12th March 2001
Original: English

ORGANISATION EUROPÉENNE POUR LA RECHERCHE NUCLÉAIRE
CERN EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH

BUILDING THE LHC COMPUTING ENVIRONMENT

Executive Summary

The LHC computing needs have been reviewed[1] last year with a view to identifying the requirements for full exploitation of LHC data by all participating institutes and physicists. In the attached paper, the Management summarises the findings of the review and their first reactions.

The size and complexity of this task requires a worldwide distributed computing system to which all institutes and funding agencies participating in the LHC experimental programme are asked to contribute.

The demands for the remaining developments and the installations required at CERN cannot be satisfied within the present budget. The Management therefore proposes a two-phased approach to the problem, covering the years 2001 to 2007:

?? Development and prototyping from 2001 to 2004, requiring mostly expert manpower, and the establishment of a production prototype;

?? Installation of the full initial production system in the years 2005 to 2007, requiring continued manpower efforts and substantial material resources.

The associated resource requirements are summarised in the table below:

Summary - Additional Resources needed	*R&D Phase (Phase 1)*				*First Production System (Phase 2)*			*Maint-enance*	*Total R&D 2001-04 (Phase 1)*	*Total First System 2005-07 (Phase 2)*
year	*2001*	*2002*	*2003*	*2004*	*2005*	*2006*	*2007*	*2008*		
Services required at CERN										
Additional personnel (person-years)	**16**	**41**	**42**	**50**	**50**	**50**	**46**	**21**		
Cost if employed as CERN staff (MCHF)	2.4	6.2	6.3	7.5	7.5	7.5	6.9	3.2	*22.4*	*21.9*
Additional materials (MCHF)	2.1	6.6	10.1	10.7	30	33.4	32.4	22.6	*29.5*	*95.8*
Service funding required at CERN (MCHF)	**4.5**	**12.8**	**16.4**	**18.2**	**37.5**	**40.9**	**39.3**	**25.8**	*51.9*	*117.7*
Core software for LHC (in IT Division) - see note										
Additional s/w professionals (person-years	**6**	**6**	**6**	**6**	**6**	**6**				
Cost if employed as CERN staff (MCHF)	**0.9**	**0.9**	**0.9**	**0.9**	**0.9**	**0.9**				

Note: It is assumed that the core software effort will be provided for by the collaborating institutes and their funding agencies

Summary of additional resources needed

Member States are asked now to contribute to the first phase of funding of LHC computing as part of the base programme lying beyond the present budget limits.

As the work of the R&D phase proceeds, the estimates for the construction of the first production system will become more precise and Member States will be asked to contribute to this second phase in due course.

[1] CERN / LHCC / 2001-004 – Report of the Steering Group of the LHC Computing Review

Building the LHC Computing Environment

1. The LHC Computing Review

The LHC Computing Review was initiated in 1999 and mandated to report the status of computing preparations of the LHC experiments and IT Division and to make recommendations. It was organised in 3 panels and a steering group. The panels and the steering group represented high level users of computing services, managers of computing efforts and those responsible for all aspects of LHC computing (listed in Annex 1). Each panel reported to the steering group, which in turn prepared the final review report.

The final review report was the result of a large collaborative effort. Its contents are the responsibility of the Steering Group and represent a broad consensus of all parties involved.

The outcome and recommendations of the computing review may be summarised as follows:

1.1 The LHC computing model

The main parameters of the LHC computing requirements as estimated by the review are listed in Annex 2 of this paper.

After critical assessment, the review accepted the scale of the resource requirements as submitted by the four experiments.

A *multi-Tier hierarchical model* similar to that developed by the MONARC project should be the key element of the LHC computing model.

Computing Fabrics containing tens of thousands of components (processors, disks, tape drives and network switches) and optimised for data intensive operations will be at the heart of the centres of this multi-Tier hierarchical model. A particularly large Computing Fabric installed at CERN will serve as a common Tier 0 centre for all experiments and handle the raw data reconstruction and storage.

Grid Technology will be used to interconnect the distributed computing fabric in an attempt to contribute solutions to this model that provide a combination of efficient resource utilisation and rapid turnaround time.

In order to provide the required *bandwidth* of the wide area network between nodes in the distributed computing fabric, it will be vital to ensure a good, high-performance Research Networking infrastructure is in place.

1.2 Software

Joint efforts and *common projects* between the experiments and CERN/IT are recommended to minimise costs and risks in building the software.

Data Challenges of increasing size and complexity must be performed as planned by all the experiments until LHC start-up. These are needed in order to subject the steadily developing prototype computing environment to progressively more realistic tests.

CERN should sponsor a coherent programme to ease the transition of the bulk of the physics community from Fortran to Object Oriented (OO) programming.

Further identified areas of concern are the limited maturity of current planning and resource estimates for the production of the software, the development and support of simulation packages and the support and future evolution of analysis tools.

1.3 Management and Resources

Current cost estimates are based on the forecast evolution of price and performance of computer hardware.

On this basis, the *hardware costs* for the initial set-up of the LHC distributed computing centres, including Tier0 to Tier2 structures for all experiments, are currently estimated to be *240 MCHF*. The CERN-based common Tier0+Tier1 centre is estimated to make up about 1/3 of the overall computing capacity.

The total investment for the initial system is due to be spent – in approximately equal portions – in the years 2005, 2006 and 2007, assuming LHC starts up in 2006 and then runs reliably for the planned periods.

The *Core Software teams* of all four experiments are currently *seriously understaffed*. The additional human resources for the development of the core software of the collaborations are summarised in Table 1[2], as the difference between the requirement and the number of people available in 2000. About 40 additional software professionals are required during the development phase between now and 2005.

---

[2] The data is taken from Table A3.12 of the report of the LHC Computing Review.

Table 1 - Resources required for the core software of the collaborations

Required human resources to write the Core Software	available	---------- resources required (person-years) ----------					
year	*2000*	*2000*	*2001*	*2002*	*2003*	*2004*	*2005*
collaboration							
ALICE	12.0	17.0	17.5	16.5	17.0	17.5	16.5
ATLAS	23.0	31.0	36.0	35.0	30.0	28.0	29.0
CMS	15.0	25.0	27.0	31.0	33.0	33.0	33.0
LHCb	14.0	19.0	25.0	24.0	23.0	22.0	21.0
Totals	**64.0**	**92.0**	**105.5**	**106.5**	**103**	**100.5**	**99.5**
Required additional resources		**28.0**	**41.5**	**42.5**	**39.0**	**36.5**	**35.5**

It is assumed that the collaborating institutes and their funding agencies will provide for these persons. Such a shortfall was already noted in the report to the Resources Review Board (RRB) meeting of October 2000. As argued in section 4.1 below, some of these persons will have to be located at CERN under the guidance of IT Division.

The collaborations are working under the assumption that their physicists will take care of the other software aspects of the experiments such as preparing for the physics analysis.

The *staffing level of CERN/IT* as envisaged under the current CERN-wide staff reduction plan is incompatible with an efficient running of the CERN-based LHC computing system and software support.

The construction of a *common prototype* of the distributed computing system should be launched urgently as a joint project of the four experiments and CERN / IT, along with the major Regional Centres.

Software and hardware should be periodically integrated into increasingly realistic versions of a prototype of the computing environment and *Data Challenges* of increasing size and complexity must be performed. These Data Challenges are already foreseen in the preliminary planning of the experiments.

An agreement should be reached amongst the partners in this project, in which the construction, the cost sharing, the goals and the technical solutions of the common prototype are laid down.

The resource estimates used represent today's best knowledge of the conditions expected at the time that LHC will run. These estimates carry significant *intrinsic uncertainties;* they must be continuously monitored and updated according to actual knowledge.

Initial funding is needed urgently, in order to provide resources for prototype developments. The serious shortfall of currently available human resources for software development and support of the computing infrastructure is a major concern. At the same time, it should be noted that the additional development staff need not necessarily be permanent.

The approach to *Maintenance and Operation* of the LHC computing system includes the strategy of replacing approximately one third of the equipment annually, within a constant budget.

All those concerned are asked to undertake the necessary steps to ensure timely availability of sufficient computing resources to permit the LHC to realise its enormous potential in terms of physics outcome.

1.4 Specific Proposals

An *LHC Software and Computing Steering Committee* (*SC2*), composed of highest level software and computing management in experiments, CERN/IT and Regional Tier1 Centres, must be established to oversee the deployment of the entire LHC hierarchical computing structure and operate within the overall review process of the experiments.

Under the auspices of this committee, *Technical Assessment Groups* must be established to prepare common activities amongst experiments and Regional centres.

Technical Software Projects to execute common activities identified by the Technical Assessment Groups should be adequately resourced and managed under a dedicated Project Leader.

Each collaboration must prepare, on a common pattern, a *Memorandum of Understanding* (MoU) for LHC computing.

2. First Assessments of the CERN Management

The report summarised above describes the model adopted by the collaborations for collecting, storing and analysing the LHC data. It provides detailed estimates of the human and material resources that will be required for the development of the applications and the construction of the computing environment at CERN and in Regional Centres distributed throughout the world. The report also describes the R&D that will be necessary to build the software and tools required to manage and operate the distributed environment.

The report will now be discussed in the CERN scientific committees (LHCC, Research Board), in the divisions and experiments concerned, with the funding agencies participating in the LHC experiments (RRBs) and with the CERN Committees to arrive at a phased plan of providing the LHC computing environment.

The current status of the planning for the LHC computing environment may be summarised as follows:

- The requirements are clear, within the margin of uncertainties inevitable at this stage, and agreed by all of the collaborations.

- The computing model that has been developed is a worldwide, distributed facility with major potential for technology transfer to other sciences and industry. A number of European particle physics institutes have already begun to collaborate with other scientific disciplines on the European Commission DataGRID project[3], to develop and demonstrate some of the novel software, or *middleware*, that is required for LHC. This project, led by CERN, is collaborating closely with parallel complementary activities in several Member States and non-Member States.

- The implementation of the model will include large computing facilities at CERN and in national and regional centres, some of which have already been identified and funded to some degree.

- A major programme must be started now to -

  Carry out the R&D to develop the necessary management and organisational software required.

  Build a production quality prototype on a scale sufficiently large to validate the model and the technologies involved. This requires a prototype about 50% of the size of the final system needed by one experiment, and the work must be completed in 2004, prior to commencing acquisitions for of the main facility in 2005-7.

- Resources must be identified within the LHC collaborations to bring the teams working on the development of the core application software up to the required strength.

With this background, CERN-IT Division has prepared a preliminary plan and detailed estimates of the resources required to develop and build the part of the computing environment at CERN and provide support for application software development. When finalised, the plan will describe the development phase, the prototype, the initial installation, the further developments of the CERN LHC computing facilities and the co-ordination tasks with the outside computing centres. It will have the characteristics of a phased development activity within the rapidly developing field of computing, giving precise milestones and activities for the immediately following year(s) and becoming progressively less certain for the more distant future.

An approach with distinct execution phases is proposed to provide a framework for contributions to be made by the Member States.

3. Proposed Scope

Two phases can be distinguished:

---

[3] See http://www.eu-datagrid.org

1) Phase 1

The immediate requirement is the implementation of the development and prototype phase, to be completed in 2004. The goals are:

a) Prototype construction

Build a prototype of the LHC computing facility, using simulated LHC data and also real data and requirements of currently running experiments:

- Develop the system software and middleware required to manage a very large-scale computing fabric located at a single location.
- Develop the grid middleware to organise the interaction between such computing fabrics installed at geographically remote locations, creating a single coherent computing environment.
- Acquire experience with high speed wide-area network and data management technologies, developing appropriate tools to achieve required levels of performance and reliability for migration, replication and caching of large data collections between computing fabrics.
- Adapt LHC applications to exploit this environment.
- Progressively deploy a half-scale prototype of the LHC computing facility for a single experiment at CERN and in a number of future Tier 1 and Tier 2 centres, demonstrating the required functionality, usability, performance and production-quality reliability.
- Define the characteristics of the initial full production facility.

b) Core Software

Complete the development of the first versions of the physics application software and enable it for the distributed computing model (grids):

- Develop and support common libraries, tools and frameworks to support the development of the application software, particularly in the areas of simulation and analysis.
- In parallel with this, the LHC collaborations must develop and deploy the first versions of their core software.

This work will be carried out in a close collaboration between CERN, Tier 1 and Tier 2 centres and the LHC Collaborations.

2 Phase 2: Deploy the initial full production facility

Construct the initial full production facility (2005-2007) according to the experience gained in the years of prototyping.

For this construction phase, the resources required are still not known to a sufficient degree of precision but will be defined as part of the Phase 1 activity.

It is assumed that national funding agencies in the Member States and non-Member States will ensure the construction of the prototype Tier 1 and Tier 2 centres and their share of the distributed computing infrastructure. It is further assumed that the experimental collaborations will ensure the software requirements. For some aspects of the core software CERN will ask for resources to perform such tasks under CERN guidance.

It is of course natural to foresee that, as for past generations of experiments, the worldwide production facilities will have to be upgraded regularly according to the needs of the experimental programme.

4. Resources required at CERN

4.1 Human Resources

The human resources required at CERN for the provision of computing services are summarised in Table 2. This covers all of the services of the Information Technology Division. The requirements for physics support are shown separately for:

Baseline services – Management and operation of the basic services for physics, including the program development environment for LHC experiments and the residual support for non-LHC experiments.

LHC services & software operation – Management and operation of the computing facilities used for the storage, simulation and analysis of data for LHC experiments; management of the specialised high performance wide area network for physics data; provision of common software libraries and tools. Note that a contractor assures the basic operation of the computing fabric, the cost of which is included in the materials estimates.

Fabric & Grid R&D – The development of the systems software required for the LHC computing facility; the development of the prototype system.

Software & applications R&D – The development of common software libraries, tools and frameworks.

Engineering & accelerator services include the implementation of the network for the LHC machine. This activity was transferred to IT Division from the Accelerator Sector at the end of 2000, and was therefore not included in the tables prepared by the LHC Computing Review. From 2007 it is planned to consolidate support for the accelerator network with the operation of the campus local area network, included in the *Infrastructure* line.

Infrastructure & base computing services include the provision of support for desktop computing, the on-site network, the basic wide-area network services, electronic mail, Web services, databases, and general user support. Note that a contractor provides a large part of the basic operation of the desktop computing services.

Table 2 - Summary of the human resources required for computing services at CERN

Human Resource Requirements for Computing Services at CERN

all figures person-years, unless otherwise stated			*R&D Phase*				*First Production System*			*Maint-enance*
year	***1999***	***2000***	***2001***	***2002***	***2003***	***2004***	***2005***	***2006***	***2007***	***2008***
Physics										
Baseline physics services (inc. non-LHC support)	*10*	*10*	*10*	*10*	*10*	*10*	*10*	*10*	*10*	*10*
LHC services & software - operation	*19*	*14*	*16*	*16*	*16*	*25*	*30*	*30*	*35*	*25*
Fabric & Grid management software R&D	*1*	*4*	*20*	*32*	*32*	*18*	*10*	*5*	*4*	*3*
Software & applications R&D	*14*	*14*	*17*	*26*	*27*	*21*	*14*	*12*	*11*	*8*
Total physics services	*44*	*42*	63	84	85	74	64	57	60	46
Engineering & accelerator services (see note)	*23*	*23*	27	27	26	28	28	28	23	22
Infrastructure & base computing services	*125*	*123*	121	117	113	111	111	112	112	102
Total personnel required	*192*	*188*	**211**	**228**	**224**	**213**	**203**	**197**	**195**	**170**
Complement (see note) - actual for past years	*192*	*188*	185	177	172	163	153	147	149	149
DataGRID funding			10	10	10					
Additional personnel required			**16**	**41**	**42**	**50**	**50**	**50**	**46**	**21**
Cost of additional personnel (MCHF), assuming employed as CERN staff			2.4	6.2	6.3	7.5	7.5	7.5	6.9	3.2

note: The numbers in these rows correspond to the situation at end-January 2001. They are higher than the figures used in the LHC Computing Review report due a shift of responsibilities to IT from the Technical and Accelerator sectors

For comparison, the equivalent staffing levels are shown for 1999 and 2000. The total estimated requirement rises above the 1999 level by 20-30 posts during the R&D phase, falling again to the 1999 level by 2006. The staff complement available for these services is shown, declining according to the general CERN model until 2007. Note that the 'complement' level in 2007 is little more than that required for infrastructure and engineering support. Additional person-years not achieved in 2001 will have to be added, with justification, in future updates of this plan. The situation for the period of 2008 and later will become clearer after the first runs of the experiments.

For each of the next three years the DataGRID project provides funding for 10 posts at CERN. The shortfall in personnel numbers is calculated as the difference between the estimated requirement and the sum of the complement and DataGRID funding. Finally, the cost of the additional personnel required is estimated using an average cost per person of 150 kCHF, equivalent to the gross cost of employing a young engineer as a staff member at CERN. The most convenient model would be to employ such persons at CERN as staff with limited duration contracts. Other models of employment for work at CERN can be considered. It would be highly desirable if, amongst the additional personnel required, some could be experts detached temporarily from Member States to manage certain tasks.

In addition to the above, at the interface of the core software of the four experiments, the organisation of the distributed analysis and the grid middleware, 6 persons would be required at CERN IT Division (counted from the additional resources summarised in Table 1). These people will be needed until the data analysis in the distributed environment is operating satisfactorily. Clearly both these people and the core software effort in the experiments will be needed to perform successfully the prototyping work. It will become possible in due course to estimate better the effort needed for this activity during the construction phase of the production facility.

4.2 Materials

The estimated materials costs for providing the computing services at CERN are shown in Table 3, separated into the R&D phase (2001-4), the construction of the initial production service (2005-7) and the first year of the "maintenance" phase, in 2008. As explained in the report of the LHC Computing Review, maintenance of the facility involves replacing the commodity components at the end of their useful life (3-5 years, depending on the type of component) with more powerful, higher capacity components. The annual cost of maintenance is expected to decline in the years following 2008.

The lines *Base physics services*, *Engineering and accelerator services*, and *Infrastructure* cover the same areas as the equivalent lines in Table 2, except that the material costs of the LHC accelerator network are not included.

The costs for LHC computing are shown separately for the construction of the prototype, the upgrading costs for the computer centre[4], and the high bandwidth wide area network capacity for interconnecting CERN with the regional centres. The investment costs at CERN for the Tier 0 centre (data storage and reconstruction service) and Tier 1 centre (analysis service) are shown separately, along with the estimated cost of the outsourced systems administration and operation.

The funding allocated to provide these services in the current CERN Medium term Plan is shown, and the required additional funding is calculated. The total LHC costs are given for the prototype (R&D) phase and for the period of construction of the first production system.

---

[4] To modernise the power supply and distribution infrastructure, to provide additional cooling capacity, and to extend the area available for the installation of the large LHC computing fabric.

Table 3 - Summary of the estimated materials costs for computing services at CERN

Estimated Costs for Computing Services at CERN										
all numbers MCHF	*R&D Phase*				*First Production System*			*Maint-enance*	*Total R&D*	*Total First System*
year	*2001*	*2002*	*2003*	*2004*	*2005*	*2006*	*2007*	*2008*	*2001-04*	*2005-07*
Tier 0 investment					*16.8*	*18.4*	*17.4*	*18.3*		*52.6*
Tier 1 investment					*6.3*	*6.9*	*6.5*	*4.9*		*19.7*
Outsourced administration & operation					*4.4*	*6.2*	*7.0*	*6.5*		*17.6*
Prototype	*1.9*	*4.0*	*6.1*	*6.0*					*18.0*	
Computer centre refurbishment	*0.0*	*2.5*	*4.5*	*4.0*					*11.0*	
Physics WAN	*0.2*	*0.2*	*0.2*	*2.0*	*4.0*	*4.0*	*4.0*	*3.4*	*2.6*	*12.0*
Totals - LHC Operation	2.1	6.7	10.8	12	31.5	35.5	34.9	33.1	*31.6*	*101.9*
Base physics services	5.4	5.7	5.5	5.6	6.9	7	6.9	6.8		
Engineering & accelerator services	2.5	2.5	2.5	2.5	2.5	2.5	2.5	2.5		
Infrastructure	15.3	14.6	14.7	14.8	14.8	14.7	14.5	14.3		
Total materials required	25.3	29.5	33.5	34.9	55.7	59.7	58.8	56.7		
Funding available (Medium Term Plan)	23.2	22.9	23.4	24.2	25.7	26.3	26.4	34.1		
Additional funding required	**2.1**	**6.6**	**10.1**	**10.7**	**30**	**33.4**	**32.4**	**22.6**	*29.5*	*95.8*

The many difficulties in estimating reliably the cost of computing equipment more than a few years ahead must be emphasised. The cost estimates for the production system and maintenance phases must be regularly reassessed and will be better defined as a result of the prototype phase.

4.3 Summary of additional resources

Table 4 summarises the additional human and materials resources at CERN for the required services and for the work at the interface with the core software of the experiments. The additional human resources are expressed both as person-years, and as the gross cost of employing an equivalent number of people as young CERN staff members. It is assumed that those listed under core software will be provided for by the collaborating institutes and their funding agencies.

Table 4 - Summary of additional resources needed

Summary - Additional Resources needed										
	R&D Phase (Phase 1)				*First Production System (Phase 2)*			*Maint-enance*	*Total R&D 2001-04 (Phase 1)*	*Total First System 2005-07 (Phase 2)*
year	*2001*	*2002*	*2003*	*2004*	*2005*	*2006*	*2007*	*2008*		
Services required at CERN										
Additional personnel (person-years)	**16**	**41**	**42**	**50**	**50**	**50**	**46**	**21**		
Cost if employed as CERN staff (MCHF)	2.4	6.2	6.3	7.5	7.5	7.5	6.9	3.2	*22.4*	*21.9*
Additional materials (MCHF)	2.1	6.6	10.1	10.7	30	33.4	32.4	22.6	*29.5*	*95.8*
Service funding required at CERN (MCHF)	**4.5**	**12.8**	**16.4**	**18.2**	**37.5**	**40.9**	**39.3**	**25.8**	*51.9*	*117.7*
Core software for LHC (in IT Division)										
Additional s/w professionals (person-years)	**6**	**6**	**6**	**6**	**6**	**6**				
Cost if employed as CERN staff (MCHF)	**0.9**	**0.9**	**0.9**	**0.9**	**0.9**	**0.9**				

5 Proposal

Member States are asked to contribute to the first phase of funding of LHC computing as part of the base programme lying beyond the present budget limits. The work is intended to enable worldwide access to data and applications for a very large, distributed community of scientists. It will involve a pioneering large-scale implementation of a Grid fabric and should be of immediate value in other application fields. The first phase of the programme covers the development of the prototype during the years 2001-4. As this phase develops, and the cost estimates for the deployment phase become clearer, a second phase will be proposed.

The proposal will allow contributions to be made as direct funding for personnel to be hired by CERN, or for materials necessary to build the prototype. Part of the contributions can be made in the form of personnel detached to CERN as associates working under project control for two or three years.

While the major result of the programme will be the realisation of the LHC computing facility, software and applications, CERN will provide a rich training ground for staff during the coming years. When they return to their institutes or home countries they will bring back significant practical experience, providing a very real form of technology transfer. The programme can also be an important factor in enabling young staff of potential Tier 1 and Tier 2 centres to gain hands on experience during the period prior to the operation of LHC.

Annex 1

Membership of the review

Steering Group

Members:	S. Bethke (MPI Munich)	Chair	
	H.F. Hoffmann (CERN)	CERN Director for Sc. Computing	
	D. Jacobs (CERN)	Secretary	
	M. Calvetti (INFN Florence)	Chair of the Mgmt and Resources Panel	
	M. Kasemann (FNAL)	Chair of the Software Project Panel	
	D. Linglin (CC-IN2P3/CNRS)	Chair of the Computing Model Panel	
In Attendance:		**Representative**	**Alternate**
	IT Division	M. Delfino (CERN)	L. Robertson (CERN)
	ALICE	F. Carminati (CERN)	K. Safarik (CERN)
	ATLAS	N. McCubbin (RAL)	G. Poulard (CERN)
	CMS	M. Pimia (CERN)	H. Newman (CALTECH)
	LHCb	J. Harvey (CERN)	M. Cattaneo (CERN)
Observers:	R. Cashmore (CERN)	CERN Director for collider programmes	
	J. Engelen (NIKHEF)	LHCC chairman	

Worldwide Analysis / Computing Model panel

	D. Linglin (CC-IN2P3/CNRS)	Chair	
	F. Gagliardi (CERN)	Secretary	
Expt. Reps:		**Representative**	**Alternate**
	ALICE	A. Masoni (INFN Rome)	A. Sandoval (GSI Darmstadt)
	ATLAS	A. Putzer (U. Heidelberg)	L. Perini (U. Milan)
	CMS	H. Newman (CALTECH)	W. Jank (CERN)
	LHCb	F. Harris (U. Oxford)	M. Schmelling (MPI Heidelberg)
Experts:	Y. Morita (KEK)	C. Michau (UREC STIC/CNRS)	

Software Project panel

	M. Kasemann (FNAL)	Chair	
	A. Pfeiffer (CERN)	Secretary and CERN-IT representative	
Expt. Reps:		**Representative**	**Alternate**
	ALICE	R. Brun (CERN)	A. Morsch (CERN)
	ATLAS	D. Barberis (U. Genoa)	M. Bosman (U.A. Barcelona)
	CMS	L. Taylor (Northeastern U.)	T. Todorov (IN2P3 Strasbourg)
	LHCb	P. Mato (CERN)	O. Callot (LAL Orsay)
Experts:	V. White (FNAL)		

Management and Resources panel

	M. Calvetti (INFN Florence)	Chair	
	M. Lamanna (INFN Trieste and CERN)	Secretary	
Expt. Reps:		**Representative**	**Alternate**
	ALICE	P. Vande Vyvre (CERN)	K. Safarik (CERN)
	ATLAS	J. Huth (U. Harvard)	H. Meinhard (CERN)
	CMS	P. Capiluppi (INFN Bologna)	I. Willers (CERN)
	LHCb	J. Harvey (CERN)	J.P. Dufey (CERN)
Experts:	F. Étienne (IN2P3 Marseilles)	J. Gordon (RAL)	L. Robertson (CERN)
	F. Ruggieri (INFN Bari)	T. Wenaus (BNL)	K. Woller (DESY)
	G. Wormser (IN2P3 Paris)		

Annex 2

Summary of the approximate LHC computing requirements

General Parameters

Size of a recorded p-p event: 1 MB (up to 40MB for a Pb-Pb event in ALICE)

Data taking rate: 10^2 (Hz), down from 10^9 Hz p-p collisions, after several trigger levels.

Recorded p-p events per day: 10^7 (out of 10^{14})

Data taking: 10^7 seconds/y, or ~116 days/y (except Ion runs for ALICE ~15 days/y).

Recorded p-p events per year: 10^9

Storage per experiment

3 to 10 PB on tape, total ~28 PB (with 2/3 more per year beyond), Raw Data storage is ~ 1/3 of this total.

1 to 6 PB of disk, total ~11 PB (with 1/3 more per year beyond)

CPU (off-line) per Experiment

Best guesses today range from ~1 M SI-95 in LHCb, to ~2 M SI-95 for each of ALICE, ATLAS and CMS. Uncertainties are at least a factor 2. Estimates are the sum of Tier0, all Tier1 and all Tier2.

Networking

6-12 Gbps between main centres (1.5-3 Gbps per experiment)