

LARGE HADRON COLLIDER COMMITTEE

Review of Computing Technical Design Reports for ALICE

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, and 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 $2 \times 10^{32} \text{cm}^{-2} \text{s}^{-1}$. The luminosity is expected to rise steadily until it reaches the design luminosity of $10^{34} \text{cm}^{-2} \text{s}^{-1}$ 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 be 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 compiled 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 reportⁱ 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 (ALICEⁱⁱ, ATLASⁱⁱⁱ, CMS^{iv}, LHCb^v) 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^{vi}, ATLAS^{vii}, CMS^{viii}, LHCb^{ix}, LCG^x. 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 descopeing 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 ALICE:

As a general comment, the committee was pleased to notice that the ALICE collaboration has made important progresses in implementing software compatible with, or reusing standard LCG components. While this effort toward standardization should be pursued wherever possible, the committee feel that the ALICE computing is now close to a good balance between the implementation of common LCG software and the necessary experiment specific customization.

Most of the 2005 milestones have been met. The milestones which have been delayed are those related to the Service Challenge 3 (SC3), the delay allowed the ALICE collaboration to improve the compatibility and integration with LCG. The combined test with SC3 is ready to go and the distributed analysis will be demonstrated at the Super Computing 05 conference. The committee is confident that all the remaining 2005 milestones will be met.

The ALICE analysis will rely on several components from which two look particularly important.

1. Events of interest to the physicists will be located using a grid collector based on compressed bitmap index technology and ROOT I/O. This piece of software has been already developed and tested by the STAR collaboration and seems to be efficient and scalable.
2. Analysis will be performed using the Parallel ROOT Facility (PROOF). The simplest approach consists in using PROOF in a single site. A more complex possibility is to run analysis on a set of PROOF systems distributed over multiple sites.

Both components look very promising and seem to provide a good basis for an efficient grid based analysis. The committee recommends to test both of them in a realistic environment in order to verify their functionalities and scalability in the ALICE framework. A full scale test of a PROOF system would be especially useful to specify the hardware and related architecture to be deployed in the CERN and remote Analysis Facilities.

Historically ALICE has always been in favor of the so-called “cloud model” implementation of the GRID, where sites contribute on an equal basis. The LCG approach is radically different and based on the hierarchical “MONARC” model. Following recommendations from previous computing reviews, ALICE is now well engaged on an implementation of the now standard MONARC model.

Due to the tremendous complexity of Heavy Ion (HI) events, it will not be possible to reconstruct these data “online” at the Tier-0 and a 4-month delay is expected. ALICE is not willing to push some of these data outside CERN as this would impair ongoing physics analyses. The committee is worried about the consequences of this delay and recommends to implement a strategy in order to make sure that important physics results and feedback on the detector performances will be available to the collaboration soon after data taking.

In 2007 ALICE is expecting 40% pp and 20% HI of a standard year. However the event rate will be kept at nominal by loosening the triggers. ALICE states that the full

rate of looser triggers is needed to allow the discovery physics to be done, making use of the cleaner event sample with low pile-up occurring during the initial low luminosity run. The committee recommends that this question be further considered and that a careful assessment of the balance between physics benefits and computing cost be done. A smaller event sample would allow to delay some computing equipment buying and then to save money.

ALICE has been developing its own set of grid services called ALIEN. The LCG grid implementation is now progressing and is offering more and more common services. ALICE is replacing ALIEN components by LCG ones everywhere it is possible. ALIEN is now restricted to an interface for ALICE users to the heterogeneous and fast evolving grid environment. The advantage of such an interface is that it offers a simple interoperability layer to access all present and future grid flavors. The committee feels that ALIEN is now a mature software package and that the redundancies with LCG have been removed.

The manpower dedicated to the ALICE core computing is 26.6 FTE in 2005 and is expected to decrease in the next years, down to 14.5 FTE in 2009-2010. ALICE is justifying this relatively limited amount of manpower by designing a system with a high level of automation in the production tasks. The committee feels that this choice may be risky as initial difficulties in the grid reliability may lead to inefficiencies that would seriously impair any automation.

ALICE remaining computing milestones

MS4-July 2005: PDC05 Start of combined test with SC3 (phase 2). **Delayed to December 05**

MS5-September 2005: PDC05 Start of distributed analysis (phase 3). Phase 3 of PDC05 (see Chapter 3 of computing TDR) is combined with the services test of LCG SC3. It includes the AliEn services and gShell, the ALICE user interface to the AliEn services. **Delayed to January 2006**

MS6-September 2005: Metadata prototype ready.

MS7-December 2005: Condition infrastructure deployed.

MS8-December 2005: Preliminary implementation of algorithms for alignment and calibration ready for all detectors.

MS9-January 2006: Release of AliRoot framework in preparation of the PDC06. Final prototype of alignment and calibration ready to be tested. Alignment and calibration algorithms prepared for all detectors. Global alignment and inter-calibration of the detectors at the prototype stage.

MS10-January 2006: Start of PDC06. Test of the full production chain, including calibration and alignment. Size approximately equal to 20% of the real data of a standard data-taking year. Distributed batch analysis of the data. Evaluation of the distributed interactive analysis framework.

MS11-June 2006: End of PDC06. Evaluation of the data challenge results. Planning of most urgent activities for next year.

MS12-June 2006: Final implementation of algorithms for alignment and calibration ready for all detectors. This includes inter-calibration and alignment algorithms.

MS13-December 2006: ALICE computing environment ready for data taking.

Appendix A

Committee Representation for October 7-8 review

Chair: P. McBride

Representatives from the LHCC: K. Borrás, F. Forti, S. de Jong, M. Martínez-Pérez, 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

ⁱ Review of Computing Resources for the LHC Experiments CERN/LHCC/2005-006

<http://committees.web.cern.ch/Committees/LHCC/lhcc-2005-006.pdf>

ⁱ ALICE Computing Model CERN-LHCC-2004-038/G-086, draft: 05-Jan-05, updated 04-Feb-05

ⁱ The ATLAS Computing Model CERN-LHCC-2004-037/G-085

ⁱ The CMS Computing Model CERN-LHCC-2004-035/G-083

ⁱ LHCb Computing Model CERN-LHCC-2004-036/G-084 (CERN-LHCb-2004-119)

ⁱⁱ ALICE Technical Design Report of the Computing CERN-LHCC-2005-018, ALICE-TDR-012

ⁱⁱⁱ The ATLAS Computing Technical Design Report CERN-LHCC-2005-022, ATLAS-TDR-017

^{iv} CMS The Computing Project Technical Design Report CERN-LHCC-2005-023, CMS-TDR-007

^v LHCb Computing Model CERN-LHCC-2004-036/G-084 (CERN-LHCb-2004-119)

^{vi} ALICE Technical Design Report of the Computing CERN-LHCC-2005-018, ALICE-TDR-012

^{vii} The ATLAS Computing Technical Design Report CERN-LHCC-2005-022, ATLAS-TDR-017

^{viii} CMS The Computing Project Technical Design Report CERN-LHCC-2005-023, CMS-TDR-007

^{ix} LHCb Computing Model CERN-LHCC-2004-036/G-084 (CERN-LHCb-2004-119)

^x LHCb Computing Model CERN-LHCC-2004-036/G-084 (CERN-LHCb-2004-119)