Report of the Computing Resources Scrutiny Group

C.Bozzi (Italy), C.Diaconu (France), D.Espriu (Spain, *Chairman*), J.Flynn (UK), M.Gasthuber (Germany), D.Groep (The Netherlands), A.Lazzarini (USA), B.Panzer-Steidel (CERN), , W.Trischuk (Canada), B.Vinter (Nordic Grid), H.Renshall (CERN/IT, *Scientific Secretary*)

CRSG Mandate:

On behalf of the C-RRB we scrutinize:

- The resource accounting figures for the preceding year
- The use the experiments made of these resources
- The match between the refereed requests and the pledges from the Institutions
- The overall request for resources for every experiment for the following year and forecasts for the subsequent two years
- We also make recommendations concerning apparent under-fundings

RESOURCE ACCOUNTING AND USAGE BY THE EXPERIMENTS IN 2009

Report of the Computing Resources Scrutiny Group

WLCG accounting report for 2009

http://lcg.web.cern.ch/LCG/accounting/Tier1/2009/december-09/Master_accounting_summaries_December2009.pdf ,

EGEE accounting portal at CESGA.ES

http://www3.egee.cesga.es/gridsite/accounting/CESGA/tier1_view.html ,

and the reports that the experiments have provided to the CRSG

*We ask the experimental collaborations to provide complete reports by **1st MARCH** every year*

Overall usage 2009

	Site(s)	**Used/Available**
CPU	CERN	**26 %**
	T1	**69 %**
	T2	**N/A**
Disk	CERN	**71 %**
	T1	**81 %**
	T2	**N/A**
Tape	CERN	**41 %**
	T1	**64 %**

Report of the Computing Resources Scrutiny Group

ALICE

Resource	Site(s)	Used/Available
CPU	CERN	**16 %**
	T1	**29 %**
	T2	**53 %** (*)
Disk	CERN	**19 %**
	T1	**29 %**
	T2	**16 % (*)**
Tape	CERN	**17 %**
	T1	**10 %**

ATLAS

Resource	Site(s)	Used/Available
CPU	CERN	**30 %**
	T1	**97 %**
	T2	**154 % (*)**
Disk	CERN	**96 %**
	T1	**76 %**
	T2	**38 % (*)**
Tape	CERN	**58 %**
	T1	**39 %**

CMS

Resource	Site(s)	Used/Available
CPU	CERN	**36 %**
	T1	**43 %**
	T2	**59 %**
Disk	CERN	**60 %**
	T1	**82 %**
	T2	**40 %**
Tape	CERN	**47 %**
	T1	**87 %**

LHCb

Resource	Site(s)	Used/Pledged
CPU(kHS06*yr)	CERN	**52 % (*)**
	T1	**41 % (*)**
	T2	**69 % (*)**
Disk(TB)	CERN	**52 %**
	T1	**28 %**
	T2	**N/A (*)**
Tapes (TB)	CERN	**27 % (*)**
	T1	**22 % (*)**

Report of the Computing Resources Scrutiny Group

CPU usage: average over 2009

Disk + Tape: utilization at the end of the reported period.

In **RED**: percentage calculated by the CRSG

(*) used/pledged: the installed capacities were not available

Installed CPU and disk installed/used at the Tier 2 is not properly centrally accounted yet.

The CRSG requests the WLCG collaboration to take steps in the direction of implementing a complete accounting of the Tier 2 resources.

Tier 1 and Tier 2 disk: A metric that takes into account a disk access pattern would be most appropriate, allowing to identify "hot spots", namely data that are frequently accessed. We are aware of the technical difficulties involved and of the timescale needed to define and implement such a metric.

For the time being, we require that the disk utilization in the next reports from the experiments should be given in terms of the various data types involved and how frequently they were changed/replaced on disk, whenever possible.

Report of the Computing Resources Scrutiny Group

Percentage of use of the T1 +CERN resources by experiment in 2009

Collaboration	% of tape in T1+CERN used at end of period	% of disk in T1+CERN used at end of period	% of CPU in T1+CERN used	of which at CERN
ALICE	6%	4%	10%	29%
ATLAS	30%	57%	55%	9%
CMS	61%	34%	26%	35%
LHCb	3%	6%	9%	20%

Report of the Computing Resources Scrutiny Group

Percentage of use of the T2 CPU resources by experiment in 2009

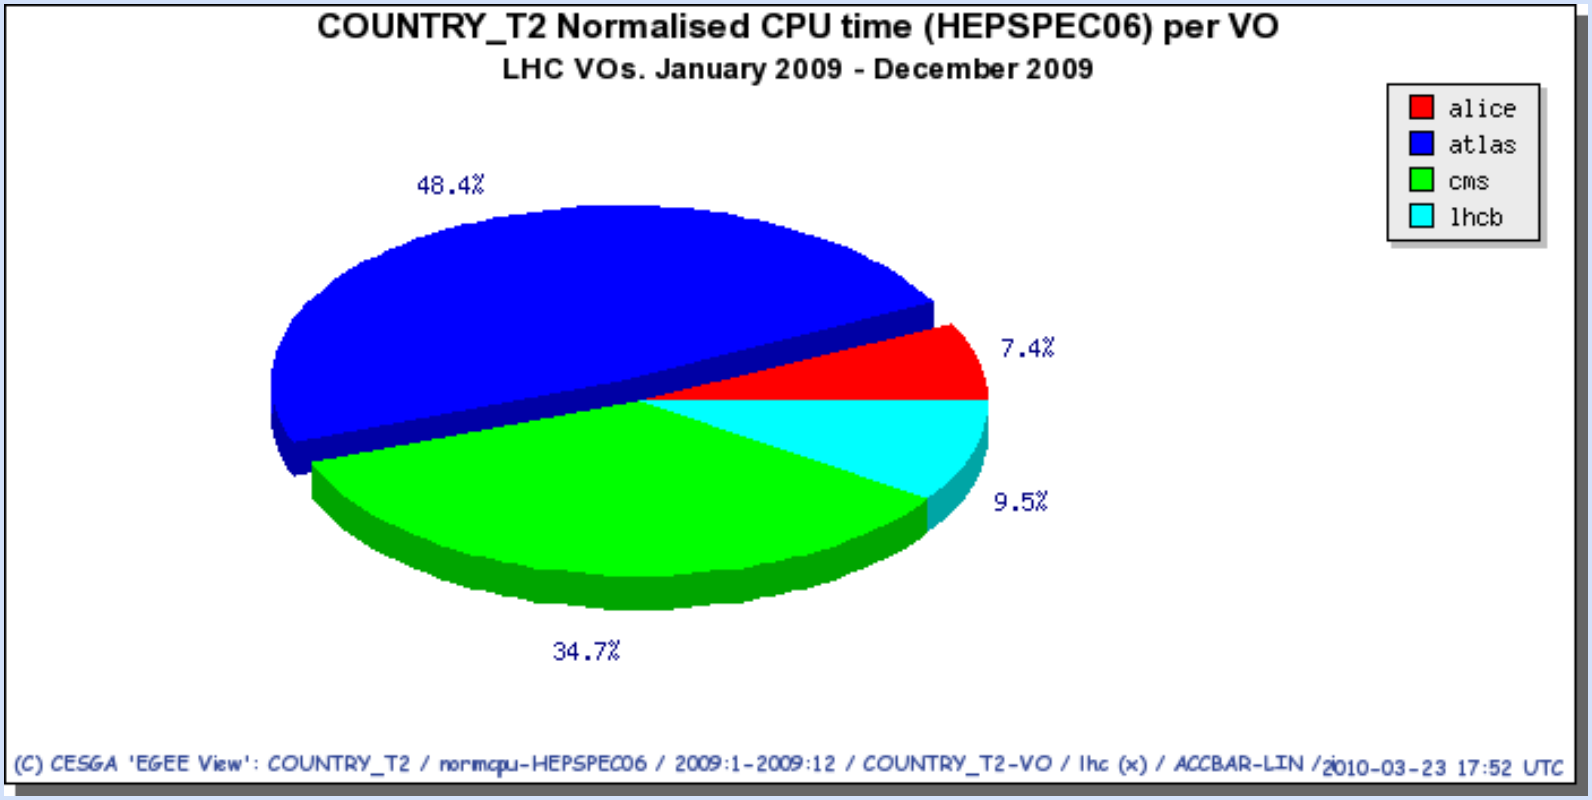

Report of the Computing Resources Scrutiny Group

T2 efficiency is assumed when computing the resources needed to be **60%.** However the following efficiencies are now reported (but note that the efficiency under chaotic access has not been really tested yet)

ALICE	**65%**
ATLAS	**85%**
CMS	**66%**
LHCb	**90%**

We ask the experimental collaborations and the sites to keep track of the efficiency of the CPU and disk usage to improve the use of the resources, and in order to adjust, as much as possible, the resources available to the needed ones.

No information has been provided so far on the efficiency of the disk usage so the **70%** estimate remains untested.

Report of the Computing Resources Scrutiny Group

The re-commissioning of the LHC in 2009, following the incident on September 2008, was very successful but the total amount of LHC collision data delivered was small. In addition the experiments accumulated a much larger amount of cosmics data, used for purposes of calibration and alignment of the detector systems.

The experiments have not used all resources made available to them, but the usage of the WLCG resources has been both extensive and intensive.

The tools and middleware appear now to be sufficiently mature to warrant a fruitful exploitation of the grid. *However more work is needed to improve the resource brokering code as some of the larger Tier 1 have been underused.*

This readiness combined with the enormous interest and enthusiasm in the first real LHC events, even if most of the data was really useful for commissioning only, triggered a noticeable increase in the use of the WLCG resources. This is particularly remarkable in the case of some of the collaborations that have shown a spectacular increase in the usage of the WLCG resources.

One of the experimental collaborations is clearly below the average in using the requested resources but we believe that is far too early to draw consequences.

DELIVERED VS. PLEDGED

Report of the Computing Resources Scrutiny Group

Available versus pledged resources

Resource	Site(s)	Available / pledged
CPU	CERN	**120 %**
	T1	**96 %**
	T2	**N/A**
Disk	CERN	**77 %**
	T1	**83 %**
	T2	**N/A**
Tape	CERN	**100 %**
	T1	**118 %**

CPU figures correspond to the average usage in 2009

Disk and Tape figures correspond to the resources in place at the end of the reporting period

Report of the Computing Resources Scrutiny Group

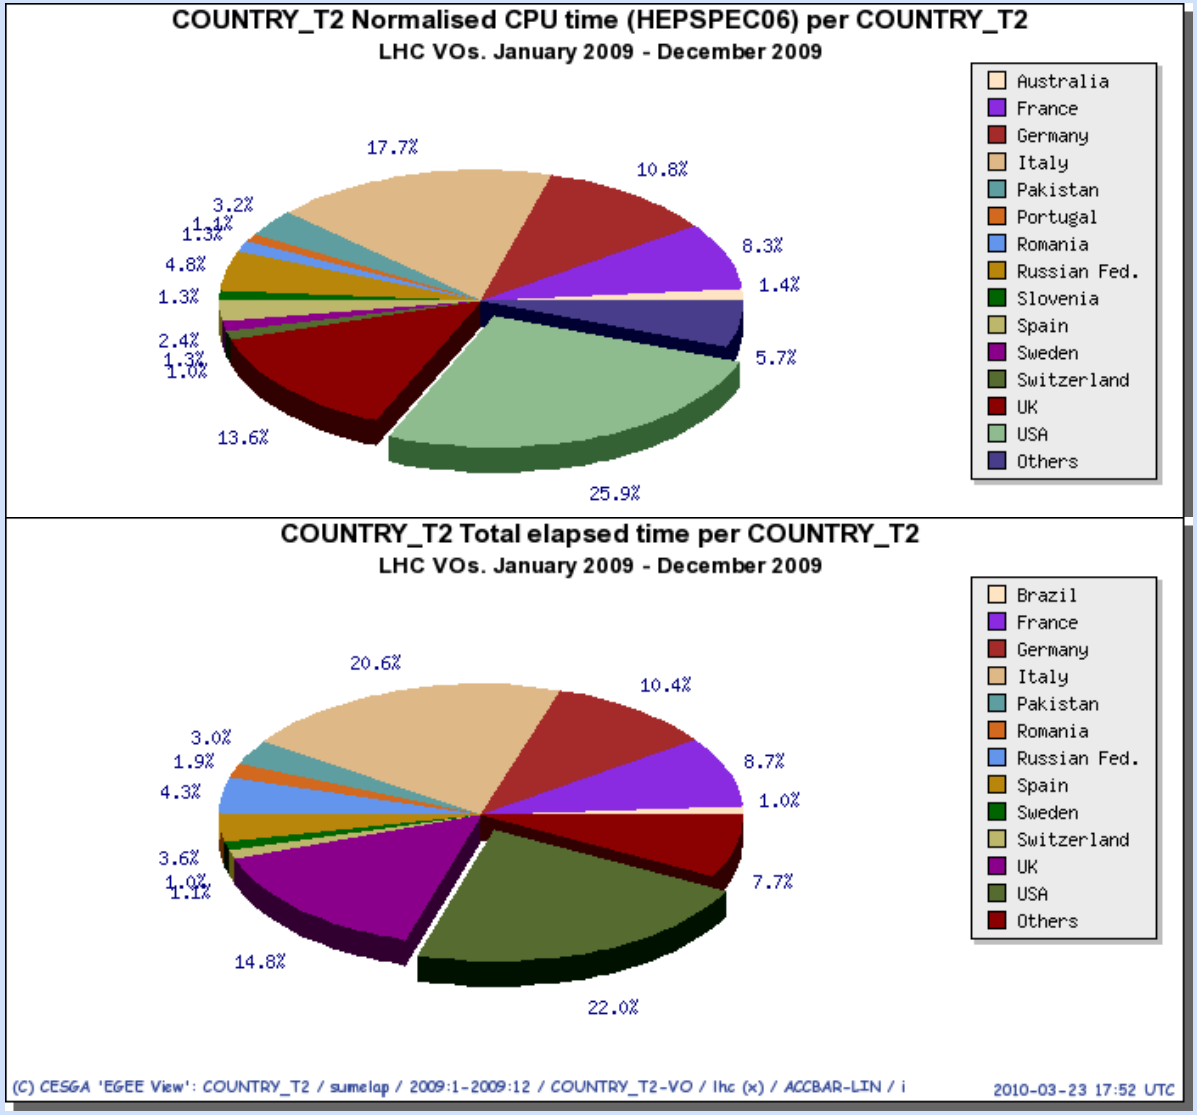

Report of the Computing Resources Scrutiny Group

At the end of the reporting period two Tier 1s were below the 90% ‘green line’ in terms of installed CPU: **TW-ASGC** and **IT-CNAF**.

The mismatches are more visible in disk capacity, which is expected to some extent as disk is the most expensive commodity. At the end of the reporting period only four of the sites were above the 90% level, and five were actually below the 80% line: **IT-CNAF, BNL**, **NL-NIKHEF/SARA, CA-TRIUMF, CERN** and **UK-RAL** (listed in order of decreasing mismatch with the pledge).

Tape storage was more in line with the pledges generally speaking, but a notable exception was **NL-NIKHEF/SARA** that had only installed a mere 13% of the pledges. **UK-RAL** was barely above the 50% ‘red line’.

None of these mismatches jeopardized physics analysis or simulated data production in view of the low turnout of the LHC, and to some extent they have been triggered in some cases by the evidence that the agreed pledges were over committing resources. Yet there is a danger in this attitude and it is desirable that in the future the pledged resources are made available within the agreed WLCG deadlines.

GENERAL ISSUES TO BE MONITORED BY THE LHCC & CRSG

Report of the Computing Resources Scrutiny Group

Over the last months a very fluid dialogue has been established between the CRSG and the LHCC including now joint meetings.

During 2010 no specific issues appeared for which we thought it was necessary to refer to the LHCC.

Some issues to be followed are listed in the next slide.

Report of the Computing Resources Scrutiny Group

Experiments should make an effort to reduce the raw event size (and of all subsequent derived formats) and event processing times by establishing a reduction profile without jeopardizing the physics. The experiments should also be prepared to make some compromises by revising their data distribution policy.

The potential proliferation of different data formats serving the same purposes should be watched closely. Strong overlap between different primary data sets may also be a matter of concern.

Care should be taken that the worldwide LCG resources are used as much as possible as there may be a tendency by collaborations to place heavier demands on CERN resources or suggest that a larger than originally planned part of their analysis should be done at CERN.

Both CMS and ATLAS plan to reprocess at sustained rates around 400Hz – 500Hz in coming years. In view of the duty cycle of the LHC this may not be strictly necessary. This requirement may have a large impact on cost.

The CRSG encourages a close collaboration of the different Tier centres with the experiments to achieve an efficient and cost-effective access to data. Intelligent storage management requires close information exchange and coordination between the computing fabric and the experiment management tools. Generally speaking, the CRSG encourages pre-staging from tape for centrally organized activities as a cost-effective measure.

Different collaborations still estimate the impact of pile-up due to the foreseen running conditions differently and uniformity in this respect should be sought.

SCRUTINY FOR 2011

Report of the Computing Resources Scrutiny Group

The expected beam time in 2010, 2011 and 2012 was subject to several revisions during the last months of 2009 and the early months of 2010. The final schedule was decided after the Chamonix meeting, in February 2010.

The LHC will run for the best part of 2010 and 2011, accumulating an integrated luminosity of 1 inverse fb @ 3.5 GeV + 3.5 GeV. About 10% of the time is expected to be dedicated to HI physics.

After this long run a long shutdown will follow to enable the machine to reach the design energy in 2013.

RRB year	RRB year start	RRB year end	Months (max) data taking	Total live time (in Ms)	pp	HI
2009		May '10	3	2.2	2.2	0
2010	June '10	March '11	8	5.8	5.1	0.7
2011	April '11	March '12	8	5.8	5.1	0.7
2012	April '12	March '13	0	0	0	0
TOTAL			19	13.8	12.4	1.4

Report of the Computing Resources Scrutiny Group

The experiments were asked to provide their request by 1st MARCH but the last request was received on 17th MARCH and some documents arrived as late as 30th MARCH. We would like to stress that the deadline for requests to RRB should be scrupulously followed in the future. *For the mid-October C-RRB we establish a **1st SEPTEMBER** deadline*

The experiments requests contained a limited number of changes with respect to previous iterations of the scrutiny procedure. This is an indication of the maturity of the computing models

Generally speaking we endorse the changes due to larger sizes and/or processing times as they are unavoidable in the starting period but we are of the opinion that there is room for optimization as more experience is gained. In any case, *we ask the experiments to reach a compromise in the requested resources and to try to establish a reduction profile*

As agreed with the ATLAS and CMS management in 2009, the scrutiny procedure for these two experiments is done by a common team of referees, using analogous techniques and methods.

The interactions with the experiments are quite fluid and we thank the respective management for their openness and collaboration. However the CRSG is nevertheless concerned about the short time it was allowed to develop its scrutiny.

Report of the Computing Resources Scrutiny Group

The CRSG recognizes the special characteristics of this combined run, extending over two years, and the uncertainties that remain in the current computing models and the need for redundancy/contingency. The situation for 2010 should indeed be regarded as somewhat exceptional, and to a large extent driven by the transient character of the start-up months, but we understand that commissioning activities will be largely over by 2011 and have scrutinized the resource requests in view of this.

The CRSG commits itself to undertake a revision of the present scrutiny ahead of the October 2010 C-RRB in the light of the usage of the resources up to that moment.

Report of the Computing Resources Scrutiny Group

CERN resources	2011	
	Scrutiny	*ALICE*
Tier0 CPU (kHS06)	47	*48*
CAF CPU (kHS06)	12	*14*
Tier0 disk (PB)	5.0	*5.8*
CAF disk (PB)	0.5	*0.5*
CERN tape (PB)	6.7	*6.8*
Non-CERN resources	**2011**	
	Scrutiny	***ALICE***
Tier1 CPU (kHS06)	121	*117*
Tier1 disk (PB)	7.2	*8.5*
Tier1 tape (PB)	13	*13*
Tier2 CPU (kHS06)	111	*121*
Tier2 disk (PB)	6.0	*6.7*

Report of the Computing Resources Scrutiny Group

CERN resources	2011	
	Scrutiny	*ATLAS*
Tier0 CPU (kHS06)	30.2	*30.0*
CAF CPU (kHS06)	45.0	*45.0*
Tier0 disk (PB)	0.9	*0.9*
CAF disk (PB)	5.5	*6.1*
CERN tape (PB)	12.0	*12.2*
Non-CERN resources	**2011**	
	Scrutiny	***ATLAS***
Tier1 CPU (kHS06)	208.1	*226.0*
Tier1 disk (PB)	23.1	*24.8*
Tier1 tape (PB)	30.5	*30.1*
Tier2 CPU (kHS06)	277.5	*278.0*
Tier2 disk (PB)	34.2	*38.4*

Report of the Computing Resources Scrutiny Group

CERN resources	2011	
	Scrutiny	*CMS*
Tier0 CPU (kHS06)	71.9	*72.0*
CAF CPU (kHS06)	34.1	*34.1*
Tier0 disk (PB)	0.9	*0.9*
CAF disk (PB)	3.6	*3.6*
CERN tape (PB)	21.6	*21.6*
Non-CERN resources	**2011**	
	Scrutiny	***CMS***
Tier1 CPU (kHS06)	139.5	*150.7*
Tier1 disk (PB)	18.0	19.5
Tier1 tape (PB)	49.4	52.4
Tier2 CPU (kHS06)	305.1	*319.5*
Tier2 disk (PB)	18.1	*19.9*

Report of the Computing Resources Scrutiny Group

CMS HI request

Resource	2011
Tier 0 CPU	*No additional request*
Tier 0 Disk	*No additional request*
Tier 0 Tape (PB)	***2***
Reconstruction CPU (kHS06) (Tier 1 services)	***5.8***
Reconstruction and data serving disk (TB) (Tier 1 services)	***875***
Tier 1 Tape (TB)	***800***
Tier 2 CPU (kHS06)	***9***
Tier 2 Disk (TB)	***500***

Report of the Computing Resources Scrutiny Group

LHCb resources requests for 2011

Date	Site	kHS06	Disk (TB)	Tape (TB)
April 2011	CERN	21	1510	2480
	Tier-1	65	3500	3470
	Tier-2	36	20	0

Report of the Computing Resources Scrutiny Group

After several interactions the scrutiny and the experiments requests show a good degree of convergence.

The CRSG was provided with simplified spreadsheets by the experiments that made the work of the referees easier. These simplified spreadsheets agreed consistently with the ones made of our own, constructed from the information contained in the computing models and the written requests submitted to us.

Some experiments request resources that are systematically higher than the ones resulting from their simplified models or from the CRSG understanding of the respective computing models. Unfortunately we have never been able to scrutinize in detail the origin of these differences as specific information is not provided. *The experiments and the CRSG will work towards a clarification as complete as possible of the requests*.

.

Report of the Computing Resources Scrutiny Group

General recommendations to the collaborations

The experiments should incorporate the running conditions into their models in a uniform way, in particular the implications of pile-up .

The experiments are asked to actively pursue the policy of reducing the size of their raw events and other derived formats in future years as much as possible as detectors become better understood. Reprocessing and simulation times also have a major impact and reduction should be sought. Experiments should also be prepared to make compromises.

We recommend the experiments make maximal use of the distributed resources in the GRID.

In the case of CERN resources, we advocate for a very clear separation between the contributions used for calibration and first pass reconstruction and central analysis, and those used to perform physics analysis by the CERN based physicists.

As far as data distribution is concerned, we notice different strategies which try to optimize the total CPU power required to analyse data, ranging from the maximization of replicas among the computing centres to the usage of streamlined primary datasets. Each approach has its own advantage and disadvantages. We recommend that the experiments use the upcoming data taking period to determine which strategy optimizes physics output while keeping resource requirements at a reasonable and sustainable level.

Report of the Computing Resources Scrutiny Group

On the CRSG membership

Two members of the CRSG left since the last C-RRB meeting. We have welcomed two new members: Martin Gasthuber from DESY who provisionally replaces Holger Martin and Bernd Panzer-Steidel from CERN who provisionally replaces Jürgen Knobloch, also from CERN. The appointment of these individuals has been made by the respective funding agencies but it is contingent on the approval of the C-RRB.

Following the April 2010 C-RRB three new members should join the CRSG in order to comply with the mandate of the WLCG MoU.

Jonathan Flynn from the University Southampton has a rather serious accident and has been unable to actively participate in the present scrutiny. The members of the CRSG wish him a quick and complete recovery.

Report of the Computing Resources Scrutiny Group

The 2011 requests of the four LHC experiments have been scrutinized after the initial period of data taking in 2009.

The usage report for 2009 has also been scrutinized in detail.

While we still find some potentially troublesome issues, our interactions with the experiments have led to an very good understanding of their models and we have reached a reasonable degree of convergence.

The CRSG is nevertheless concerned about the long term sustainability of the computing models.

The CRSG believes that the different computing models have largely proven their validity and we have no doubt that they will pass the test of the long 2010+2011 run.

We congratulate the four experiments for their excellent performance in the first days of the LHC and, of course, the machine people for their outstanding work.