg LHCb_Reconstruction_2010_Clean.cfg  
LHCb_Reconstruction_2012_Clean.cfg LHCb_Simulation.cfg LHCb_Users.cfg  
LHCb_Charm_2011.cfg
```

2. Summary of 2010 resource usage

A full review of the usage of computing resources by LHCb during 2010 is presented in the note “LHCb Computing Resource usage in 2010”, LHCb-PUB-2011-008. This section presents a summary of those aspects more relevant for the purpose of this document.

We have reported on the usage of computing resources by LHCb during the period from 1st April 2010 to 31st January 2011. The larger event size (and trigger rate at design value throughout the year, despite an order of magnitude less luminosity) has led us to make compromises on disk space in order to fit into available resources. CPU peak power has been fully utilized to achieve reprocessing in a reasonable time. We are using more than the pledged peak power for the current MC10 simulation campaign, this will have to be smoothed in future. Integrated CPU work has been adequate.

The following tables from LHCb-PUB-2011-008, summarize its content.

LHCb 2010 RAW data		
SE	Size (TB)	# of Files
CERN-RAW (T1D0)	97.4	87233
CERN-RDST (T1D1)	83.4	76711
CERN	180.8	163944
CNAF-RAW (T1D0)	19.5	17847
GRIDKA-RAW (T1D0)	27.5	25141
IN2P3-RAW (T1D0)	34.4	31666
NIKHEF-RAW (T1D0)	33.8	31024
PIC-RAW (T1D0)	9.1	8520
RAL-RAW (T1D0)	30.6	27978
Tier1s	154.9	142176

Table 2-1: RAW data collected by LHCb during 2010 and its distribution between the different SEs. Only physics data (155 TB in total) is replicated to Tier1s.

Disk Summary (20/12/2010)	Pledge (TB)	Seen by SLS			Seen by LHCb	
		TB			TB	
		Total	Used	Avail.	Used	Pledge-Used
FZK	495	500	331	169	339.9	155.1
IN2P3	610	641	334	304	320.7	289.3
CNAF	450	463	392	71	391.6	58.4
NL-T1	560	563	339	224	254.5	305.5
PIC	240	255	138	117	138.3	101.7
RAL	505	791	562	229	453.3	51.7
Tier1s	2860	3213	2096	1114	1897.5	962.5
CERN	1135	1175	922	253	763.6	371.4

Table 2-2: Snapshot of the Disk Storage usage at the different LHCb Tier0/1s. The view of the site, provided by SLS, and the view of LHCb taken from the LFC are shown and compared to the pledge.

Tape Summary (20/12/2010)	Pledge (TB)	Seen by SLS			Seen by LHCb	
		TB			TB	
		Total	Used	Avail.	Used	Pledge-Used
FZK	350				160.7	189.3
IN2P3	555				188.4	366.6
CNAF	265				126.4	138.6
NL-T1	420				161.1	258.9
PIC	130				65	65
RAL	380				201.6	178.4
Tier1s	2100				903.2	1196.8
CERN	1635				844.7	790.3

Table 2-3: Snapshot of Tape Storage usage at the different LHCb Tier0/1s. SLS only provides usage of the disk cache so it is not included. The LHCb view is taken from the LFC and compared to the pledges.

HS06	Norm CPU	Fraction	Norm Elapse	CPU Eff
CERN-Tier0	7166	14.5%	8401	85%
IT-Tier1	3817	21.0%	4598	83%
DE-Tier1	4039	22.3%	4725	85%
FR-Tier1	2437	13.4%	2806	87%
NL-Tier1	3188	17.6%	3775	84%
ES-Tier1	1412	7.8%	1652	86%
UK-Tier1	3255	17.9%	3853	84%
Tier1s	18148	36.6%	21409	85%
FR-Tier2	2553	10.5%	2680	95%
DE-Tier2	2114	8.7%	2142	99%
IT-Tier2	5739	23.6%	6685	86%
RO-Tier2	238	1.0%	250	95%
RU-Tier2	2610	10.8%	2841	92%
ES-Tier2	830	3.4%	558	149% ³
CH-Tier2	771	3.2%	863	89%
UK-Tier2	9420	38.8%	10450	90%
Tier2s	24273	48.9%	26469	92%
All	49587	100.0%	56279	88%

Table 2-4: Average CPU power used and efficiency⁴ as reported by WLCG Accounting portal.

³ At least one of the sites publishes “unnormalized” elapse time.

⁴ When fractions are given for individual Tier1s or Tier2s, they are with respect to the sum of all Sites of the same Tier level and not the total; i.e., the fractions for all Tier1s add up to 100% and the same for all Tier2s.

	Ref.Jobs	Frac.	Ref.CPU	Norm	All CPU	All CPU	Frac.	Av.CPU Power
		%	days	HS06	days	kHS06*day	%	HS06
Tier0	6303	2.6	2265	7.7	265480	2032	16.9	6640
DE-T1	29562	31.8	10249	7.9	115914	920	23.6	3005
ES-T1	8149	8.8	2654	8.4	38782	328	8.4	1070
FR-T1	2610	2.8	1032	7.0	80709	561	14.4	1834
IT-T1	31668	34.1	11574	7.5	99219	747	19.1	2440
NL-T1	5045	5.4	1756	7.9	76726	606	15.5	1981
UK-T1	15886	17.1	5555	7.9	89448	704	18.0	2299
Tier1s	92920	38.6	32820	7.8	500798	3900	32.4	12745
CH-T2	2598	2.3	1062	6.7	20048	135	2.9	441
DE-T2	2297	2.1	7834	8.1	20860	168	3.6	549
ES-T2	6417	5.7	2575	6.9	34903	239	5.1	782
FR-T2	6434	5.8	2333	7.6	87298	662	14.0	2164
IT-T2	27048	24.2	10903	6.8	126560	864	18.3	2822
PL-T2	1097	1.0	404	7.5	19646	147	3.1	479
RO-T2	9	0.0	7	3.6	8570	31	0.7	101
RU-T2	8753	7.8	3019	8.0	63523	507	10.7	1656
UK-T2	57215	51.1	21901	7.2	277408	1993	42.3	6514
Tier2s	111868	46.4	42988	7.2	658815	4716	39.2	15411
Others	29899	12.4	11123	7.4	195724	1447	12.0	4729
Total	240990	100.0	89196	7.4	1620817	12045	100.0	39364⁵

Table 2-5: Raw CPU work (in days), Normalized CPU work (in kHS06-day), and average CPU power in HS06 for all LHCb Jobs. Normalization is calculated based on the known requirements of a Reference Monte Carlo simulation producing 24 million events requiring 660 kHS06-day.

⁵ The difference between this number, 39.5 kHS06, and the WLCG estimation, 49.5 kHS06, is due to the different normalization procedures. LHCb is defining a single procedure and applying it for all sites, while WLCG allows some freedom for each site to define its own normalization. More details can be found in LHCb-PUB-2011-008.

3. Changes in the LHCb computing model

The main changes introduced in the LHCb computing model are due to the increase in the average number of visible proton-proton collisions per bunch crossing, ν . The nominal ν for LHCb was 0.4 that maximizes the number of events with a single visible interaction. In order to increase the delivered luminosity LHC proposed to increase this nominal value during 2010. In an initial period ν went up to 1, then to 1.5 and for most of the luminosity delivered LHC fills started with value of 2.5 with exponential decay to 1.5 and average value of 2.0. LHCb has proved with 2010 data that it is able to fulfill its physics cases under these new conditions, although at the cost of additional CPU requirements to handle the extra complexity of the events and extra storage requirements to cope with increased event sizes. For 2011 and 2012 we expect that LHC will be able to provide data at the LHCb interaction point in similar conditions. Therefore the new ν average of 2.0 has been taken to determine the per event CPU and storage needs for the different applications and data types. The updated values, compared to the old nominal ones are presented in Table 3-1.

Process	CPU (HS06-s/evt)		Data Type	Storage (kB/evt)	
	New	Old		New	Old
Data Taking			RAW	50	30
Reconstruction	25	12	SDST	40	25
Stripping	1.75	0.8	DST	130	80
			MDST	13	
Simulation	1700	376	DST	400	300

Table 3-1: New and old per event CPU and Storage needs for the different LHCb applications and data types relevant for this document.

These changes in the basic parameters have several immediate consequences:

- There is an important increase in the CPU needs to process the new “nominal” event, that will produce a larger output.
- Simulation needs, for the same amount of produced events, will require a large increase in the amount of CPU resources.

After the experience from 2010, these variations have driven to the following changes in the computing model:

- The number of replicas kept for the last and previous passes of the reconstruction (DST/MDST) of a given data sample is reduced from 7 (one at each Tier0/1) to 4. These 4 replicas are distributed in the following way: **2 “master” replicas on T1D1** (one at CERN and the other distributed among the Tier1s) and the other **2 “extra” replicas are on T0D1** (distributed among the Tier1s).
- For the simulation, following the same trend as in the past, most of the resources will be dedicated to **simulation of signal events and dedicated backgrounds** (b/c samples) while combinatorial backgrounds will largely be estimated from the data itself.

Other global changes derived from 2010 experience are the following:

- **CPU efficiency for Monte Carlo simulation**, including human errors and intrinsic efficiency of the system is increased from 70 to **80%**.
- **2 “archival” replicas on TXD0** (one on T2D0 at CERN and one on T1D0 distributed among the Tier1s) have been added to the model for DST/MDST (for real and simulated data). Experience has shown that it is not possible in practice to migrate data from T1D1 to T1D0 as final step in the life time

of a given sample as foreseen, therefore “archival” replicas are created at the same time the new data becomes available.

A summary of other considerations taken into account while producing the new estimates:

- Users activity has shown during 2010 large short term fluctuations that are properly handled by the DIRAC priority mechanism even under large loads due to simulation or data processing activities. Therefore for the model only average usages are defined. They include a flat background activity plus a higher activity during the “summer” period from June to October (going from the first partial reprocessing of the current year’s data in June, to the first partial re-stripping of the current year’s data in September, see Section 4). This activity has an associated data produced for which 2 T0D1 replicas are foreseen. As starting point, 250 TB of Disk storage (with 2 replicas and the corresponding Storage efficiency to be folded in) has been allocated for all users. The level of the user activity and data sizes are taken from 2010 experience with some reasonable increase due to the additional available luminosity.
- Concerning older versions of reconstructed data, the 2 most recent versions are kept as described above. The number of “master” and “extra” replicas is reduced to half for the next older version (if existing), and replicas for even older versions (if existing) are completely removed. While at the moment this is implemented exactly in this manner, it is foreseen that during 2011 the tools are developed to use the space for “extra” replicas in a more dynamic manner based on some popularity measurement, allowing more than 2 replicas for “popular” data samples and reducing the below 2 replicas less used data.
- Low luminosity data taken during 2010 has shown the big potential of LHCb for charm physics. This has driven the experiment to study the possibility to increase the nominal trigger rate from 2 to 3 kHz, dedicating the extra bandwidth for this purpose. For 2011, this extra contribution is included in Section 4, however the consequences of running in the old 2kHz mode have been analyzed. For 2012, the base line presented in Section 5 already includes this extra data. It is beyond the scope of this document to justify the physics case for this charm sample.
- The number of re-processing and re-stripping passes for each data taking year will be detailed in the corresponding section.
- Considering the expected uniformity in the data from 2011 and 2012, Monte Carlo samples are assumed to be mostly compatible for the period described in this document. Therefore Monte Carlo produced during 2010 (approximately 50% of the amount we expect for 2011 or 2012) will be kept until 2012. When the 2012 sample is fully available 2010 replicas will be reduce to half and fully removed when 2013 samples are available. The same applies with a one-year delay for 2011 samples.

For completeness, the settings that have not been changed from previous estimates are:

- **CPU efficiency of 70 %** for centrally driven **data processing** activities, reconstruction, re-processing, stripping and re-stripping.
- **CPU efficiency of 60%** for user activity.
- **Disk Storage efficiency of 70%.**
- **Tape Storage efficiency of 100%.**

4. Re-assessment of 2011 needs

Apart from parameters described in Section 3, the following settings have been introduced to describe the 2011 data taking conditions. Firstly the LHC schedule has been divided in the following periods:

- **Startup:** 5 weeks from March 14th with a reduced duty cycle.
- **Ramp up:** 6 weeks from April 18th with nominal duty cycle and increasing luminosity.
- **Nominal:** 24 weeks from May 30th with nominal duty cycle and nominal luminosity.

This amounts to a total of **5·10⁶ seconds** of LHC collisions that at a nominal trigger rate of 2 kHz will produce a total of **10·10⁹ events**, and a size of **500 TB⁶ of RAW data**. After reconstruction it will produce **400 TB of SDST** (reduced DST format) and after stripping with an average retention of 10% it corresponds to **130 TB of DST** per full pass.

The new **1 kHz of charm data corresponds to 5·10⁹ events, 250 TB of RAW data and 65 TB of MDST data** for each reconstruction.

The full Tier2 pledged CPU power dedicated to simulation would allow to **produce 750 M Monte Carlo events and with size of 300 TB**.

The data from the detector is fully reconstructed and stripped quasi-online, following distribution of the RAW data to the Tier1s. The following re-processing and re-stripping passes are foreseen:

- Partial re-processing and re-stripping at the beginning of June of the data taken in the “Startup” and “Ramp up” periods, once these data have been used to re-optimize the reconstruction code for the 2011 detector and accelerator conditions.
- Partial re-stripping at the beginning of September of all data taken so far.
- Full re-reprocessing and re-stripping of 2011 sample at the end of the data-taking period. In order to be in time for winter conferences this pass must be completed within 2 months.
- Before the start of the 2012 data-taking period a full re-stripping of the full data sample is foreseen. This pass will be the “final” one for physics analysis based on 2011 data.

For the data taken in 2010 the following passes are foreseen during 2011:

- Full re-processing and re-stripping shortly before the end of 2011 data-taking, to bring to the latest version of the reconstruction the data from 2010.

With all these ingredients, as well as the user and simulation activities described in Section 3 and the efficiency numbers quoted there the total CPU work consumed during 2011, the peak CPU power necessary during the full re-processing, and the Disk and Tape Storage usages are estimated. They are presented in a tabular form in Table 7-1, Table 7-2, Table 7-3 and Table 7-4, respectively.

The corresponding CPU pledges from the Tier0, the Tier1s and the Tier2s for the 2011 are 21 kHS06, 70 kHS06 and 48 kSH06 respectively. **2011 CPU pledges should be able to cover the 2011 estimated needs with a 20% head room at the Tier0/Tier1s that will need to provide a 30% extra peak during the full reprocessing at the end of the data-taking period in order to finalize it within 2 months**, as shown in Figure 7-1 and Figure 7-2. For Tier2s the full pledged power is dedicated to Monte Carlo simulation.

For Disk the re-assessed estimates fall 1.9 PB short of the 2011 pledges. This is due mostly to the extra charm data but also in part for the large event sizes delivered by LHC, even if the total number of replicas has been reduced to compensate this effect. The Tape request for CERN represents an important increase

⁶ In this document TB refers to 10¹² bytes. In contrast to TiB = (2¹⁰)⁴ bytes = 1.1·10¹² bytes. In the same way PB refers to 10¹⁵ bytes.

due to the move from T1D0 to T2D0 for “archival” replicas. With the new charm data Tape need at Tier1s slightly goes over the current pledge.

For a comparison of these re-assessed estimates with the estimated needs without the extra charm data, see Section 7.

5. Estimates of 2012 needs

For 2012 the changes introduced with respect to 2011 are:

- 1 kHz of trigger rate dedicated to charm physics. This extra data increases by **250 TB the size of the RAW** sample and by **65 TB of MDST** for each reconstruction pass. This request for 2012 is independent of whether the same extra change is approved for 2011.
- The foreseen re-processing and re-stripping schedule for 2012 data is the same as for 2011 data the previous year.
- 2010 and 2011 data will be fully re-processed and re-stripped, together with 2012, data at the end of the data-taking period.
- A moderate increase of 30% in user activity is assumed, with the corresponding increase in the amount of data.

As for 2011, these parameters are put together and the resulting estimates are included in Table 7-1, Table 7-2, Table 7-3 and Table 7-4 for the yearly CPU work, the peak CPU power, the Disk and Tape usage.

The yearly CPU work represents only a very moderate increase with respect to 2011 pledges, approximately 15%. As in 2011, a peak of computing power is expected towards the end of the data-taking period, when the full 2010-12 data sample get re-processed and re-stripped. This peak power corresponds to a 40% increase with respect to the yearly average. This peak power is only necessary when approaching the end of the data-taking period and thus **sites could be asked to deploy the CPU power by September 2012 instead of April 2012.**

A similar increase is foreseen for Disk and Tape usage, but in this case the 2011 pledge was not able to cover the re-assessed needs, and therefore the increase fully goes into an increase of the requests to the sites. With respect to the deployment schedule of Disk Storage, it must be noted that for performance reasons a continuous increase following the need is far from optimal. Experience shows that this causes data produced simultaneously to populate the same disk servers. Since this data is likely to be accessed simultaneously later on, this creates bottlenecks in the access to the Disk. However, **a deployment schedule for Disk in 2 steps, in April and September, will most likely be OK** while still allowing some cost optimization.

6. Forecast of 2013 needs

An attempt has been done to forecast the needs for 2013. The basic assumptions going into the estimates presented in this document are that:

- There is no LHC activity.
- The full 2010-2012 samples are re-stripped for the summer conferences and fully re-processed and re-stripped with final calibration and alignments in time for the 2014 winter conferences.
- Extra Monte Carlo samples are produced at the Tier0/Tier1s using the resources that on previous years have been dedicated to quasi-online data processing. This extra power dedicated to simulation allows to complement the contribution from the Tier2s, going from 750 M to 1000 M simulated events.

The resulting estimates are included, together with those from 2011 and 2012, in Table 7-1, Table 7-2, Table 7-3 and Table 7-4. With the given assumption the CPU work and peak power estimates are quite similar to those from 2012, while the storage requirements increase to make space for the newly produced simulation (although partially recovered by the removal of older data) and the re-processing of the full data sample. This increase is about 15% for Disk and 25% for Tape with respect to 2012 requests.

7. Summary

The following tables and figures summarize the results of the estimates discussed in the previous sections for the period between 2011 and 2013.

Work	2011		2012		2013	
	kHS06*y	%	kHS06*y	%	kHS06*y	%
Tier0	17	14	24	15	22	15
Tier1	58	47	80	52	74	51
Tier2	48	39	48	33	48	34
MC	48	39	48	33	65	45
Physics	36	29	46	32	46	32
Reco	17	14	17	12	0	0
Repro	23	19	40	24	33	23
All	124	100	151	100	144	100

Table 7-1: LHCb 2011-13 CPU work estimated need grouped by Tier level and by Activity.

Power	2011		2012		2013	
	kHS06	%	kHS06	%	kHS06	%
Tier0	27	16	34	17	33	17
Tier1	90	55	113	58	110	58
Tier2	48	29	48	25	48	25
MC	48	29	48	25	65	34
Physics	44	27	57	29	57	30
Reco	30	18	30	15		
Repro	90	55	90	46	90	47
All	165	100	195	100	190	100

Table 7-2: LHCb 2011-13 CPU peak power estimated need grouped by Tier level and by Activity.

These numbers are to be compared with the current pledges for 2011 that add up to 21 kHS06 for CERN, 70 kHS06 for the sum of all the Tier1s and 48 kHS06 for the sum of all the Tier2s.

These tables are produced from the detailed usage profiles obtained from the new simulation tool. Figure 7-1 and Figure 7-2 show the obtained profiles for the Tier0 and the sum of the Tier1s. Tier2s are dedicated to simulation with a flat profile and are thus not shown.

For 2012, an increase on the CPU work requirements of 40% with respect to 2011 is foreseen. The integral estimated work is above the 2011 pledge from the sites by 15%, while the peak power need during the last part of 2012 is presents just an increase of 25% with respect to the same estimate for 2011. In both cases this goes over the requested yearly average due to the need of fully reprocess the data sample in time for the winter conferences.




Figure 7-1: Profile of the CPU power usage at the Tier0 for the different activities described in the text.




Figure 7-2: Profile of the CPU power usage at the Tier1s for the different activities described in the text.




Figure 7-3: Profile of the CPU power usage for the different activities described in the text. Contributions from all Tiers are included.

Concerning usage of Disk and Tape Storage resources, Table 7-3 and Table 7-4 present the estimates for the years 2011 to 2012. These estimates correspond to the maximum of the usage obtained from the detailed profiles from Figure 7-4 and Figure 7-5. As reference the 2011 pledges for Disk at the Tier0 and Tier1s are 1.5 and 3.8 PB, respectively. For Tape, the 2011 pledges are 2.5 and 3.9 PB for the Tier0 and the Tier1s respectively.

As can be seen there is a shortage of almost 2 PB in the Disk for 2011 and important deficit of more than 3 PB at the Tier0 for Tape in 2011. The source of the storage shortage comes from 2 reasons, the increase event size (partially recovered by the reduction in the number of replicas), and the new charm samples. The additional tape shortage comes from the move from T1D0 to T2D0 as basic “archival” media at CERN.

The increase between 2011 and 2012 is almost 5 PB for Disk and 9 PB for Tape, i.e. an increase of approximately 80% for Disk and 90% for Tape.

Disk	2011		2012		2013	
	PB	%	PB	%	PB	%
Tie0	1.9	26	3.5	26	3.7	26
Tier1	5.3	74	9.6	73	10.4	74
RAW	0.1	1	0.2	1	0.2	1
SDST	0.1	1	0.1	1	0.1	1
DST	2.3	32	4.1	31	5.8	41
MDST	1.4	20	3.9	30	2.5	18
MC	2.4	33	3.6	28	4.6	32
User	0.9	13	1.1	9	1.3	9
All	7.2	100	13.0	100	14.1	100

Table 7-3: LHCb 2011-2013 Disk Storage estimated need grouped by Tier level and type of Data.

Tape	2011		2012		2013	
	PB	%	PB	%	PB	%
Tie0	5.6	57	10.7	56	13.4	55
Tier1	4.3	43	8.5	44	10.8	45
RAW	2.7	27	4.8	25	4.8	20
SDST	1.2	12	2.6	14	3.5	15
DST	2.7	27	5.0	26	6.7	28
MDST	1.3	13	3.4	18	4.2	18
MC	2.1	21	3.4	17	4.8	20
All	10.0	100	19.2	100	24.1	100

Table 7-4: LHCb 2011-2013 Tape Storage estimated need grouped by Tier level and type of Data.




Figure 7-4: Disk Storage usage profile for each data type.




Figure 7-5: Tape Storage usage profile for each data type.

For sake of completeness Table 7-5 shows the 2011 result of the 2011 reassessment if no additional trigger bandwidth for charm is allowed.

No Charm	CPU		Disk	Tape
	kHS06*y	kHS06	TB	TB
Tier0	15	23	1.5	4.5
Tier1	50	78	4.2	3.5
Tier2	48	48		
All	113	149	5.8	8.0

Table 7-5: Re-assessment of 2011 needs without an extra 1 kHz of trigger rate for charm physics.

Therefore the re-assessment of 2011 period shows an increase of the peak CPU power needs beyond the current 2011 pledges of the Sites, necessary to reprocess the full data sample in time for the 2012 winter conferences. The addition of an extra 1 kHz of triggers for charm physics motivates the increase. For 2012 a moderate increase is necessary to handle the additional data samples not foreseen in older estimates. The large increase in the event size has been partially address by reducing the number of replicas of the data, after successful experience from 2010. Still the additional charm data, and the move from T1D0 to T2D0 as archival media at CERN, produce important increases in 2011 and subsequently also in 2012. These increases have a well-defined reason, based on LHCb physics case, and a profile that would advise to split the deployment of new resources in 2 points of the year. For 2013 a first attempt has been done to determine the extra needs and, provided the current LHC schedule is kept, expect for extra Tape, very little or almost no increase is foreseen for CPU or Disk.