Worldwide LHC Computing Grid Project
Project Status Report
Resource Review Board – 20th April 2010

Ian Bird
Version 1.0
5th April 2010

This status report covers the period from October 2009 – early April 2010. Further details on progress, planning and resources, including accounting and reliability data for CERN and the Tier 1 centres, and detailed quarterly progress reports, can be found in the documents linked to the LCG Planning Page on the web.

1. The WLCG Service

This reporting period is notable for the re-start of the LHC in November 2009 and first physics in March 2010. Real data is passing through the WLCG service as anticipated, with data being reconstructed, distributed, re-processed and analysed regularly. Analysis data from these early runs is being delivered to Tier 2 centres within hours of being acquired. The fact that this success of the computing system has gone largely unremarked is due to the breadth and depth of preparations and challenges over the past years, and the readiness of the Tier 1 and Tier 2 sites to accept data.

From 2009, in terms of data throughput Figure 1 shows the features of STEP'09 where large data rates were achieved of up to 4 GB/s, with overall throughputs of close to 1 PB/week. Later in the year in preparation for LHC startup there is a clear peak of traffic when the experiments were taking cosmic ray data and performing other specific tests. Finally at the end of the year we see the effect of the week or so of data taking followed by campaigns of processing and reprocessing.

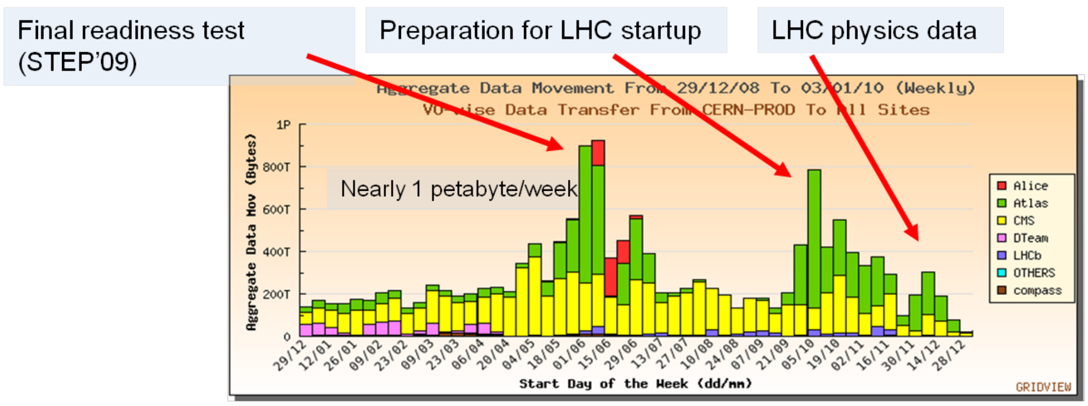

Figure 1: Data throughput 2009

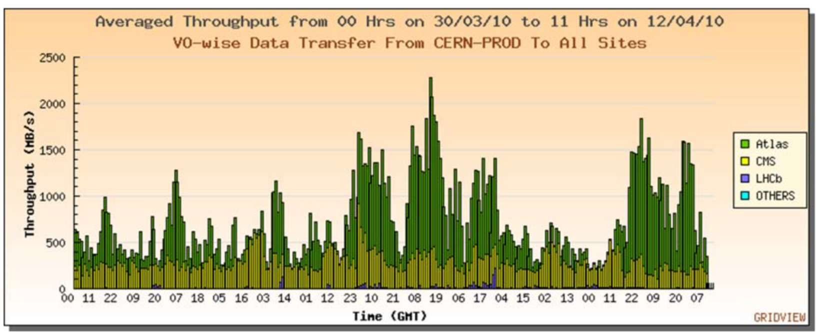

Figure 2: Data throughputs: data taking 2010

In Figure 2 is shown the transfer rates from 7 TeV collisions starting March 30 until mid-April, and shows the increasing rates – again up to peaks in excess of 3 GB/s. These data rates are matched by those seen in the Castor Tier 0 mass storage systems which are required to accept and serve data. Figure 3 illustrates the performance of Castor during the first week of April 2010 with real data taking. It is accepting data at around 4 GB/s and serving at rates up to 13 GB/s. These rates are higher than have been seen in Castor before, but the systems have not shown any performance related problems.

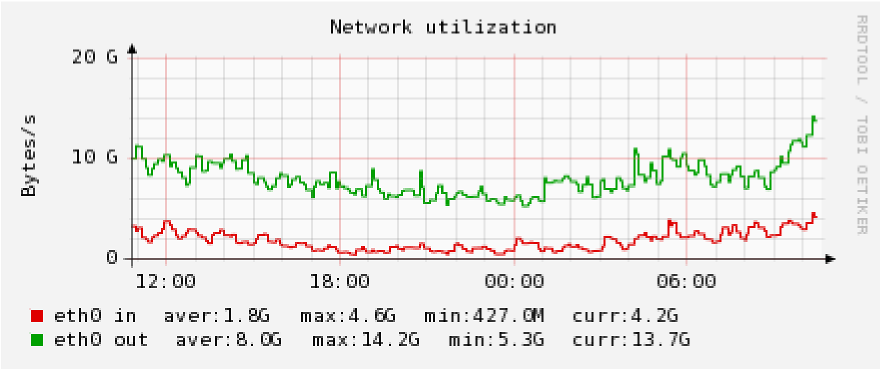

Figure 3: Data rates in Castor during 1st week of data taking

Another aspect of the scale of the WLCG service is that of job workloads. Figure 4 shows the numbers of jobs processed globally by WLCG in 2009, showing continually increasing workloads, which will no doubt increase even further with the real data. Well over half a million jobs per day are continually supported.

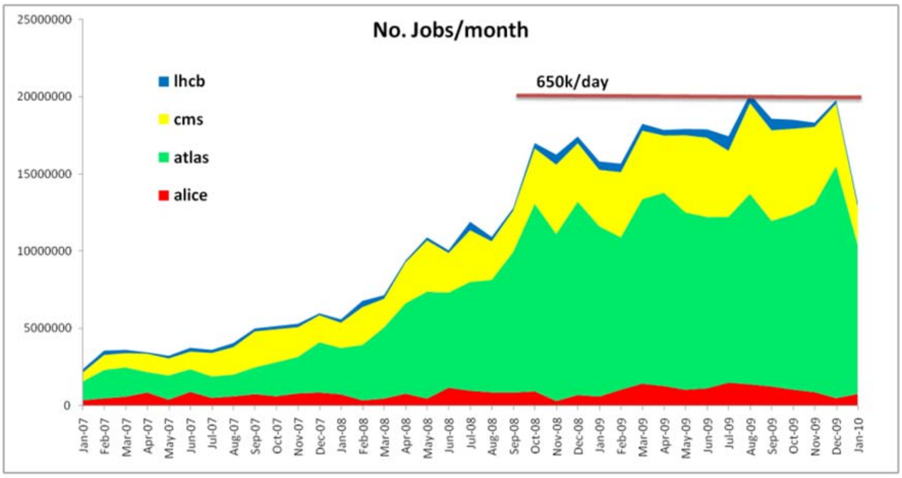

Figure 4: Job workloads in 2009

In terms of CPU time delivered the following 3 Figures show the total amount of CPU delivered to WLCG worldwide (Figure 5) and shows that this regularly now reaches an amount equivalent to approximately 100k CPU-days per day – which is interesting as 100k CPU is what was stated in the original project plans as the scale of computing required for LHC data analysis. The other two Figures illustrate for EGEE and OSG that while the CPU delivered to WLCG dominates, nevertheless in both cases the amount delivered to other sciences remains significant and increasing.

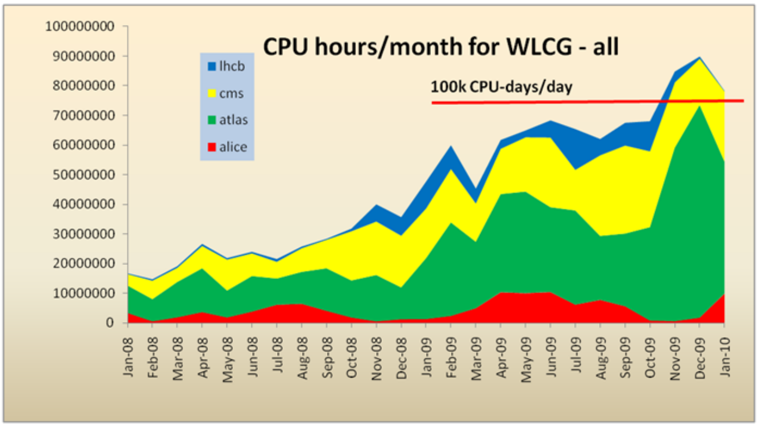

Figure 5: Total CPU time delivered in WLCG worldwide

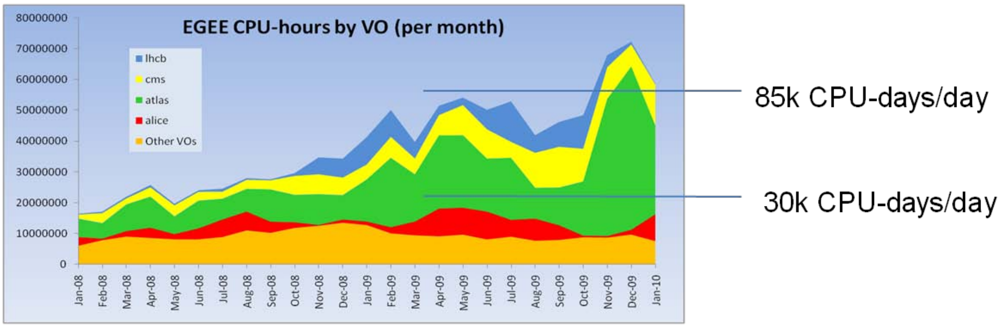

Figure 6: CPU time delivered by EGEE showing the fraction used by other sciences (orange area at bottom)

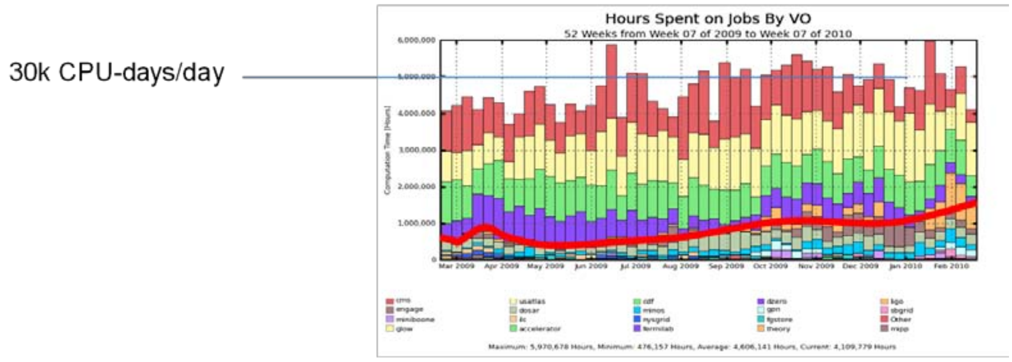

Figure 7: CPU time delivered by OSG showing the fraction used by other science (red line)

It is important to understand that during the entire period of testing, cosmic ray data taking (in Autumn) and real data taking (November/December and April) the system has continued to be fully used by ongoing simulation productions of the experiments. In recent months also there has been a significant increase in the numbers of users doing analysis – a number measured in hundreds of individuals for each experiment now. There are now significant analysis support activities within the experiments and with support from the project to ensure that these users are able to be productive in the analysis tasks. This effort will become a focus in the next months as more and more people start to look at the data.

Ongoing Service status

As agreed, significant service interruptions require a documented follow up (Service Incident Report – SIR). The full list for this period (summarised in the Table below) including the full incident reports can be seen as a summary in each Quarterly Report, or consulted on line at https://twiki.cern.ch/twiki/bin/view/LCG/WLCGServiceIncidents. These are followed by the Management Board, with the goal being that lessons are learned and disseminated to other sites. It is still the case that often power and cooling problems are the cause of these significant degradations.

Table 1: Incidents for which a report and follow-up was required

Site	Date	Duration	Service	Impact
CERN	3 Mar	18 hours	DB Replication	Replication of LHCb conditions Tier0->Tier1, Tier0->online partially down
IN2P3	15 Feb	4.25 hours	Batch	Local worker nodes lost network connectivity
PIC	10 Feb	7 hours	Spanish-CA CRLs expired at CERN	Complete blackout of services involving grid certificates either personal or host from Spanish CA at CERN:VOMS, FTS, etc.
CERN	7 Feb	4 hours	Batch Tier-0 Atlas RTT cluster	Degraded service on RunTimeTester cluster due to misconfiguration
CERN	30 Jan	2 days	CASTORATLAS	The xroot daemon was looping on the castoratlas name server because of a bug and slowing down all normal name server calls which was causing the migrator policy to fail
RAL	29 Jan	5 days	CASTOR - all instances	A scheduled outage to migrate the Castor Databases back to their original disk arrays encountered significant problems resulting in an extended outage.
ASGC	18 Jan	2 days	power system	power surge for one second and most services were restarted
GridKa/KIT	13 Jan	26 hours	site BDII and lcg-CE	site BDII query problems and missing lcg-CE information
IN2P3	4 Jan	6 hours	Batch	Local batch system database server overload.......................
PIC	19 Dec	4.5 hours	Cooling	Most of Tier-1 services shutdown to avoid increasing temperature due to cooling failure
IN2P3	8 Dec	1.5 hours	Networking	Grid services unavailability caused by load balancing mechanism failure
CERN	2 Dec	2 hours +	Site wide power cut	Most CC services down
RAL	30 Nov	n/a	Storage	LHCb Data Loss Incident at RAL
CERN	20 Nov	1h	SRM/ATLAS	SRM high failure rate and restart after thread exhaustion
CERN	18 Nov	10h	CMS	Performance degradation

			Dashboard	
IN2P3	12 Nov	n/a	Storage	CMS Data Loss Incident at FR-CCIN2P3
IN2P3	3 Nov	4h	Many	Many services have been disturbed due to automatic reboot of machines
IN2P3	14 Oct 2009	13h	batch	only very short jobs able to run
CERN	13 Oct 2009	1-2h	CASTOR nameserver sick	All CASTOR services dead
RAL	9 Oct	n/a	Storage (Castor)	data loss from Castor
IN2P3	8 & 10 Oct 2009	11h (8 Oct) and 6h (10 Oct)	SRM crashed	SRM service interrupted
RAL	4-9 Oct 2009		disk failures -> Oracle problems	CASTOR, LFC and FTS services down

A new Tier 1 Service Coordination meeting has been started improving information flow between service providers at the Tier 0 and Tier 1 sites and the experiments. The meeting covers a standing list of items, such as Data Management and Database Service issues (driven by the experiments' input), conditions data access, as well as topical service issues.

In preparation for service interventions a more formal change assessment has been introduced, in particular for critical services such as Castor. These assessments provide a basic risk analysis of the change, set out fallback scenarios, and allow the experiments to understand the implications of a change. This kind of process will be more important now that we anticipate accelerator running for an extended period without interruption, as significant changes may need to be made in this time.

During this reporting period there has been a marked increase in the regular attendance of sites at the daily operations meetings, surely encouraged by the presence of real data.

Figure 8 shows the experiments' measures of site availability for the first 2 weeks of data taking, and shows a very marked improvement in overall availability of all sites for all experiments – showing that support staff promptly reacting to service issues as they arise.

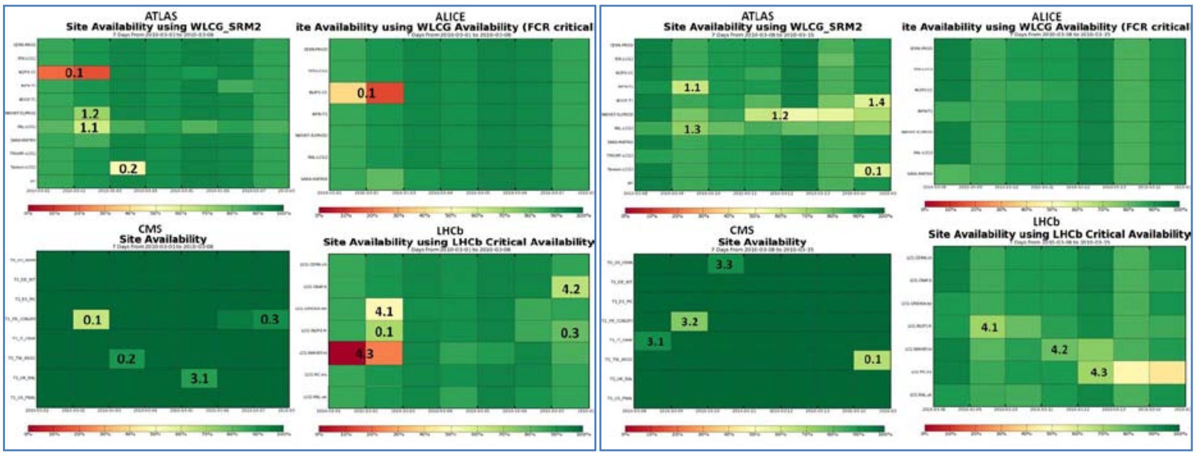

Figure 8: Site availability as measured by experiments for 2 weeks at start of data taking (left 4 – just prior to data, right 4 – 1st week of data)

In these plots the reliability of the sites as reported by the experiment tests are reported. Each box represents 1 day, and the colours represent the availability of the site according to the experiment tests.

Middleware services

There are no specific middleware issues, but it is important to note that the process of providing middleware upgrades, security patches, and new services has been shown to work well during challenges and data taking. This process is non-disruptive and allows sites to make upgrades on their own timescales, after significant effort by the middleware teams to reduce dependencies. Again this is important, as we enter an extended production phase of the service.

LHCOPN

The LHCOPN group proposed treating all T1-T1 links as part of the OPN and forming a group to work on data flows to/from T2s as well. This was approved by the CB and MB and representatives of experiments and non-T1 sites sought.

Security Patching

A large number of sites took too long to install an urgent security patch that was made public last summer. Many only patched when EGEE threatened them with suspension in October. A subsequent threat later in the quarter was controlled more quickly by issuing the suspension threat sooner. A new EGEE policy was adopted of suspending sites who do not patch within seven days of the Security Officer requiring it.

2. Applications Area

The SPI project has extended the AA software stack into new areas, such as the successful migration to the Intel icc compiler suite. More compiler flavours, such as the experimental llvm suite, are foreseen for the near future. The nightly build system, which is performing continuous building and testing of the AA project software, has been released in a new version which uses the full potential of multi-core build machines by parallelizing the builds on several levels. Moreover new features such as code coverage testing results and a new overview webpage have been deployed and are very much appreciated by the user community.

Several new versions of all Persistency Framework (PF) projects have been released. LCG_56d (Jan 2010) is based on ROOT 5.22 and was requested by ATLAS, while LCG_57a and LCG_58 are based on ROOT 5.24 and ROOT 5.26 respectively and were requested by LHCb. The three releases include several enhancements specific to PF projects, such as a COOL performance fix for CLOB data access and the CORAL move to the "light" version of the Oracle instant client, both requested by ATLAS. Reconnecting to an Oracle database after a connection glitch has been made more robust in CORAL, following many support requests of production users in the experiments at the time of the LHC start up. The POOL fast file merge feature implemented in an earlier release has also been validated by ATLAS.

Progress was made in improving monitoring and performance for the CORAL server software, but these enhancements have not yet been fully tested, therefore their release and deployment has been postponed to avoid disruptions to the ATLAS online system. New issues have been reported in the Oracle client libraries, caused by text relocations on SLC5 with SELinux enabled, and are being followed up with Oracle support. The port of CORAL to the icc compiler in Q4 2009 was useful to further investigate this problem, as the same symptoms have been observed in the CORAL libraries built using an old version of icc.

In the simulation are the investigations carried out on the transition between different hadronic models have been beneficial and relevant improvements have been made in several models: FTF, Bertini and CHIPS. The new public release of Geant4 - Geant4 9.3, was delivered on schedule; CPU speedups (both at initialisation and run-time) and improved memory management are included in this new release, as the result of strong cooperation with the ATLAS and CMS teams.

3. Site Reliability

The reliabilities for the last 6 months for CERN and the Tier 1 sites are shown in Table 2. In addition to the general reliability testing reported in this table, the experiment-specific measurements are published monthly together with the general reports. The regular reporting for the Tier 2s now also provides an overall Tier 2 federation reliability which is the average of the sites in the federation weighted by the number of CPU reported in the information system where that number is published.

All of the availability and reliability reports for all sites can be consulted at: http://lcg.web.cern.ch/LCG/reliability.htm.

Table 2: WLCG Site Reliability

	Average of the 8 best sites (not always same 8)					
	Sep-09	Oct-09	Nov-09	Dec-09	Jan-10	Feb-10
	99	99	99	99	99	98
	Average of ALL Tier 0 and Tier 1 sites					
	Sep-09	Oct-09	Nov-09	Dec-09	Jan-10	Feb-10
	96	95	97	98	97	97
Detailed Monthly Site Reliability						
Site	Sep-09	Oct-09	Nov-09	Dec-09	Jan-10	Feb-10
CA-TRIUMF	99	99	99	98	96	98
CERN	100	100	99	99	100	100
DE-KIT (FZK)	98	95	94	97	96	99
ES-PIC	99	99	100	98	98	98
FR-CCIN2P3	99	99	98	99	97	90
IT-INFN-CNAF	99	99	92	100	100	99

NDGF	97	72	90	91	96	99
NL-T1	91	96	98	98	96	90
TW-ASGC	83	99	100	99	99	97
UK-T1-RAL	93	83	99	100	88	97
US-FNAL-CMS	100	99	99	100	100	97
US-T1-BNL	100	99	100	99	100	98
Target	*97*	*97*	*97*	*97*	*97*	*97*
Above Target (+ >90% Target)	9+2	9+1	9+3	11+1	7+5	10+2
Colours:	Green > Target	Orange > 90% Target	Red < 90% Target			

4. Level-1 Milestones

A full report on milestones and progress can be found on the WLCG web at http://lcg.web.cern.ch/LCG/milestones.htm. Several of these have been mentioned in sections above.

As reported at the previous meeting, many of the milestones are now complete, as is appropriate as we enter an extended data taking production period, where stability is important. Providing a fuller set of metrics by which to assess performance will be a focus in this time. Nevertheless there are still a number of milestones where progress has been made. These are:

- Support for multi-user pilot jobs. This was the most important outstanding development to provide the services that support the experiments' ability to run pilot jobs that execute workloads from different users. The issue with this is the policy requirement that individual user actions at a site be traceable. The services required for this are "gLexec" and a back-end service that provides the "yes/no" decision on whether a specific user is permitted to access resources. In EGEE this back-end service is called SCAS, while in OSG this is called GUMS. There has been a long process to review the experiment frameworks in this context, to review the gLexec code and to thoroughly test it, and in EGEE, to provide the SCAS service. This process has progressed significantly and these services are now available for deployment and are being deployed. However, the deployment is very slow with many sites unwilling to deploy these services without there being some prior experience. Nevertheless, the experiments need to be able to run jobs in this mode for analysis. Thus the Management Board agreed a time-limited suspension of the policy to allow the experiments to use this facility, while the supporting services are widely deployed. This situation is regularly reviewed.
- Automated gathering of installed capacity data. The mechanism and process for data publication are agreed, and most sites publish data now. The process of validating this data is ongoing but takes some time as the data is complex. Tools and tests are available and in production to help in this process and to provide consistency checking.
- CREAM CE deployment. This is now much further advanced with all Tier 1s having instances of this service deployed. The intention is that this CE should replace the old LCG-CE as it will provide better scalability and performance.

5. Update of Experiment Requirements

With a new schedule for accelerator running in 2010 and 2011, a re-assessment by the experiments of their computing and storage requirements has been made. This assessment also provided requirements for 2012 under the assumption that the accelerator would not be taking data in that year. These requirements have been reviewed by the LHCC referees and the C-RSG referees.

The assumptions that have been used as input to this reassessment are the following:

- Similar schedule in both 2010 and 2011, with physics running from mid February until end November (although in 2010 data taking did not start until end March), and no running in 2012. Running in 2010 and 2011 will be essentially continuous, with short technical stops for 3 days per month only.
- There will be 1 month of Heavy Ion (HI) running in November each year, with a few weeks preparation period in October.
- From the Accelerator groups and CERN management, the following parameters of the running are agreed as the base assumptions:
    - Availability of the machine for physics = 70%; the remainder is technical stop and Machine Development,
    - Efficiency for physics (fraction of the 70%) where colliding beams are available = 40%,
    - Thus we obtain around 200 hours effective physics running per month.
- It is recognised that the 2010 requirements have been previously agreed and have been closely matched by the existing pledges and actual installations. Thus, no change for 2010 is foreseen in this exercise.

The numbers presented in the tables below have been presented in the WLCG Management Board and to the LHCC referees. The result of the C-RSG scrutiny of these updated requirements will be available to this C-RRB meeting.

There are several important points to make about these requirements:

- The numbers shown for 2010 are the numbers presented in the last RRB for both requirements and pledges (although the pledges have been updated with information arriving after the meeting) for comparison.
- For 2011 there are modest increases in need, reflecting the additional data and storage requirements. There will be an important analysis load in that year.
- In 2012 there are only minor increases in requirement to support continuing analysis. However, it is important to understand that this must not be seen as a reason not to invest in computing resources during that time, as all sites will need to be able to replace and renew equipment.
- It is not yet known what the likely running in 2013 will be, but if there will be higher energy, and higher luminosity running, it will be important that we are able to provide the appropriate level of resources for that, without having to have a single major purchase.

Table 3: Updated requirements for 2011, 2012. The 2010 numbers are those presented and agreed at the previous RRB

CMS	2010	2010 pledge	2011	2012	ATLAS	2010	2010 pledge	2011	2012
CERN CPU	96.6	96.6	105	105	CERN CPU	67	67	75	28
CERN disk	4.1	4.1	4.5	4.5	CERN disk	3.9	3.9	7	7.5
CERN tape	14.6	14.6	21.6	24.6	CERN tape	9	9	12.2	12.5
T1 CPU	100.5	105	150	150	T1 CPU	192	216	226	223
T1 disk	13.4	12.5	23.4	25.1	T1 disk	21.9	22.3	25	27
T1 tape	23.3	24.1	52.4	64.7	T1 tape	14.2	15.5	30	39
T2 CPU	195	197	320	320	T2 CPU	240	217	278	295
T2 disk	9.2	13.7	21	23	T2 disk	24.8	21.3	38	44
ALICE	2010	2010 pledge	2011	2012	LHCb	2010	2010 pledge	2011	2012
CERN CPU	46.8	46.8	62.3	65.7	CERN CPU	23	23	21	21
CERN disk	5.5	5.5	6.7	9.1	CERN disk	1.29	1.3	1.5	1.7
CERN tape	6.3	6.3	12.5	9.6	CERN tape	1.8	1.8	2.5	3
T1 CPU	57.6	46.3	102.5	146	T1 CPU	44	44.7	65	65
T1 disk	10.8	6.3	14.4	13.1	T1 disk	3.29	3.4	3.5	3.7
T1 tape	16.3	8.7	25.2	23.6	T1 tape	2.4	3.1	3.47	4.42
T2 CPU	89.6	54.1	100.2	140	T2 CPU	38	43	32	32
T2 disk	12.6	4.2	16.4	11.4	T2 disk	0.02	0.4	0.02	0.02

Table 4: Summary of the updated requirements

Totals	2010	2010 pledge	2011	2012
CERN CPU	233.4	233.4	263.3	219.7
CERN disk	14.79	14.8	19.7	22.8
CERN tape	31.7	31.7	48.8	49.7
T1 CPU	394.1	412	543.5	584
T1 disk	49.39	44.5	66.3	68.9
T1 tape	56.2	51.4	111.07	131.72
T2 CPU	562.6	511.1	730.2	787
T2 disk	46.62	39.6	75.42	78.42

6. Planning and Evolution

Tier 0

The evolution of the power needs of the Tier 0 centre have been continually evaluated over the past several years. Following aggressive replacement of older equipment with newer, more efficient machines, the point at which the available power is insufficient is now anticipated to be 2012-13. However, the practical limit of 2.9 MW usable power, may be somewhat less (closer to 2.7 MW), and it may be risky to approach this limit too closely. As seen from the figure below, there are some uncertainties on the power evolution because of uncertainty of technology evolution and actual requirements, so it is important to take action now.

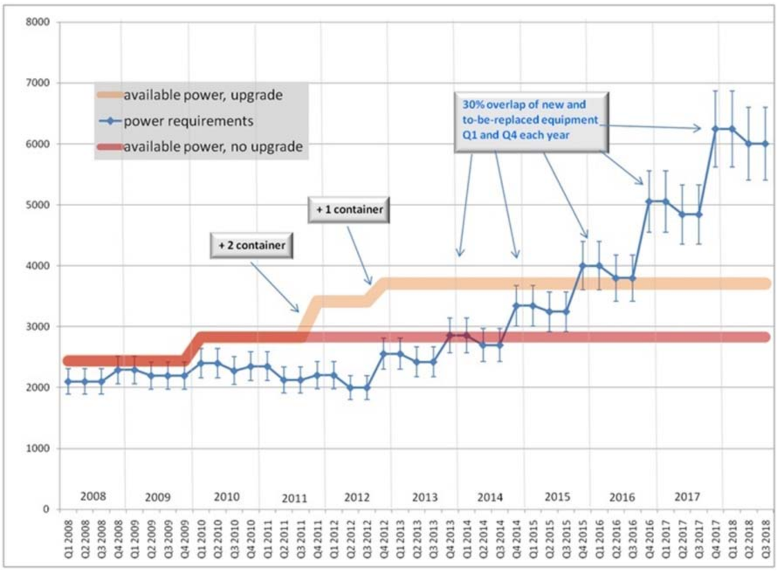

Figure 9: Evolution of Tier 0 power requirements

The needs of the experiments after 2012 are still unclear, and until there is some significant experience with real data it will be difficult to predict the evolution. It has also become apparent that there is a more urgent problem which is the total amount of backed-up (Diesel) power available for critical services. This is currently limited to a total of 300 kW, and is all used. It is difficult to add new services which require such reliability.

The strategy now is the following:

- Consolidate the power in the present building, adding 400 kW, with an increase in backed-up power to around 600 kW. This is unlikely to be available before mid-2011.
- Provide external hosting in the Geneva area for services requiring backed-up power. An agreement has been made for an initial 100 kW, to start within the next 2 months. Services are being identified as candidates to move to this facility.
- Provide containers. The initial technology assessment has been done, and a market survey started. A location on the Prévessin site has been identified. However, this requires a new service building for the power and cooling infrastructure, and this requires planning permission. The timescale is thus that container based capacity could not be available before the end of 2011.
- Remote hosting of Tier 0 services. This has been suggested by Norway and others. We are awaiting concrete proposals for such a solution.

It must be noted that the current stop-gap solutions of a combination of external hosting and container infrastructure requires a considerable yearly investment. It is clear that a decision must be made on a final strategy, between the original idea of a new Computer Centre at CERN and a mixed long term strategy with remote hosting of significant facilities. This decision must be taken soon.

EGEE to EGI transition

The EGEE project is now in its final month. The projects proposed to the EC to support the new EGI-based structure were several:

- EGI-Inspire: a project to provide the support of the main EGI organisation and overall coordination of NGIs and the global operation. It has a sub-task specifically aimed at supporting existing large heavy user communities such as WLCG.
- EMI: providing essential ongoing middleware support, and work on harmonising the different European middleware stacks. Of interest for WLCG would be the bringing together of gLite and ARC.
- A project (ROSCOE) that included a "Virtual Research Community" for High Energy Physics, that was intended to provide application community support.

These project proposals have now been reviewed. The first two – EGI-Inspire and EMI – have been successful, the former being funded at the level requested, and the middleware project approved with a 9% budget cut relative to the request. However, the ROSCOE project was not successful, and there will be no such community support project.

This situation does not represent a major risk for WLCG, as the EGEE to EGI transition has been well planned by EGEE, and is well advanced in practice. The countries representing the majority of the resources have NGIs and the Tier 1s are well placed within those countries.

The important operational tools (such as GGUS, monitoring tools, etc.) are assured by the institutes responsible even if project funding does not appear. The WLCG operational procedures are well tested and are today mostly independent of the existence of EGEE or EGI. The activity in the EGI-Inspire project for the support of existing user communities contains key application support such as Dashboards, Ganga, and has specific tasks and effort dedicated to each experiment.

The EMI project contains the essential middleware support for the middleware services that WLCG depends upon. However the available effort is at the level of ongoing support and the above-mentioned harmonisation – significant developments are excluded. Thus it becomes important that we should consider the strategy for the longer term evolution and sustainability of the middleware and other services that we rely on.

Earlier concerns over the apparent exclusion of some countries from EGI that were previously part of EGEE has been addressed to some extent by introducing the category of associate members of EGI. These countries are full partners in the EGI-Inspire project but do not have the same status in the EGI organisation.