CERN-RRB-2006-007
CERN-LHCC-2006-009
LHCC-G-109
27 November 2005

LARGE HADRON COLLIDER COMMITTEE

Review of Computing Technical Design Reports for LHCb

November 2005

Executive Summary

The LHC experiments and the LCG are assembling the computing infrastructure required to process the large amount of data expected in the years 2007-2010 and beyond. Each experiment has prepared a Computing Technical Design Report (TDR) for review by the LHCC. These Computing TDRs outline the requirements for the basic infrastructure that will be needed for LHC computing during the first few years of data taking. These computing models will most likely evolve over time as the detectors are commissioned and data analysis begins.

The review committee felt that the TDR documents from the experiments contained the conceptual designs of the computing systems. Technology is evolving rapidly, and while the overall goals are clearly defined, in many instances, choices of specific technology remain. In part, this is due to the pressure to delay decisions to be able to take advantage of technical advances and falling prices. In addition, the distributed computing model using the GRID is still undergoing intensive testing and commissioning.

The committee stresses that these computing models remain essentially untested. Some components were exercised during the data challenges and the recent service challenges. Other critical elements will be tested in the coming year. The lack of testing under battle conditions is particularly true for the distributed analysis portion of the models. The Tier-0 planning is the most advanced and the planned Tier-0 resource seemed well matched to the needs of the experiments.

The ongoing program of Service Challenges (SCs) forms the core of the conditioning tests for the Worldwide LHC Computing Grid (WLCG). SC-3 was underway at the time of the review. A list of the required WLCG baseline services and associated milestones has been established. Much work remains to be done to build and commission this global system of computers and have them operating 24 x 7. The committee applauds all the work that has been accomplished through the data challenges and the service challenges, but cautions that there are still very significant milestones to be met during the coming year.

The experiments have not yet fully determined their plans for alignment and calibration. The committee recommends that the experiments develop their calibration and alignment strategies as soon as possible - including the use of the CAF at CERN and the Tier-1 and Tier-2 centres. These plans should be developed in collaboration with the (W)LCG.

Large-scale data analysis tests are scheduled for the coming year that will use the WLCG baseline grid services. By September 2006, the major components of the computing systems will have been tested and the production system should be in operation. This is approximately the same timescale for initiating the large computing purchase required for the start of the LHC. The committee urges the computing management to proceed with caution in purchasing computing and to re-evaluate the resource planning regularly so that the computing resources are purchased only when needed. At the same time, we recognize the need to proceed with planned purchases for 2006 so that infrastructure is put in place and the large-scale system tests can advance.

A large quantity of common infrastructure software has been developed for the LHC. Much of this software development was done through external (non-CERN) funding. A plan is needed to support the computing and software infrastructure after the funding for EGEE and other GRID project funding comes to an end. This support plan should be included in the MoU process.

The computing resources at the Tier-1 and Tier-2 centres are identified through the MoU process. This process has already identified a large fraction of the required resources. It is important to keep in mind that there could still be large uncertainties in the resource requirement estimates. Moreover, the overall balance of resources amongst the experiments for resources outside of CERN seems difficult to achieve. There is an estimate of the pledged resources pledged for each experiment. At the time of the October review, ALICE had only identified about 50% of their required computing resources (CPU, disks, tapes). CMS was also lacking the required pledges to meet their resource needs. ATLAS and CMS have differing computing requirements for their Tier-1 sites. The committee finds that these differences are not fundamental but depend primarily on the details of their computing models.

The resource balance amongst the experiments is a major concern and must be resolved to ensure the physics output of all experiments. The current balancing scheme outlined in the MOU that involves descoping recommendation from the LHCC seems unworkable without some means for enforcement.

The review committee supports the first steps of the transition of the LCG organization towards management of the operation of a global LHC “computing centre”. The committee feels that the management of this global computing project still needs to be strengthened. Issues of balance of resources and global operations will require a strong organization and management team. We encourage a stronger connection between the computing planning process and the physics goals of the experiments. Therefore, we recommend that a “Computing Coordinator” be appointed to work together with CERN management, the LCG project and the four experiments.

The committee congratulates the collaborations and the LCG for their work presented in the Computing TDRs and recommends approval. More detailed comments and specific recommendations are outlined in the full document.

Introduction

Starting in 2007 the LHC is expected to produce proton-proton collisions at a center of mass energy of 14,000 GeV with an initial luminosity of approximately 2 x $10^{32}cm^{-2}s^{-1}$. The luminosity is expected to rise steadily until it reaches the design luminosity of $10^{34}cm^{-2}s^{-1}$ by 2010. It is expected that the trigger rates will not vary much with luminosity, which implies that the expected amount of data output rate will remain roughly constant.

Table 1 provides a comparison of the event size and trigger rate for the four experiments during p-p collisions. The event size, trigger rate and speed of the reconstruction and simulation codes are the fundamental parameters that are used to calculate the required production capacity. The capacity needed for analysis is more difficult to calculate and depends on the number of users, the number of analysis groups, and access patterns that are strongly analysis dependent. The AOD data will generally be available for analysis, but during commissioning access to the ESD and perhaps even the RAW data is anticipated.

p-p (HI)	**SIMU**	**Sim ESD**	**RAW**	**Trigger Rate**	**RAW Rate**	**RECO**	**AOD**	**TAG**
	MB	**MB**	**MB**	**Hz**	**MB/s**	**kB**	**kB**	**kB**
ALICE (HI)	300	2.1	12.5	100	1250	2500	250	10
ALICE (pp)	0.4	0.040	1.0	100	100	200	50	10
ATLAS (pp)	2	0.5	1.6	200	320	500	100	1
CMS (pp)	2	0.4	1.5	150	225	250	50	10
LHCb (pp)	--	0.4	0.025	2000	50	75	25	1

Table 1: The anticipated event size and raw data rates for the four experiments during proton-proton (pp) collisions at the LHC are presented in the Table. Also included are the estimates for Heavy Ion (HI) collisions at the ALICE experiment. SIMU: Simulated RAW data, RAW: data from the experiment to be recorded in mass storage, RECO (or ESD): output from the reconstruction code, AOD: reduced data format for analysis, TAG: summary for event selection.

There are large uncertainties on the estimates of processing times and how they scale with luminosity. These numbers will only be known when LHC data is available. Reprocessing of datasets primarily occurs at the Tier-1 centres. The experiments have estimated that this reprocessing could happen several times per year. It is anticipated that the number of passes be limited by the manpower available for validation and preparation of production software releases as well as by the overall compute capacity.

Planning for detector calibration and alignment is beginning for each of the experiments. In general, these schemes have not yet been worked out in detail and the impact on processing and reprocessing has not yet been fully evaluated.

The estimates for computing (cpu and disk) resources requirements for the LHC Computing system from 2007 to 2010 are presented in Figure 1 and Figure 2. These estimates include the requirements at CERN (Tier-0 and CAFs), the Tier-1 centres and the Tier-2 centres (or federations). The CPU resources are expected to grow from about 53 MSI2K in 2007 to over 330 MSI2k by 2010. At the moment 11 Tier-1 centres have been identified. The expected disk requirement for the beginning of LHC operations 2007 is about 17 PB. This expands to about 135 PB in 2010. Mass storage (tape) requirements grow from about 14 PB in 2007 to almost 140 PB in 2010.

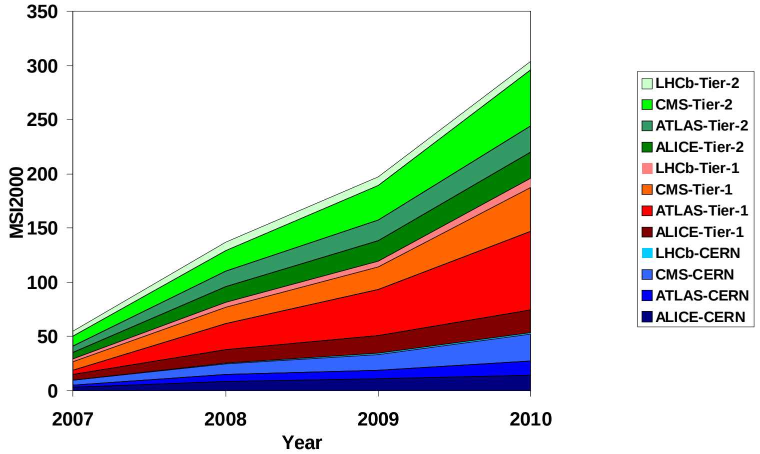

Figure 1 CPU requirements at the Tier-0, Tier-1 and Tier-2 centres for each of the four experiments. The time period is from 2007 to 2010.

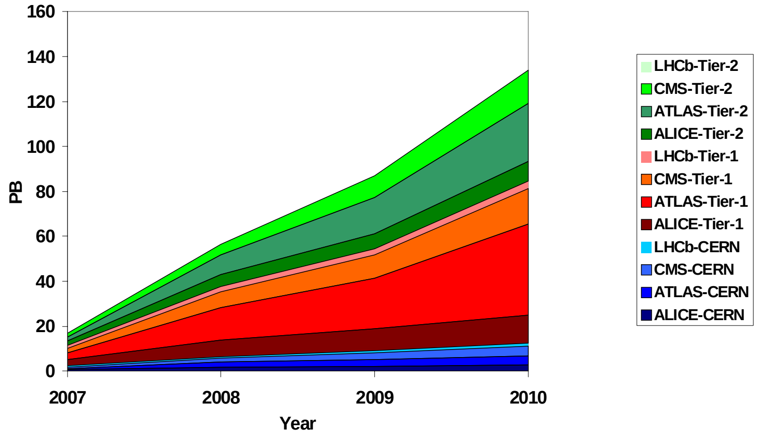

Figure 2 Disk requirements at the Tier-0, Tier-1 and Tier-2 centres for each of the four experiments. The time period is from 2007 to 2010. Note that LHCb plans to use the Tier-2 centres for MC production. Their analysis will primarily done at the Tier-1 centres.

For 2008 the total CPU requirements are 83 MSI2K and the disk requirements call for a total of 31 PB of disk at all centres. The exact amounts, of course, depend on the details of the LHC commissioning and run plan. More information on 2007- 2008 physics and calibration plans will be available during the coming year.

The Computing-Resource Review Board (C-RRB) has begun reviewing detailed MoUs regarding level of service and capacity at the Tier-1 and Tier-2 centres. The pledges collected at the time of the October meeting are presented in Table 2. While the overall level of the pledged resources for cpu, disk and tape at the centres was within about 20-30% at the time of this report, some experiments appear to have significant deficits in their identified resources at the Tier-1 and Tier-2 centres.

Tier-1 Planning for 2008		ALICE	ATLAS	CMS	LHCb	SUM 2008
CPU - MSI2K	**Offered**	6.7	22.7	12.5	4.4	46.3
	TDR Requirements	12.3	24.0	15.2	4.4	55.9
	Balance	**-46%**	**-5%**	**-18%**	**-0%**	**-17%**
Disk - PBytes	**Offered**	2.8	12.5	5.7	2.2	23.2
	TDR Requirements	7.4	14.4	7.0	2.4	31.2
	Balance	**-62%**	**-13%**	**-18%**	**-10%**	**-25%**
Tape - PBytes	**Offered**	3.2	9.1	8.1	1.9	22.3
	TDR Requirements	6.9	9.0	16.7	2.1	34.7
	Balance	**-54%**	**1%**	**-51%**	**-9%**	**-36%**

Includes current planning for all Tier-1 centres

Tier-2 Planning for 2008		ALICE	ATLAS	CMS	LHCb	SUM 2008
CPU - MSI2K	**Offered**	5.0	19.5	17.4	4.4	46.3
	TDR Requirements	14.4	19.9	19.3	7.7	61.3
	Balance	**-65%**	**-2%**	**-10%**	**-42%**	**-24%**
Disk - PBytes	**Offered**	1.4	5.9	4.5	0.8	12.6
	TDR Requirements	5.1	8.7	4.9	0.023	18.723
	Balance	**-72%**	**-33%**	**-8%**	**n/a**	**-33%**
# Tier-2 federations - included(expected)		12 (13)	20 (28)	17 (19)	11 (12)	28 (37)

Table 2 contains the Computing Resource Planning for 2008 from October 2005. This table was complied at the time of the October 2005 review and includes 28 Tier-2 centres or federations and all Tier-1 centres.

The Computing TDR Review Process

In early 2005 the LHCC reviewed the Computing Models of the four LHC experiments and released a report[1] containing an evaluation of the plans for data management and computing at the LHC. At that time, each of the four LHC experiments had produced estimates of their anticipated computing capacity requirements in terms of disks, tapes, CPUs and networks for the Tier-0, Tier-1 and Tier-2 centers. These numbers were documented in their Computing Model documents (ALICE[2], ATLAS[3], CMS[4], LHCb[5]) that were submitted to the LHCC in mid-December 2004 and reviewed in January 2005.

The review committee found that "*aside from issues of peak capacity, … the computing models presented are robust enough to handle the demands of LHC production computing during early running (through 2010.) There remains a concern about the validity of the data analysis components of the models.*" The committee did not consider costs in their evaluation, so there was not a comprehensive analysis of the scope of the computing requirements, except to say they should be sufficient to address the production requirements.

The four experiments and the LCG Collaboration submitted Technical Design Reports documents for the Computing Projects in June 2005 – ALICE[6], ATLAS[7], CMS[8], LHCb[9], LCG[10]. The LHCC received these materials and listened to a series of presentations from each of the four experiments and the LCG at the June 2005 LHCC meeting. The committee invited four external referees, one per experiment, to review the documents along with the LHCC members. In addition to the external referee, LHCC members helped to evaluate the Computing TDR of their "home" experiment. The LHCC LCG referees read and made comments on the TDR for the LCG.

The referees produced a set of questions concerning the individual TDRs. On October 7-8, 2005 the referees met together at CERN. During this two-day meeting, the leaders of the computing projects in the experiments and the LCG project were asked to supply answers to these sets of questions for the reviewers. All external referees and many LHCC members attended these sessions, which covered all five TDRs. The committee met in closed session on October 8 for discussions. This document contains a summary of these discussions and from the follow-up discussions at the October LHCC meeting. In addition to these general comments, specific comments and recommendations for each of the five TDRs are also given in this report.

The committee was impressed by the outline towards building global computing models that was outlined in the TDRs, but found that the documents contained conceptual designs or status reports rather than technical designs as in the case of the detector TDRs.

General Comments

The committee felt that the TDR documents from the experiments contained the conceptual designs of the computing systems. Technology is evolving rapidly, and while the overall goals are clear, in many instances, the technology choices have not yet been made. In part, this is due to the pressure to delay decisions to be able to take advantage of technical advances and falling prices. In addition, the distributed computing models using the GRID are still undergoing intensive testing and commissioning.

The committee stresses that the computing models remain essentially untested. Some of the components have been exercised during the experiments' data challenges and the recent service challenges. The agreement on the grid baseline service has been a good step toward global interoperability and much of this infrastructure will be tested within the upcoming year.

Testing under battle conditions – many users, chaotic access patterns - is a particular challenge before data taking begins. The distributed analysis portion of the computing models remains the least well tested. The outcome of these exercises could have an impact on Tier-1 and Tier-2 resources requirements.

Tier-0 planning is the most advanced and the resource requirements seemed well matched to the needs of the experiments.

The ongoing program of Service Challenges (SCs) forms the core of the conditioning and testing for the WLCG baseline services. SC-3 was underway though one month behind schedule at the time of the review. A list of required WLCG baseline services and associated milestones has been established and will be reviewed before SC-4. Much work remains to be done to build and commission this global system of computers and have them operating 24 x 7.

The committee applauds the work that has been done on the service challenges, but cautions that there are still significant milestones to be met in the coming year.

A CERN Analysis Facilities (CAF) is now foreseen for all experiments. The committee believes that these facilities are critical for commissioning and early calibration/alignment studies. In general, their usage should be defined better by the collaborations. The committee fears that the user community will migrate their analysis to these facilities rather than distribute their analysis to the Tier-1 or Tier-2 sites. This could create difficulties since the CERN and the CERN computing centre will not be able to house the computing power needed to accommodate this large user community.

The experiments still need to address their calibration and alignment schemes and the impact on computing resources and processing scenarios. The committee has several concerns: the CAF may be heavily loaded with analysis jobs from many users if not properly managed, and the Tier-2 centres may not be prepared to offer 24x7 coverage that may be required for these time critical tasks.

The committee recommends that the experiments develop their calibration and alignment strategies as soon as possible - including the use of the CAF at CERN and the Tier-1 and Tier-2 centres. These plans should be developed with close communication with the (W)LCG.

Large-scale data analysis tests are also scheduled for the coming year that will use the WLCG baseline grid services. By September 2006, the major components of the computing systems will have been tested and the production system should be in operation. It should be the right time to initiate large computing purchases that will put the required resources in place for the start of the LHC.

The committee urges the computing management to proceed with caution and to re-evaluate the resource planning regularly so that most computing resources are purchased only when needed. At the same time, we recognize the need to proceed with planned purchases for 2006 so that infrastructure is put in place and the large-scale system tests can advance.

A large body of software has been developed for the LHC by the LCG and by EGEE, OSG and other software collaborations and GRID projects. This software comes in many flavors. It could be middleware, an application, or an analysis tools. Many of these software products have become critical components of the infrastructure for the LHC computing systems.

A plan is needed to support the computing and software infrastructure after the funding for EGEE and other GRID project funding comes to an end. This support plan should be included in the MoU process.

Funding agencies and institutes from many parts of the globe are contributing to the LHC computing effort. There are 11 Tier-1 sites and more than 28 Tier-2 identified and many of these sites will provide services for more than one experiment. Overall, the promised resources seem reasonable.

The WLCG is a collaboration of institutes that are committed to providing LHC computing based on grid technologies. The true spirit of grid is “experiment blind”, and if this were indeed true, there would be little concern regarding balance of resources amongst the experiments. Institutions and funding agencies, however, often prefer to support local interests when assigning resources and it seems doubtful that this will change. According to the Computing MOU, it is proposed that the C-RRB and the Resource Scrutiny Group (RSG) determine the availability of global computing resources. If these available resources are not sufficient, the LHCC will then be consulted to recommend the proper balance.

Overall balance of resources amongst the experiments for resources outside of CERN will be difficult to achieve. An estimate of the missing resources for each experiment was been made. At the time of the October review, ALICE could not identify approximately 50% of their required computing resources (CPU, disks, tapes) in the Tier-1 centres.

ATLAS and CMS have similar resource requirements at the Tier-0, however they have significantly different computing needs at their Tier-1 sites. The committee finds that this imbalance is not fundamental – meaning it does not originate from detector or physics effects, but that the differences depend mainly on the details of the computing model. One should note that there could still be large uncertainties in these resource numbers.

The resource balance issues are a major concern and must be resolved to ensure the physics output of all experiments. The current balancing scheme outlined in the MOU that involves descoping recommendation from the LHCC seems unworkable without some means for enforcement.

The role and mission of the LCG must refocus as LHC operations approaches. LHC computing depends on the success of the GRID and funding agencies (and other scientists) from all over the world are monitoring the HEP commitment to GRID computing. The LCG must have a strong commitment to operating large scale distributed computing on the grid.

The LCG can no longer be seen as a development or a deployment project. The LCG needs to become an organization focused on the operation of a global service. The review committee supports the first steps of the transition of the LCG towards a global LHC “computing centre”. Daily management of WLCG service and operations remains a concern. It is not clear to us how conflicts will be managed and priorities will be established.

The committee feels that the management of this global computing project needs to be strengthened. Issues of balance of resources and global operations will require a strong organization and management team and this team has to operate on a global scale as well as be responsive to the physics goals of the experiments.

We encourage a stronger connection between the computing plans and to the physics goals of the experiments in the management. Therefore, we recommend that a "Computing Coordinator" be appointed to work together with CERN management, the LCG project and the four experiments

General Conclusions

The committee expresses its congratulations to the collaborations and the LCG for their Computing Technical Design reports. These documents form the foundation for the data production and physics analysis facilities and infrastructure at the LHC. While many of the concepts have been tested and validated, much remains to be done put achieve a distributed analysis structure based on GRID tools that can be used by a large number of LHC physicists. We advise a follow-up review of LHC Computing in Fall 2006. This could be done at the time of the LCG Comprehensive Review.

The committee finds that the Computing TDRs provide the necessary foundation for planning for data production and analysis. At this time, however, these models remain essentially untested. There remains a particular concern about the validity of the data analysis components of the models and recommends a follow-up review in Fall 2006.

Comments and recommendations for LHCb

The LHCb computing TDR describes both the LHCb software and the processing steps required through to physics results including the resulting computer resource requirements.

The committee was very impressed with the level of maturity of LHCb software as presented in the TDR and updated in the presentation. All major software components already exist and have undergone review and where necessary major reimplementation and optimisation is well underway.

The manpower dedicated to LHCb core computing is 32 FTE in 2006. This is sufficient to cover all high priority tasks, although there is a lack of software engineers for support tasks such as web master, documentation and integrity checking tools.

LHCb software developer's have made excellent collaboration with other experiments. Of particular note is the Gaudi framework that has been adopted by ATLAS and others, the high level analysis toolkit LoKi which developed out of ideas from ARGUS and Hera-b and the GANGA Grid user interface in collaboration with ATLAS.

LHCb has been successful in early exploitation of the Grid by pragmatic interfacing of its own DIRAC system. Collaboration with LCG is good and DIRAC evolves to use the latest grid software components as they become available. Preparation for user analysis on the grid is advanced with Ganga in first major rewrite after 3 years development.

The computing model of LHCb foresees CERN will cover 25% of the analysis needs and 75% will be performed at the 6 other Tier-1's. The Tier-1 centres are used for reconstruction, stripping and analysis of real and Monte-Carlo data, while Tier-2 centres perform Monte-Carlo production only and have no mass storage requirements. Fresh raw data will be distributed in quasi real-time via the Tier-0 to the Tier-1 sites for initial reconstruction with a maximum delay of a few days. This requires very high reliability of the network links to and high availability of the Tier-1 centres.

In 2008 complete analysis datasets will be available at each Tier-1, but in 2009 analyses of 2008 data outside of CERN will require the grid to fully support many user multi-Tier1 jobs. This has yet to be demonstrated.

The computing model and resulting resource requirement estimates appear justified and realistic. LHCb is making every effort to achieve the estimated processing times and event sizes as of the TDR.

Improvements in execution time of stripping from that estimated in the TDR to that actually obtained at time of the TDR have been reached by a more intelligent sequencing of algorithms and the RICH reconstruction speed has been significantly improved. The reconstruction code rewrite just mentioned in the TDR is already well advanced. The pattern recognition packages have been adapted to the new event model resulting in higher efficiency and faster code. The timing estimates of the TDR have essentially now been reached although further improvements may be required to allow for possible additional factors such as real chamber occupancies.

Currently the RAW size is larger than the target (35kB cf 25 kB) and the rDST is significantly smaller (8kB cf 25kB); however, LHCb consider the targets realistic as contingency is required for ghost tracks and further developments. These numbers may be revised at the time of DC'06. Raw data transfer out of CERN, Tier-1 reconstruction and stripping are scheduled for the milestone DC'06 in May 2006.

The current software largely assumes a perfectly aligned and calibrated detector. Estimates have been made of the effect of misalignments on tracking efficiency and resolution. An example of sub-detector alignment is available for the VELO and a workshop is scheduled to ensure calibration and alignment procedures are in place on time. Efforts are underway to develop strategies to achieve ultimate calorimeter resolution.

We note that there is some flexibility in the computing model. The HLT trigger computers are foreseen to be used for reprocessing outside the lumi period, but this could alternatively be done at Tier-1s should compute time be available, decreasing also network load. Four strippings and two reconstruction passes per year are scheduled and are likely to be necessary in the initial phase but could possibly be reduced if resources are tight or in the light of experience from the first years of data taking.

Mass storage could be reduced by recycling tapes containing processed data from previous stripping passes somewhat earlier than currently foreseen (from 2010). The rule of thumb applied, with experience from LEP, HERA etc., is that a particular dataset may be required for 3 years to allow students to complete and submit a PhD thesis. However, old dataset versions that are no longer disk resident may in practice be difficult to use.

The LHCb computing model places data in "long-term" and fast access categories, and these were mapped to MSS and disk requirements respectively. As this implies, most tape use is concerned with archiving. All frequently accessed data is foreseen to be disk bound with massive tape reading only by the centrally organized stripping passes. For Monte-Carlo data the equivalent of one-year’s production only is foreseen to be disk resident as the Monte Carlo samples will be generated for very specific studies and it is assumed the majority of these will be over within 12 months and of lesser interest in subsequent years. The simulation dominates however the CPU needs of LHCb, and the specified estimate is a best guess only, which will have to be revised when the specific studies which can be tackled by simulation can be specified.

The disk requirement estimate given in the TDR results in an increasingly larger fraction of disk at CERN compared to Tier-1s in later years. It may be prudent to distribute the disk more evenly according to the envisaged CPU fractions, to avoid that too many jobs are run at CERN. Transparent job splitting will be provided in the LHCb Grid environment to facilitate multi-Tier1 jobs.

No interactive parallel grid based analysis is currently foreseen (e.g. via PROOF). It is anticipated that final batch Grid analysis steps will involve writing datasets of only a few thousand events suitable for analysis on a laptop. A second analysis activity is the generation and fitting of toy MC samples with minimal input and has no obvious interactive requirement. This may be reviewed in 2006 in light of grid analysis tests.

The total LHCb computing requirements are small compared to the total resources required on the LHC grid. While LHCb can hope to benefit from fluctuations in the actual daily usage of the larger experiments it must be ensured that LCG can guarantee LHCb network bandwidth and cpu allocation even at times of high overall load.

LHCb sees CERN acting as the Tier-0 centre and taking on the same function as any other of their Tier-1 centres i.e. support reconstruction and analysis. It is envisaged 25% of analysis requirements will be carried out at CERN given the access to the full RAW data and the availability of the "hot stream." Approximately 30% of the CPU requirements at CERN in 2008 is for user analysis. LHCb makes no separation between the roles of "CAF" and that of Tier-0+1 centre.

LHCb remaining computing milestones

November 2005: Analysis at all Tier-1's
December 2005: End of review of reconstruction software
March 2006: Final alignment strategy
May 2006: Start data processing phase of DC'06
October 2006: Alignment and calibration challenge
January 2007: Permanent Monte Carlo production mode ready
April 2007: Production system and software ready for data taking

Appendix A

Committee Representation for October 7-8 review

Chair: P. McBride
Representatives from the LHCC: K. Borras, F. Forti, S. de Jong, M.Martinez-Perez, V. Kekelidze, B. Peyaud
External:
D. Boutigny (CC-IN2P3-Lyon; for ALICE C-TDR),
T. Haas (DESY; for ATLAS C-TDR),
C. Bozzi (INFN Ferrara; for CMS C-TDR),
Alan Campbell(DESY; for LHCb C-TDR)
LHCC Chairman and Secretary: S. Bertolucci, E. Tsesmelis

Also present at the review
PH Department: J.-J. Blaising, D. Schlatter
IT Department: J. Knobloch

---

[1] Review of Computing Resources for the LHC Experiments CERN/LHCC/2005-006 http://committees.web.cern.ch/Committees/LHCC/lhcc-2005-006.pdf
[2] ALICE Computing Model CERN-LHCC-2004-038/G-086, draft:05-Jan-05, updated 04-Feb-05
[3] ATLAS Computing Model CERN-LHCC-2004-037/G-085
[4] CMS Computing Model CERN-LHCC-2004-035/G-083
[5] LHCb Computing Model CERN-LHCC-2004-036/G-084 (CERN-LHCb-2004-119)
[6] ALICE Computing Technical Design Report CERN-LHCC-2005-018, ALICE-TDR-012
[7] ATLAS Computing Technical Design Report CERN-LHCC-2005-022, ATLAS-TDR-017
[8] CMS Computing Technical Design Report CERN-LHCC-2005-023, CMS-TDR-007
[9] LHCb Computing Technical Design Report CERN-LHCC-2005-019, LHCb-TDR-011
[10] LCG Computing Technical Design Report CERN-LHCC-2005-024, LCG-TDR-001