

LHCb Computing requirements for 2009-10

Draft document

April 6, 2009

Introduction

This document re-assesses the Computing resource needs of LHCb in view of the change of schedule of the LHC after the September 19th incident.

The LHCb Computing Model splits the computing tasks in categories and assigns them to a set of sites as follows:

- Real data recording from the experiment and distribution for data preservation out of CERN: *Tier0*
- Real data reconstruction (first pass reconstruction as well as reprocessing): *Tier0 + Tier1s*
- Physics selection (a.k.a. stripping) to reduce data samples to be further analysed by physics groups: *Tier0 + Tier1s*
- Physics analysis, based on the selected events. This analysis can be done at the group level or at the individual level: *CERN-CAF + Tier1s-AF*
- Monte-Carlo simulation, digitisation and reconstruction: *all sites with lower priority on Tier0/1.*

Assumptions for the LHC run

The LHCb data taking rate is explained in the LHCb Computing and Trigger TDRs. It is assumed to be 2 kHz, including physics triggers as well as control channels and calibration events. This allows to better control systematic errors that are essential for the high precision requirements of LHCb. The corollary is that the simulation requirements are slightly less demanding and used mainly for computing signal efficiencies and identifying main sources of background.

The period considered in this document covers April 2009 to March 2011. The LHC running model used here assumes the following pattern of physics data taking:

1. April-September 2009: no LHC running
2. October 2009-March 2010: $1.7 \cdot 10^6$ seconds of physics
3. April-October 2010: $4.3 \cdot 10^6$ seconds of physics
4. October 2010-March 2011: LHC shutdown

There are many uncertainties still in the running mode of the LHC, in particular for what concerns the ratio between luminosity and number of bunches (and therefore the amount of pile-up). LHCb had always assumed a very low pile-up, but in view of this uncertainty, we have introduced a moderate contingency on the event sizes, processing time as well as simulation time in order to allow for this larger multiplicity.

Assumptions for LHCb Computing

Period 2009-1

This period will be used for simulating events required for preparing the data taking 2009-10. Event samples will be simulated with the expected conditions of the LHC (5 TeV, 50 ns bunch spacing, intermediate luminosity). Analysis of these simulated data as well as formerly simulated data will continue, with an increasing part made on the Grid, but still a substantial part done on LXBATCH at CERN.

Period 2009-2

During this period emphasis will be put on understanding the detector (alignment, calibration). Data at very low luminosity will be taken with very loose trigger cuts in order to make Early Physics measurements (cross sections, branching ratios...). It is expected that many reprocessing passes of the raw data will be necessary (4 passes assumed). No stripping is expected, as this period will not be devoted to b-physics. However data reduction will be done by reducing the event size, using the μ DST format developed in LHCb (reduction by a factor 10).

These early datasets will be extensively analysed both by detector experts and physicists. The take off of Grid batch analysis should be higher and up to 50% of analysis is expected to take place outside CERN.

Simulation will continue, after tuning of the event generation, detector simulation and digitisation according to the first data. This simulation will be used for extracting acceptances for the first physics publications.

Period 2010-1

It is expected that LHCb will use its final and tuned High Level Trigger for b-physics, collecting as much luminosity as possible. The aim is to concentrate on a few specific channels, for which the expected integrated luminosity of 0.1 to 0.3 fb^{-1} allows to make a significant impact, equalling or improving the Tevatron results.

It is expected that several reprocessing passes will be necessary, as well as multiple stripping passes (3 passes over the whole period).

Intensive analysis of these data will take place on the Grid (60%) as well as at CERN. The average dataset size for analysis is assumed to be 10^6 events and about 1000 such jobs are assumed every week (that will in turn be split in many more real jobs).

Simulation of signal and background samples will continue for b-physics, as well as preparatory simulation at the LHC nominal settings (for 2011 data taking conditions: 7 TeV, 25 ns bunch spacing, $2 \cdot 10^{32} \text{ cm}^{-2}\text{s}^{-1}$ or higher)

Period 2010-2

During the LHC shutdown, there will be at least one full reprocessing of the available dataset, including stripping. Simulation will continue for physics publications and analysis will be at a climax in order to present results at the

2001 Winter conferences. Studies will continue for HLT and stripping with the 2011 data taking conditions (simulation and analysis).

Computing Resources

Using the above assumptions, we have computed the needs for CPU resources as well as “disk” (TxD1 service classes) and “tape” storage (T1Dx service classes).

CPU resources are expressed in the new HEP-SPEC06 (SH06) unit. We are providing the integral estimated CPU in each of the calendar years 2009 and 2010 (as this is what can be compared to previous requirements), in kSH06.years, as well as the estimated power needed at the beginning of each semester, making some assumptions on the ramp up of needs following data collection.

At CERN, an additional 2 kHEP-SPEC06 has been included in order to account for the build servers and the central DIRAC services. Similarly 1 kHEP-SPEC06 was accounted for VOBoxes at Tier1s and DIRAC services at PIC.

Date	Site	kSH06
Oct'09	CERN	17
	Tier-1	31
	Tier-2	30
Apr'10	CERN	25
	Tier-1	45
	Tier-2	38
Oct'10	CERN	28
	Tier-1	49
	Tier-2	40

Table 1: CPU Power needed in place to meet LHCb requirements for the 6 month period commencing (a) October 2009, (b) April 2010 and (c) October 2010

Disk and tape resources are estimated at the end of each semester and should be made available at the beginning at the period. Obviously the storage will be filled up progressively with time. The resources quoted here correspond to the data collected or simulated as of April 2009. It is assumed that existing data will be removed progressively from storage when they are replaced by newer data. Note that 1 TB = 10^{12} Bytes and not 1024^4 bytes (a 7% difference).

An additional best estimate of disk needs as cache for T1D0 service classes has been added, being clear that the size of the caches depends on many site-related parameters such as number of tape drives, storage configuration (e.g. distinct write and read disk pools) that can only be assessed by the sites. We have assumed a total cache size of 70 TB per site (CERN and Tier1s)

Date	Site	TB
Oct'09	CERN	780
	Tier-1	2800
	Tier-2	20
Apr'10	CERN	1470
	Tier-1	4400
	Tier-2	20
Oct'10	CERN	1470
	Tier-1	4400
	Tier-2	20

Table 2: Disk requirement needed in place to meet LHCb requirements for the 6 month period commencing (a) October 2009 , (b) April 2010 and (c) October 2010.

Date	Site	TB
Oct'09	CERN	1200
	Tier-1	1300
	Tier-2	0
Apr'10	CERN	1800
	Tier-1	2100
	Tier-2	0
Oct'10	CERN	2300
	Tier-1	2900
	Tier-2	0

Table 3: Tape requirement needed in place to meet LHCb requirements for the 6 month period commencing (a) October 2009, (b) April 2010 and (c) October 2010.

Comparison with previous requirements and pledges

Previous CPU requirements have been expressed in terms of integrated CPU power over calendar years. Therefore for the sake of comparison, we have used the same metrics for the new requirements, and converted the MSI2k into kHEP-SPEC06 using the agreed conversion factor of 4.

kSH06*year	2009	2010
CERN T0 + CAF	11.37	19.19
Tier1s	16.00	33.99
Tier2s	21.86	31.48
Total	49.23	84.66

Table 4: New integrated CPU resources needed for calendar years 2009 and 2010.

kSH06*year	Old 2009	Pledges '09	Old 2010
Online farm	3.60		3.60
CERN	4.22	4.20	6.11
Tier1s	19.88	20.22	27.36
Tier2s	45.51	35.37	45.51
Total	73.20	59.79	82.58

Table 4: Old integrated CPU resources needed for calendar years 2009 and 2010. 2009 pledges are also indicated

One can notice a substantial increase of the CPU requirements at CERN in order to cope with the observed latency of users moving to the Grid as well as the need for fast feedback to the detector in terms of calibration and alignment.

The CPU requirements at Tier1's have decreased in 2009, and marginally increased in 2010 due to the anticipated higher number of reprocessing passes. There is a clear decrease in CPU requirements at Tier2s due to less requests for simulation, following the delay of the LHC.

Date	Site	TB	Pledges '09
Apr'09	CERN	991	991
	Tier-1	2759	2709
	Tier-2	23	371
Apr'10	CERN	1278	
	Tier-1	3250	
	Tier-2	23	

Table 6: Old disk requirement needed in place to meet LHCb requirements for the one year period commencing (a) April 2009 and (b) April 2010.

The disk requirements for 2010 have increased slightly at CERN and Tier1s due to changes in the stripping and analysis strategy of LHCb: it is required to keep the full DST information for all streams while in the TDR some streams has reduced information.

Date	Site	TB	Pledges '09
Apr'09	CERN	2270	2270
	Tier-1	3070	3264
	Tier-2	0	-
Apr'10	CERN	4207	
	Tier-1	5864	
	Tier-2	0	

Table 7: Old tape requirement needed in place to meet LHCb requirements for the one year period commencing (a) April 2009 and (b) April 2010.

The tape requirements have substantially decreased at CERN and Tier1s due to the late start of the LHC. One notices essentially a shift of more than one year in tape requirements.