# LHCb Computing Resource usage in 2010

## LHCb Public Note


Issue: V4
Revision: 6

Reference: LHCb-PUB-2011-008
Created: 21th Feb 2011
Last modified: 3nd Mar 2011



**Prepared By:** LHCb Computing Project
R. Graciani/Editor

# Abstract

This notes summarizes the uses of computing resources during 2010 as support document for the report due to CRRB by March 1st 2011. Computing resources are requested to WLCG and pledged by the sites to be available as of April 1st of the corresponding year. Official WLCG accounting portals have a granularity of 1 month. Thus, the present document covers the period from April 1st 2010 until February 1st 2011.

The data used in the following sections is taken from either the WLCG accounting portal: https://www3.egee.cesga.es/gridsite/accounting/CESGA/egee_view.php, or from the LHCb DIRAC Accounting portal: http://lhcbweb.pic.es/DIRAC/LHCb-Production/visitor/systems/accountingPlots/job, in what respect usage of CPU resources. For the usage of the Storage resources snapshots of the situation as seen by the SLS sensors: https://sls.cern.ch/sls/service.php?id=LHCb-Storage, and the DIRAC StorageUsage Service: dips://lhcb-serv1-dirac.cern.ch:9151/DataManagement/StorageUsage is used.

# Document Status Sheet

Table 1-1 Document Status Sheet

| **1. Document Title: LHCb Computing Resource usage in 2010** | | | |
|---|---|---|---|
| **2. Document Reference Number: [Document Reference Number]** | | | |
| **3. Issue** | **4. Revision** | **5. Date** | **6. Reason for change** |
| Draft | 1 | 22 February 2011 | First version |
| Draft | 2 | 22 February 2011 | |
| V1 | 3 | 28 February 2011 | Got the reference LHCb-PUB-2011-008 |
| V2 | 4 | 2nd March 2011 | Various minor edits |
| V3 | 5 | 2nd March 2011 | Review of changes |
| V4 | 6 | 3rd March 2011 | Incorporate some final comments |

# Table of Contents

# List of Figures

# List of Tables

# 1. Introduction

LHCb is asked periodically to provide estimates of its computing resource needs to the Computing RRB. The requests are scrutinized by the C-RSG. Once approved, they are passed to the WLCG resource centres that present a pledge for the corresponding period. These pledges are finally approved together with a deployment calendar. Finally, the experiment is asked to justify the usage of the provided resources, and to produce updates for next period and new estimates for the following periods.

This document reports on the usage of resources from 1st April 2010 to 31st January 2011. The computing resource estimates for 2012 and 2013 and a review of the already approved request for 2011 are provided in a separate document (LHCb-PUB-2011-009).

At the Computing RRB in April 2009, LHCb presented a resource request for the 2009 running, in fact for a period covering April 2009 to April 2011. This request was then updated and a final version was approved by the C-RSG in June 2009 Table 1-1 shows the approved values for 2010.

| Date | Site | kHS06 | Disk (TB) | Tape (TB) |
|---|---|---|---|---|
| Apr'10 | CERN | 21 | 1290 | 1500 |
| | Tier-1 | 41 | 3290 | 1800 |
| | Tier-2 | 36 | 20 | 0 |
| Oct'10 | CERN | 23 | 1290 | 1800 |
| | Tier-1 | 44 | 3290 | 2400 |
| | Tier-2 | 38 | 20 | 0 |

Table 1-1: CPU power, disk and tape storage needed in place to meet LHCb requirements for the 6 month period commencing (a) April 2010 and (b) October 2010, as best estimate end of June 2009.

The LHCb shares of the pledges approved by the CCRB in April 12th 2010 are summarized in Table 1-2, taken from http://lcg.web.cern.ch/lcg/Resources/WLCGResources-2009-2010_12APR10.pdf.

| Date | Site | kHS06 | Disk (TB) | Tape (TB) |
|---|---|---|---|---|
| Apr'10 | CERN | 23 | 1290 | 1800 |
| | Tier-1 | 43 | 3254 | 3036 |
| | Tier-2 | 42 | 430 | 0 |

Table 1-2: CPU power, disk and tape storage offered by the WLCG to LHCb for the 12 month period starting April 2010, as approved by CCRB in April 2010.

This document is organized as follows: section 2 presents a summary of the data-taking conditions during 2010, section 3 presents the usage of computing resources as seen by the WLCG accounting portal, while section 4 presents the same information as seen by the LHCb DIRAC portal. The usage of Storage is shown in section 5. Finally, everything is summarized in section 6.

# 2. 2010 Data-Taking conditions

Since March 30th 2010, LHCb has been collecting real data with a quite high efficiency as shown in Figure 2-1. The data-taking conditions during this period are summarized in Figure 2-2.

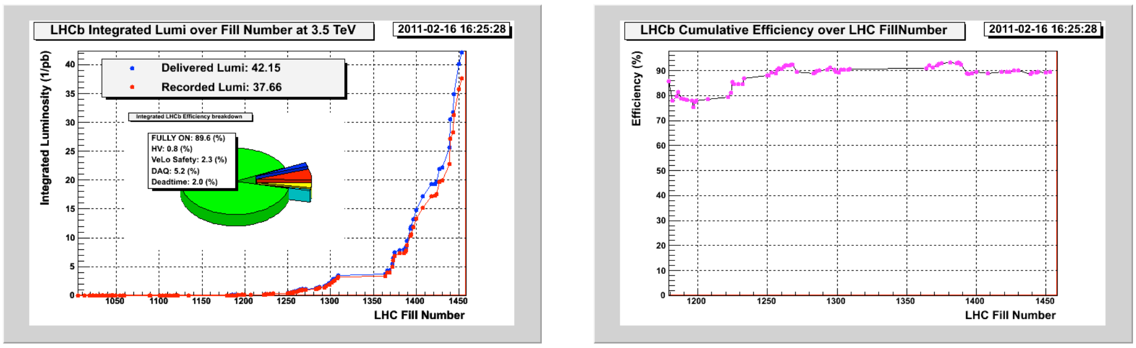


Figure 2-1: LHC delivered luminosity at LHCb interaction point and LHCb recorded luminosity for 2010 (left), and corresponding DAQ efficiency (right) as a function of the LHC fill number.

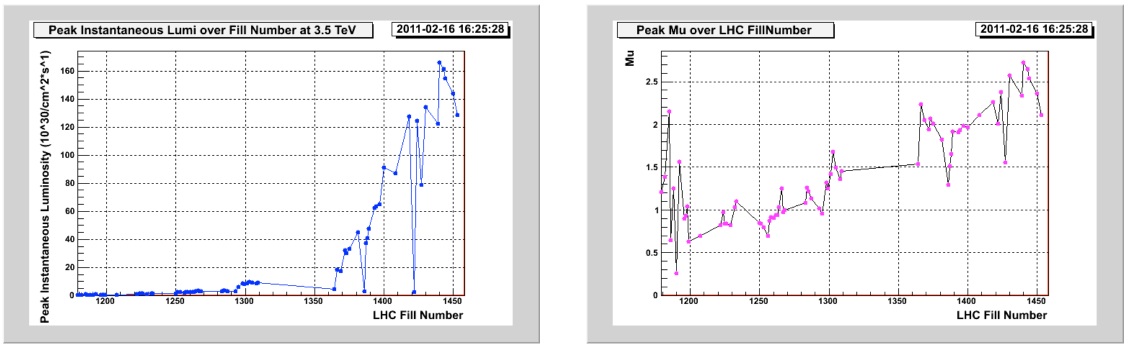


Figure 2-2: Peak instantaneous luminosity delivered to LHCb during the 2010 data-taking period (left) and the corresponding peak value of the average number of visible proton-proton collisions per bunch crossing as a function of the LHC fill number.

The increase in luminosity has been achieved by a significant increase in the average number of visible proton-proton interaction per bunch crossing (mu) with respect to the LHCb nominal operation point at mu=0.4, as shown in Figure 2-2. This has added new constraints to the LHCb computing model that will be discussed later on.

Even though the integrated luminosity is far from what was expected when the request was made, the detector DAQ rate has been most of the time around the nominal 2 kHz. This has been driven by the need to understand the performance of the detector as soon as possible, as well as by the possibility to extend the LHCb physics programme to some inclusive and charm measurements during the initial months of LHC operation at low luminosity.

The amount of collected RAW data during 2010 was 181 TB, of which 155 TB corresponds to physics data distributed to Tier1s. The rest corresponds to detector calibration data that remains at CERN. The distribution of RAW data between the different Storage Elements is shown in Table 2-1.

| 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.

The average size of the RAW files has been 1 GB during the whole 2010 data-taking period. Under nominal conditions with a mu=0.4 and 2 kHz of HLT rate, the expected RAW size for a full LHC year, with $5 \cdot 10^6$ data-taking seconds, would have been 250 TB. This is not too far from the current situation but with an integrated luminosity over 1 order of magnitude smaller than the nominal.

# 3. Usage of CPU from WLCG Accounting for 2010

The data presented in this section has been taken from the CESGA accounting portal. Figure 3-1 has been produced using the data from April 2010 until January 2011 (both included). It shows the monthly use of computing resources by LHCb in HS06·month for the Tier0, the Tier1s and the Tier2s.

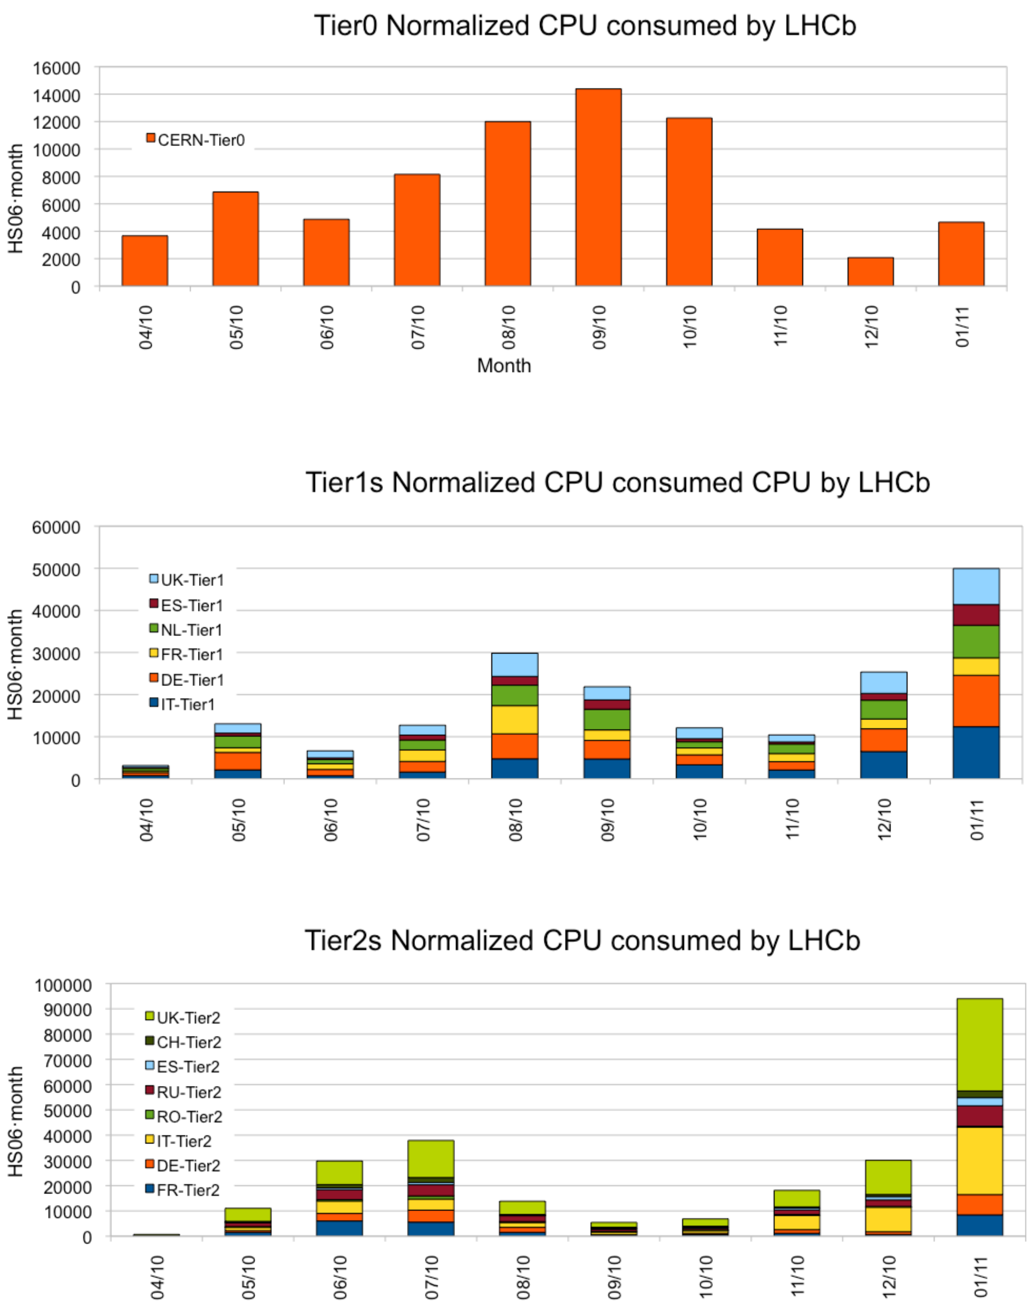


Figure 3-1: Normalized CPU consumed by LHCb from April 2010 to Jan 2011.

The average normalized CPU power consumed in the given period is summarized in Table 3-1 together with the measured average efficiency in the usage of the CPU resources.

| 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%[1] |
| 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 3-1: Average CPU power used and efficiency[2] as reported by WLCG Accounting portal.

A discussion of the time profile of the CPU usage and of the CPU efficiencies can be found in section 4.

[1] At least one of the sites publishes "unnormalized" elapse time.

[2] 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.

# 4. Usage of CPU from LHCb DIRAC Accounting for 2010

An alternative view of the same data can be obtained using the LHCb DIRAC Accounting portal. Since the portal is specific to LHCb, using it allows a more detailed classification of the different jobs types. Figure 4-1 shows the number of simultaneously running jobs from April 2010 to January 2011 classified by the type of the jobs. The figure shows the relative contributions of the different job types as well as the variation of their relative importance as a function of time. Table 1-1 shows the total amount of CPU and wall clock time consumed by the different job types.

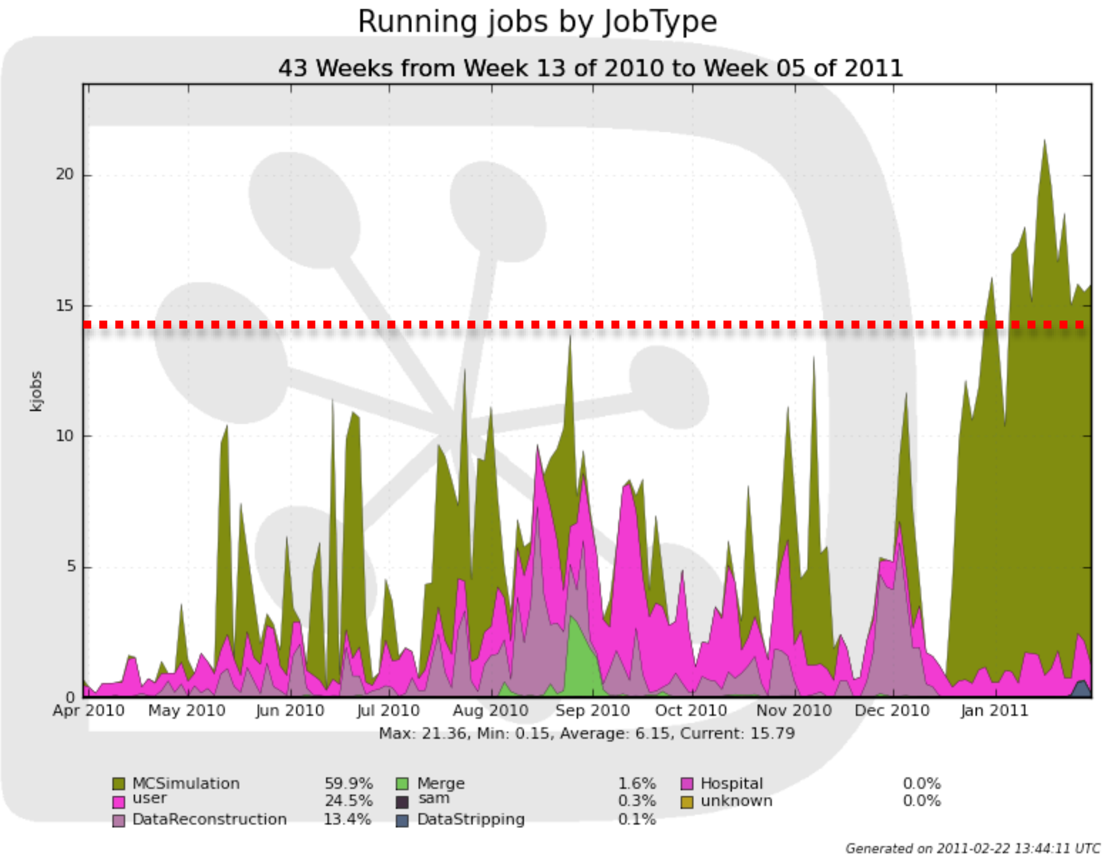


Figure 4-1: Number of concurrently running jobs as a function of time for the different LHCb jobs types. The dashed line corresponds to the expectation from the pledge and the measured average CPU power (see Table 4-2).

| | DataRecons. | DataStripping | MCSimulation | Merge | Sam | User |
|---|---|---|---|---|---|---|
| CPU | 232,930 | 1,318 | 1,071,925 | 2,143 | 1,740 | 323,644 |
| Time | 257,744 | 2,836 | 1,192,443 | 31,125 | 6,218 | 468,685 |
| Efficiency | 90.4% | 46.5% | 89.9% | 6.9% | 28.0% | 69.1% |
| Fraction | 13.2% | 0.1% | 60.9% | 1.6% | 0.3% | 23.9% |

Table 4-1: Amount of CPU days, wall clock time days and CPU efficiency for the different LHCb jobs from April 2010 to January 2011.

The ratio between CPU days and wall clock time gives the efficiency achieved in the use of the computing resources. When all job types are added together an overall efficiency of 83.4% is found. This efficiency should be compared with the 88% from the last raw of Table 3-1. These two efficiencies are calculated in a slightly different way. For the 83.4% provided by DIRAC accounting the pilot reports the raw CPU and wall clock times used by the payloads that are added for all executed jobs. The ratio of the total CPU over the total wall clock is determined. For the 88% provided by WLCG accounting each site provides normalized CPU and wall clock times as seen by the batch system, combining the payload, the DIRAC pilot and the gLite wrapper. The ratio between these normalized totals defines the WLCG efficiency.

As can be seen, during most of the year the CPU usage has not reached the requested average power.

In order to compare the contributions from the different resource centers, the number of concurrently running jobs by country at Tier0/1s are shown in Figure 4-2. The top plot corresponds to centrally driven activities, mostly real data processing activities (except for the last 45 days when these sites were used also for simulation). It shows an increase in the usage until September, due to the unexpected size of the events resulting from the additional interactions in the same bunch crossing to achieve higher luminosities. This was causing many problems to the reconstruction software that had been tune on much simpler events. At that point the re-optimized code was available and the full data sample taken so far was reprocessed. During October and November the plot shows an increasing activity corresponding to the increasing luminosity delivered by LHC. The bump around the beginning of December corresponds to the full reprocessing of 2010 data. The final part of the plot corresponds to Monte Carlo simulation, taking advantage of these resources in addition to those in use at Tier2s.

The bottom plot of Figure 4-2 shows the activity corresponding the user jobs, also in the Tier0/1s. This plot shows an important increase of the user activity during August and September, coinciding with the availability of the first partial reprocessing. The relative contribution of the different LHCb activities in those centers is 55% for centrally driven productions and 45% for user activity. Outside the Tier1s the main usage of the CPU resources is Monte Carlo simulation (as foreseen in the Computing Model).

Figure 4-3 shows the number of concurrently running jobs outside the Tier0/1s, separately for official Tier2 (top) sites and other sites (bottom). A bit more than 75% of the resources come from Tier2s while the other 23% is obtained from other sites.

For the whole 2010 period analyzed, the average contribution from Tier0/1s, shown in Figure 4-2, has been about the same as the contribution from Tier2s and other sites not explicitly bound to LHCb, shown in Figure 4-3.

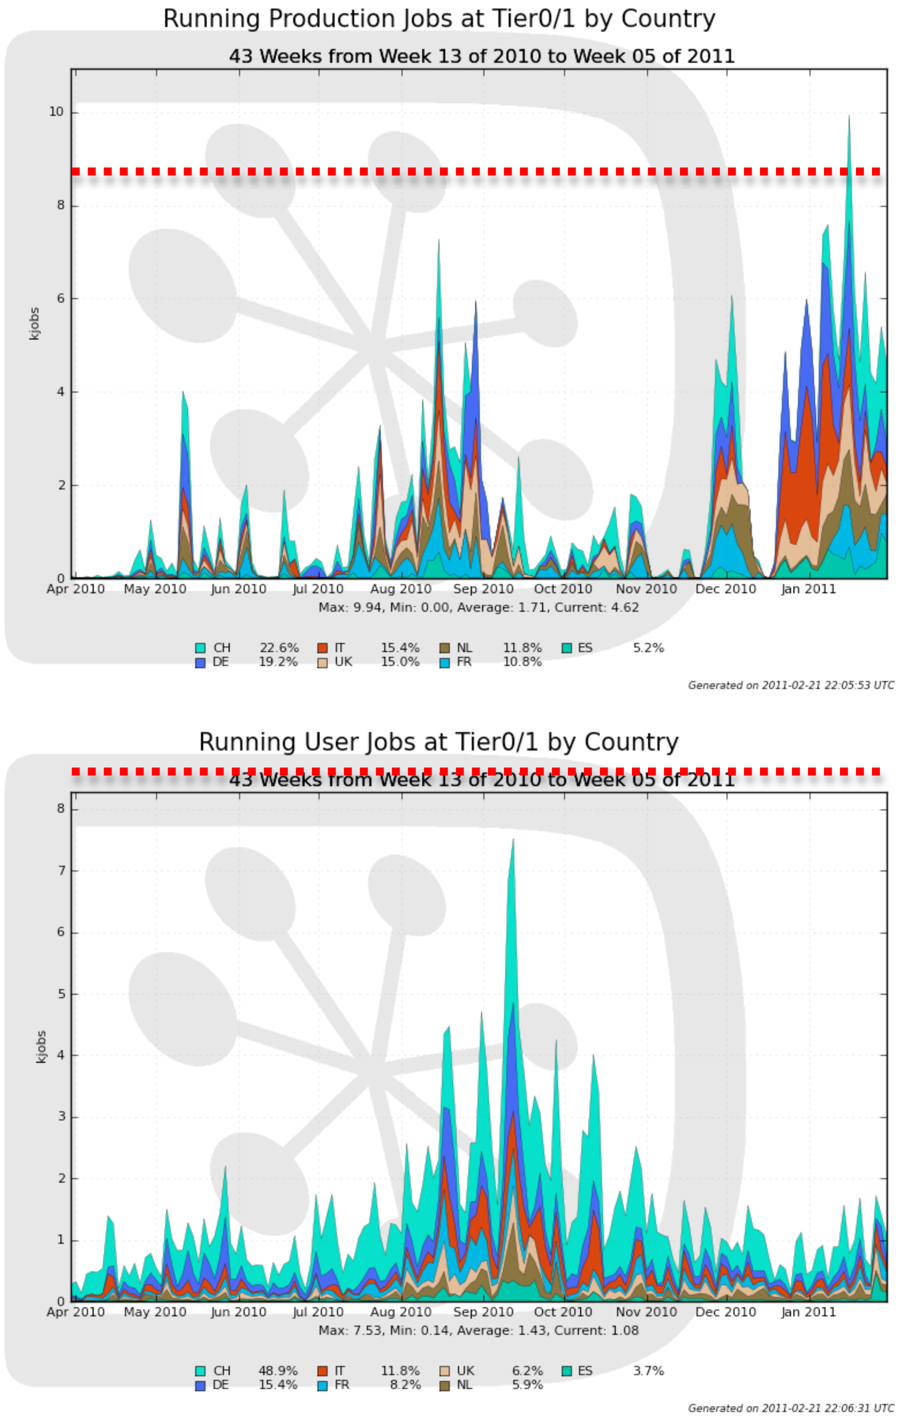


Figure 4-2: Number of concurrently running Production (top) and User (bottom) jobs by country as a function of time at the LHCb Tier0/1 centers. The dashed line corresponds to the expectation from the pledge, using the measured average CPU power (see Table 4-2).

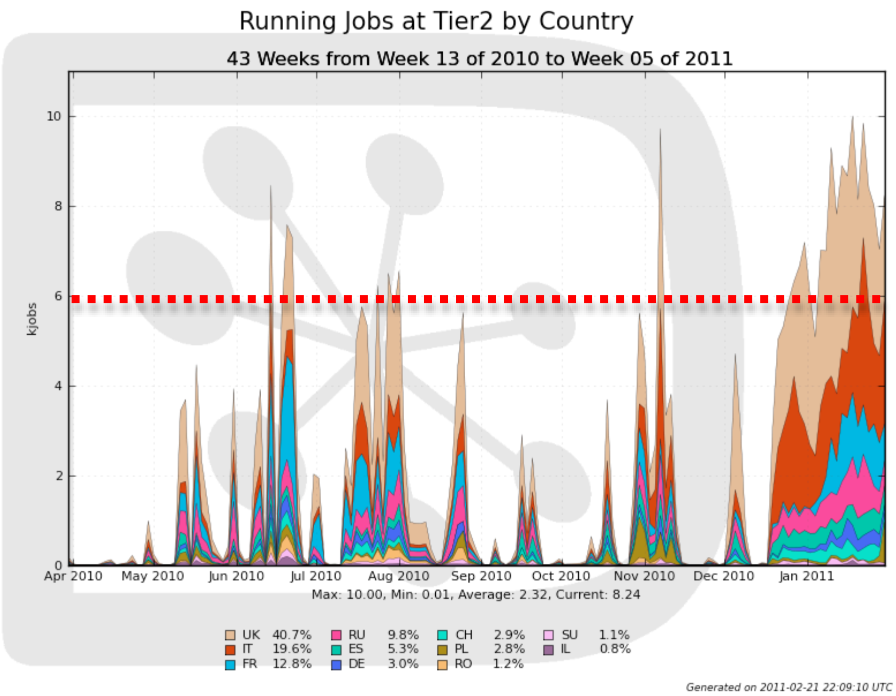


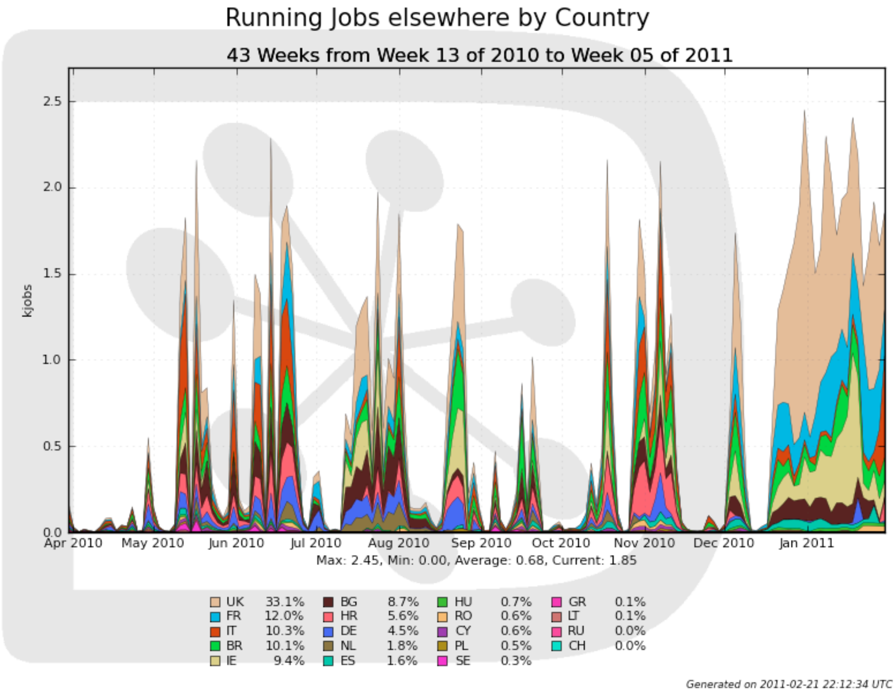


Figure 4-3: Number of concurrently running jobs as a function of time by country at the LHCb Tier2 centers and elsewhere. The dashed line corresponds to the expectation from the pledge, using the measured average CPU power (see Table 4-2).

Taking the measured CPU requirement of 2.6 kHS06·s/evt, we use some large Monte Carlo simulation productions (in particular the high momentum Di-muon samples with 24 million simulated events) to determine an absolute average normalization for the raw CPU consumed by all LHCb jobs at each site. The results are shown in Table 4-2.

| | Ref.Jobs | Frac. | Ref.CPU | Norm | All CPU | All CPU | Frac. | Av.CPU Power |
|---|---|---|---|---|---|---|---|---|
| | | % | day | HS06 | day | 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** | **39364[3]** |

Table 4-2: Raw CPU work (in day), 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.

For the Tier0 and Tier1s an alternative normalization can be determined using the last reconstruction pass over the full 2010 data sample. The input files are randomly distributed among the sites with fractions based on their CPU pledges for 2010. So, even if the event size has changed during the data-taking period, on average the samples should be comparable for all sites. The total CPU work for this task is estimated to be 540 kHS06·day·c, based on a reference measured value for the last part of the sample, with <mu> = 2, where c is a constant, smaller than one, taking into account the difference between the average mu of the full sample and the average mu of the reference sample. The result is summarized on Table 4-3. Comparing with Table 4-2, the value of the constant is estimated to be 75%.

| Ref.Jobs | Frac. | Ref.CPU | Norm | All CPU | All CPU | Frac. | Av. CPU Power |
|---|---|---|---|---|---|---|---|

[3] 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.

| | | % | day | HS06·c | day | kHS06·day·c | | HS06·c |
|---|---|---|---|---|---|---|---|---|
| **Tier0** | **31561** | **25.1** | **13114** | **10.4** | **265480** | **2760** | **34.1** | **9019** |
| DE-T1 | 16282 | 17.3 | 6822 | 10.3 | 115914 | 1195 | 22.4 | 3905 |
| ES-T1 | 5050 | 5.4 | 1808 | 12.1 | 38782 | 468 | 8.8 | 1529 |
| FR-T1 | 20129 | 21.3 | 8813 | 9.9 | 80709 | 796 | 15.0 | 2602 |
| IT-T1 | 11174 | 11.8 | 4255 | 11.3 | 99219 | 1126 | 21.1 | 3678 |
| NL-T1 | 23892 | 25.3 | 9272 | 11.1 | 76726 | 854 | 16.0 | 2791 |
| UK-T1 | 17844 | 18.9 | 7367 | 10.5 | 89448 | 936 | 17.6 | 3058 |
| **Tier1s** | **94371** | **74.9** | **38337** | **10.6** | **500798** | **5325** | **65.7** | **17402** |
| **Total** | **125932** | **100.0** | **51451** | **10.6** | **766278** | **8102** | **100.0** | **26476** |

Table 4-3: Raw CPU work (in day), Normalized CPU work (in kHS06·day·c), and average CPU power in kHS06·c for all LHCb Jobs. Normalization is calculated based on the known requirements of a Reference Full reconstruction pass processing $2.35{\cdot}10^9$ events and requiring 540 kHS06·day·c. Only Tier0 and Tier1s are included. Comparing with Table 4-2, the value of the constant c is found to be 75%.

A comparison between the results in Table 4-1, Table 4-2 and Table 4-3 is presented in Table 4-4. The absolute measurements from Table 4-2 are presented in the LHCb columns, those from Table 4-1 in the WLCG column and those from Table 4-3 in the LHCb*. The comparison between the LHCb and LHCb* columns shows a very good agreement that allows to determine the missing constant as c = 75%. Since the sample of jobs used to determine the calibration constants is totally different the results can be used to determine the systematic error of the method. The relative differences between LHCb and LHCb* calibrations show an RMS of 7%, which is an estimate of the systematic error of the method.

| | LHCb | | WLCG | | LHCb* | | Comparison | |
|---|---|---|---|---|---|---|---|---|
| | HS06 | Site/Tier0 | HS06 | Site/Tier0 | HS06·c | Site/Tier0 | (W-L)/L | (L*-L)/L |
| **Tier0** | **6640** | **100.0%** | **7166** | **100.0%** | **9019** | **100.0%** | **7.9%** | **35.83%** |
| DE-T1 | 3005 | 45.3% | 4039 | 56.4% | 3905 | 43.3% | 34.4% | 29.95% |
| ES-T1 | 1070 | 16.1% | 1412 | 19.7% | 1529 | 17.0% | 32.0% | 42.90% |
| FR-T1 | 1834 | 27.6% | 2437 | 34.0% | 2602 | 28.9% | 32.9% | 41.88% |
| IT-T1 | 2440 | 36.7% | 3817 | 53.3% | 3678 | 40.8% | 56.4% | 50.74% |
| NL-T1 | 1981 | 29.8% | 3188 | 44.5% | 2791 | 30.9% | 60.9% | 40.89% |
| UK-T1 | 2299 | 34.6% | 3255 | 45.4% | 3058 | 33.9% | 41.6% | 33.01% |
| **Tier1s** | **12745** | **191.9%** | **18148** | **253.3%** | **17402** | **192.9%** | **42.4%** | **36.54%** |
| CH-T2 | 441 | 6.6% | 771 | 10.8% | | | 74.8% | |
| DE-T2 | 549 | 8.3% | 2114 | 29.5% | | | 285.1% | |
| ES-T2 | 782 | 11.8% | 830 | 11.6% | | | 6.1% | |
| FR-T2 | 2164 | 32.6% | 2553 | 35.6% | | | 18.0% | |
| IT-T2 | 2822 | 42.5% | 5739 | 80.1% | | | 103.4% | |
| RO-T2 | 101 | 1.5% | 238 | 3.3% | | | 135.6% | |
| RU-T2 | 1656 | 24.9% | 2610 | 36.4% | | | 57.6% | |
| UK-T2 | 6514 | 98.1% | 9420 | 131.5% | | | 44.6% | |
| **Tier2s** | **15411** | **232.1%** | **24273** | **338.7%** | | | **57.5%** | |

Table 4-4: Comparison between LHCb and WLCG measurements of the normalized CPU work consumed. The HS06 (and HS06·c) columns show the measured value, the adjacent column the relative contribution with respect to Tier0 (CERN). The last columns present the difference between measurement using the absolute LHCb as reference, the WLCG on the left and the alternate LHCb* on the right.

On the other hand, when comparing the LHCb and WLCG measurements differences are rather large. Just for the Tier1s the WLCG value is over 40% higher in a direct comparison, or 30% if we use CERN as reference to correct for a possible overall scale factor. The DE, ES and FR Tier1s show a very similar scale factor, larger than CERN, while the IT, NL and UK Tier1s show an even larger deviation. For Tier2s the situation is even worst. While ES and FR show a factor similar to CERN, other large contributors like CH, RU and UK show larger deviations while for DE and IT the differences are huge.

## 4.1. Job Failures

In order to have a real measurement of the efficiency in the usage of the CPU we should also take into account the resources invested on failed jobs.

Figure 4-5 shows the fractions of Simulation jobs for different final states. It can be seen how during an initial period, July and August, a large fraction of jobs were failing due to infinite loops in the simulation code when moving from 32 to 64 bit precision. After the problem was detected the failure rate has gone down significantly and since September the failure rate is less than 5%.

Figure 4-5 shows the fractions of Reconstruction jobs for different final states. Again, during July and August there was a significant rate of jobs failing due to the larger event sizes, making the duration of the reconstruction jobs go beyond the time limit of the execution queues at the sites. Once the reconstruction code was optimized for the new environment, the failure rate dropped to the 5% level.

Figure 4-6 shows the fractions of User jobs for different final states. In this case there are two kinds of dominant error conditions. On the one hand, the most important contribution is from jobs exceeding the wall clock limit of three days. On the other hand, the second most frequent reason of failing user jobs is their long execution time that causes them to go beyond the limit of the execution queues at the sites. These problems add up to about 30% of the used resources, and are very much related to the higher multiplicity of the events that cause the combinatorics of some analysis code to explode. One can assume that users will slowly improve their code, reduce the input data samples of the jobs and eventually reduce the impact of these problems. The remaining contributions add up to less than 10%.

It is interesting to notice that for Simulation and User Jobs the actual Job failure rate is largely reduced with respect to the fraction of the resources those jobs are consuming. Successful Simulation jobs ("Execution Complete" + "Pending Request") consume 82% of the resources while they represent 92% of the simulation jobs, and successful User jobs consume 63% of the resources but represent 84% of the User jobs. While for Reconstruction, the successful jobs consume 85% of the resources and are just 82% of the jobs. This means that User and Simulation jobs tend to fail after having used more resources than the average jobs while the case is the opposite for Reconstruction jobs that are more likely to fail in the initialization phase or else go through with fewer problems. For instance "Application Finished with Errors" is most often caused by problems with the shared area while trying to setup the environment for the application. While "Stalled" jobs typically have consumed a lot of resources before being killed when reaching the limit of the execution queue at the sites.

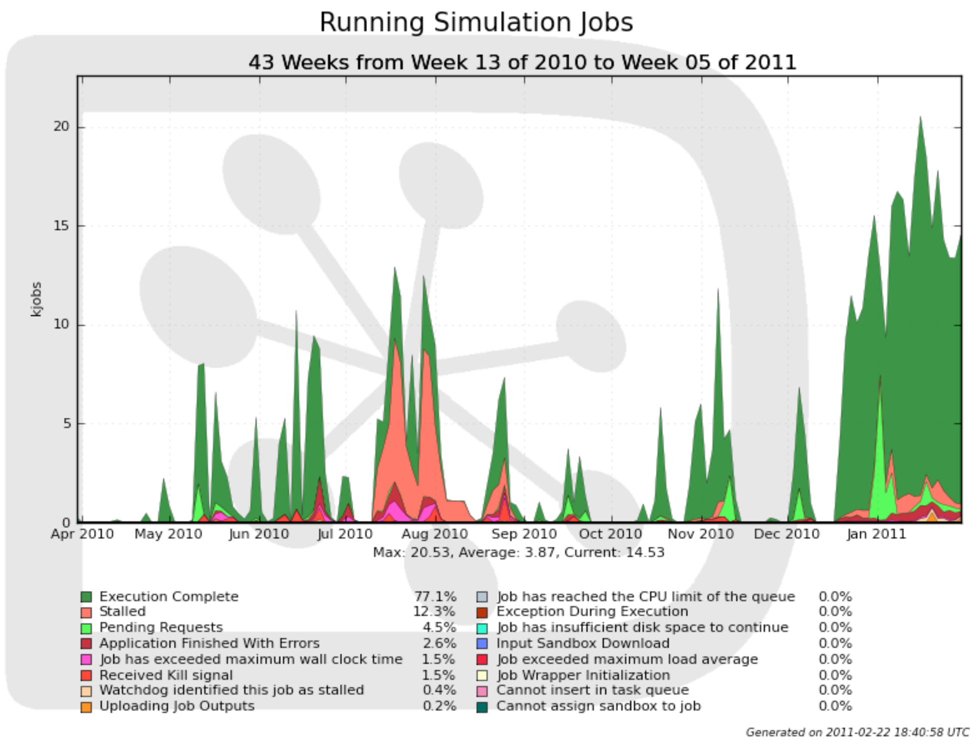


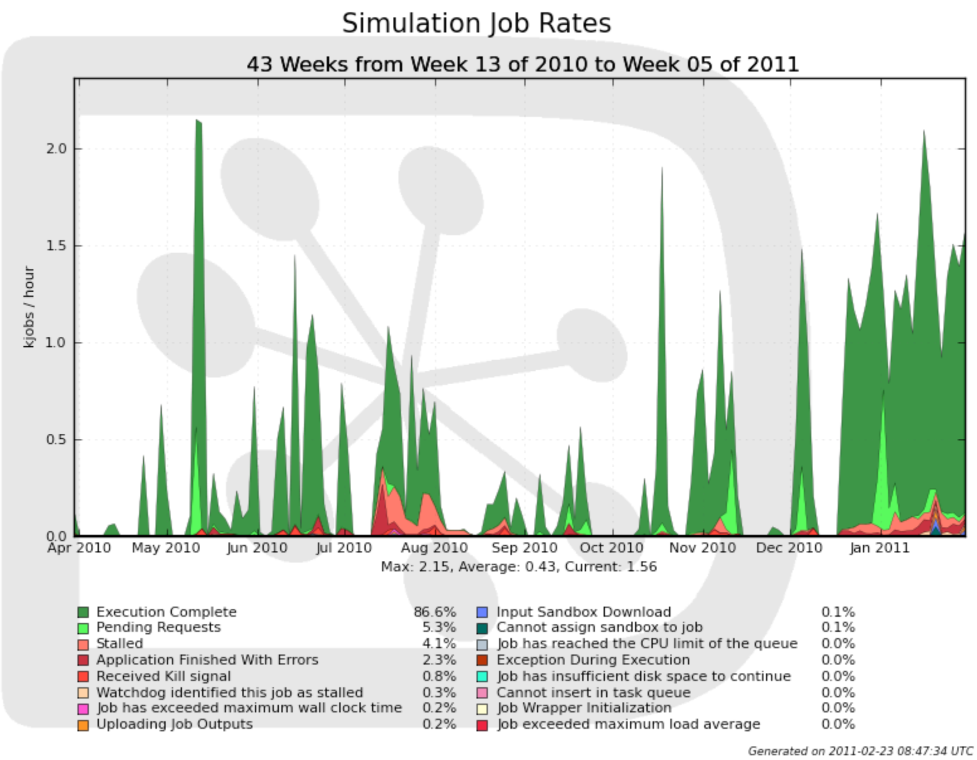


Figure 4-4: Number of simultaneously running Simulation Jobs grouped by Final state (top) and rate of Simulation Jobs for the different Final states (bottom).

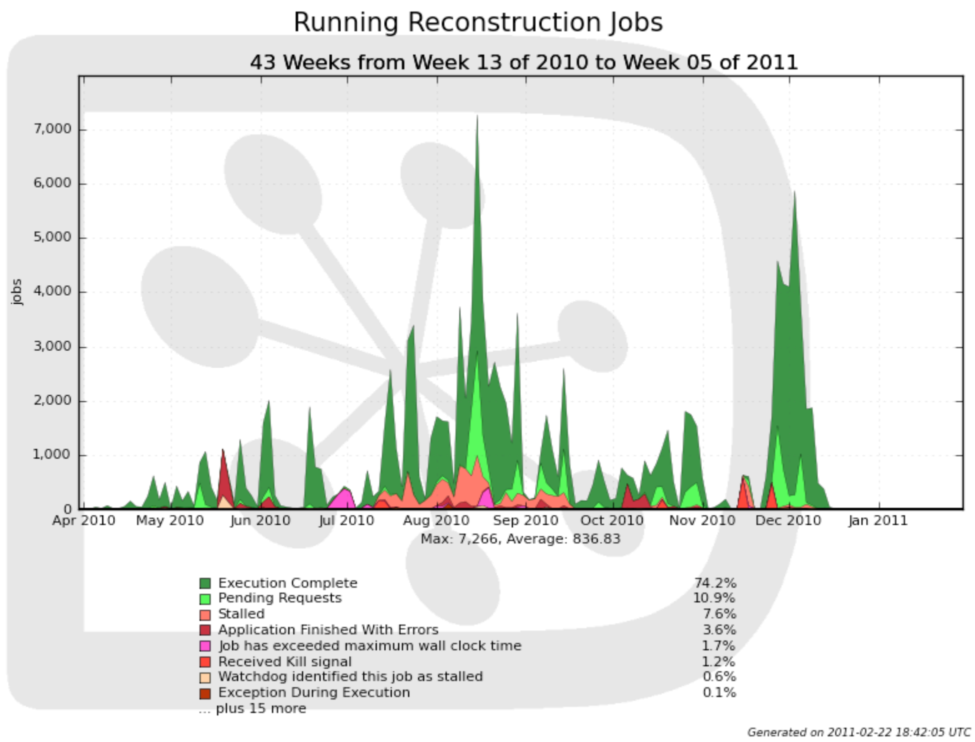


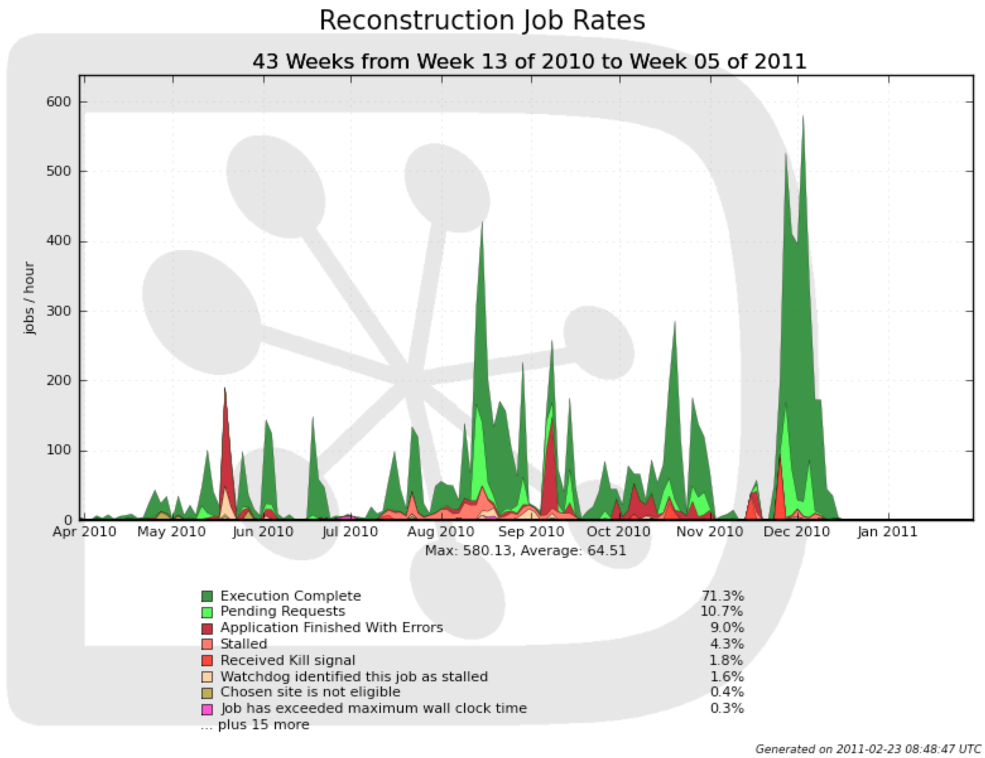


Figure 4-5: Number of simultaneously running Reconstruction Jobs grouped by Final state (top) and rate of Reconstruction Jobs for the different Final states (bottom).

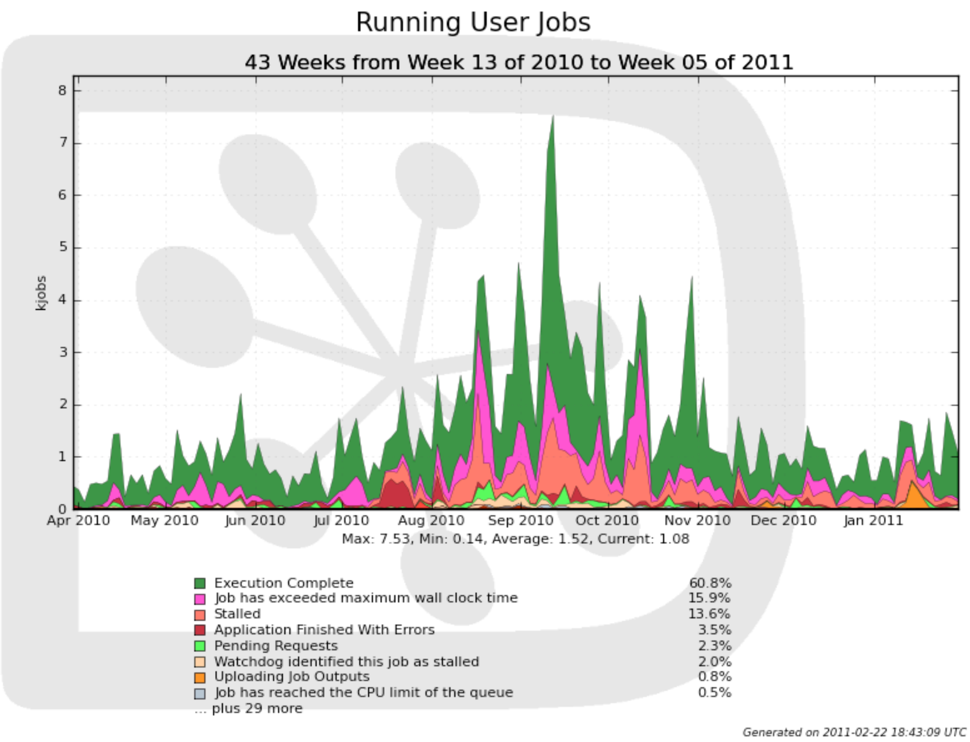


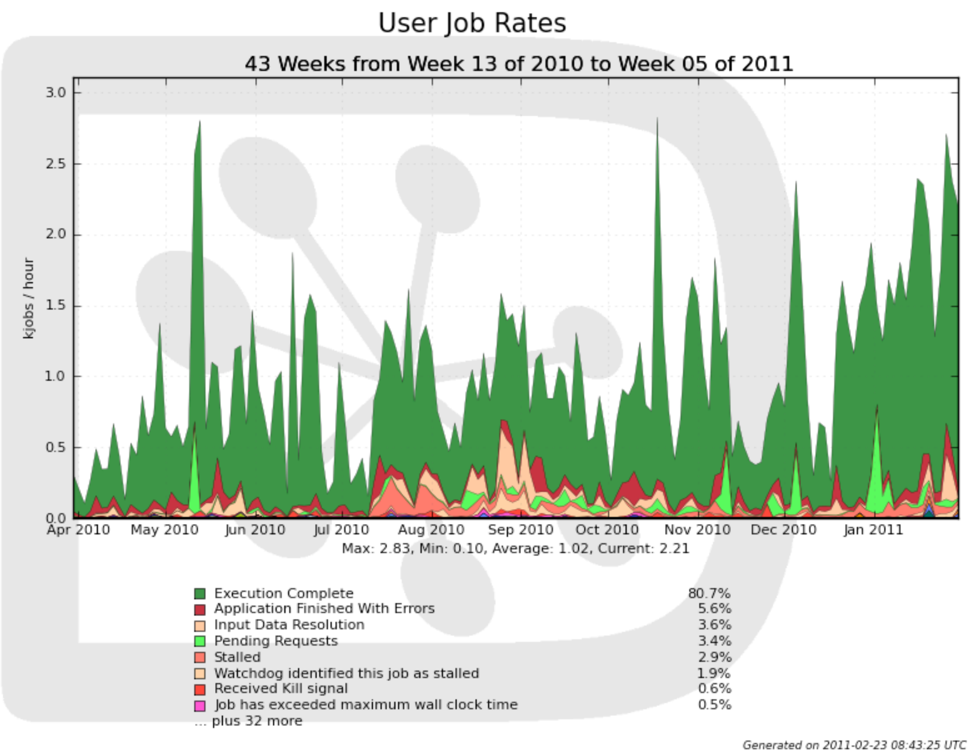


Figure 4-6: Number of simultaneously running User Jobs grouped by Final state (top) and rate of User Jobs for the different Final states (bottom).

# 5. Usage of Storage resources

The usage of storage resources was for 2010 mostly an accumulation of new data samples, since we started from a relatively empty situation. Soon we realized that the nominal model with 7 replicas of the reconstructed data samples would not fit in the available storage at the sites, because of the increased size of the events due to higher number of collisions per bunch crossing. Thus we revisited the model and decided on a compromise solution of 2 replicas on T1D1 plus another 2 replicas on T0D1.

In this document we report the status of the different SEs at the end of 2010 and at the beginning of February 2011 after a partial clean up of older versions in preparation for the 2011 data-taking. At the same time a large simulation activity started at the end of 2010 and was still on going in March 2011.

The situation of the Disk usage on December 20$^{th}$ 2010 is summarized in Table 5-1. Although there was still some space available, the high fragmentation of the storage due to explicit fractions being assigned at each site to real data, simulated data and user data (and a further split into T1D1 and T0D1 for the first two cases) has caused some SRM space tokens to become full despite there still being space available on other tokens, either at the same site or at other sites.

| **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 5-1: 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.

The overall usage went up to about 70%, but it is not uniform. Some discrepancies between the SLS and the LFC are observed and are currently being addressed.

For Tape the situation on December 20$^{th}$ 2010 is summarized in Table 5-2. SLS only reports on the usage of the disk cache in front of the Tape system and thus is irrelevant for the purpose of this document. One can see how overall the usage is about 50% of the total requested.

| **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 5-2: 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.

Table 5-3 summarizes the situation with Disk Storage on February 8$^{th}$ 2011. Two competing processes are taking place simultaneously. On the one hand old versions of real and simulated data are being removed and at the same time a large simulated data sample corresponding the running conditions of 2010 is being produced.

| **Disk Summary** (08/02/2011) | Pledge (TB) | Seen by SLS | | | Seen by LHCb | |
|---|---|---|---|---|---|---|
| | | TB | | | TB | |
| | | Total | Used | Avail. | Used | Pledge-Used |
| **FZK** | 495 | 500 | 330.2 | 169.8 | 336.3 | 158.7 |
| **IN2P3** | 610 | 638.5 | 298.7 | 338 | 284.4 | 325.6 |
| **CNAF** | 450 | 462.7 | 386 | 74.7 | 384.1 | 65.9 |
| **NL-T1** | 560 | 563 | 350.4 | 209.6 | 244.9 | 315.1 |
| **PIC** | 240 | 255 | 149.5 | 100.5 | 149.2 | 90.8 |
| **RAL** | 505 | 1174.9 | 453.5 | 183.9 | 443.9 | 505 |
| **Tier1s** | **2860** | **3594.1** | **2359.7** | **1224.5** | **1840.2** | **1019.8** |
| **CERN** | **1135** | **1175.6** | **861.2** | **314.4** | **736.8** | **398.2** |

Table 5-3: 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.

# 6. Summary

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 reasonable time. We are using more than peak power for the current MC10 simulation campaign, this will have to be smoothed in future. Integrated CPU work has been adequate.