Worldwide LHC Computing Grid Project
Project Status Report
Resource Review Board – 13th Oct. 2009

Ian Bird
Version 1.0
2nd Oct. 2009

This status report covers the period from April – September 2009. 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

The significant feature of the service during this period was the preparation for and execution of the STEP09 (Scale Testing for the Experimental Program 2009) exercise at the end of May and into June. Since then the WLCG service has been running in normal production mode supporting sustained workloads for the experiments, although since June the data transfer rates have been at a fairly modest level.

STEP'09

As reported at the last meeting this exercise was the only real opportunity this year to perform a significant scale test with all the experiments scheduling tests in synchrony with each other. The main goals were to test specific aspects of the computing models that had not previously been fully tested. The two most important tests were the reprocessing at Tier 1s with experiments recalling data from tape at the full anticipated rates, with at least ATLAS and CMS doing this together; and the testing of analysis workflows at the Tier 2s and Tier 1s where appropriate. A post-mortem workshop with more than 100 attendees was held in July to summarise the results from the challenge (http://indico.cern.ch/conferenceTimeTable.py?confId=56580). The overall results from the exercise were very encouraging with many sites meeting and exceeding the metrics set by the experiments. A few Tier 1s and several of the Tier 2s showed problems in certain areas. These are the subject of follow-up actions to understand the problems and in some cases to repeat the tests. A full follow-up test on the same scale as STEP'09 is not possible due to the scheduling constraints of the experiments and their own programs of work.

One thing to note is the much higher level of data transfer rates that were achieved during this period, with rates from CERN to the Tier 1s getting up to 4 GB/s – well in excess of the stated requirements of the experiments. This is illustrated in the Figure below. This is even more remarkable since there were multiple fibre cuts in the LHCOPN, which, while causing some decrease in bandwidth to several sites was not a cause of failure. During the challenge there were also a number of other scheduled and unscheduled service interventions. This was probably a realistic demonstration of the environment that must be expected during data taking.

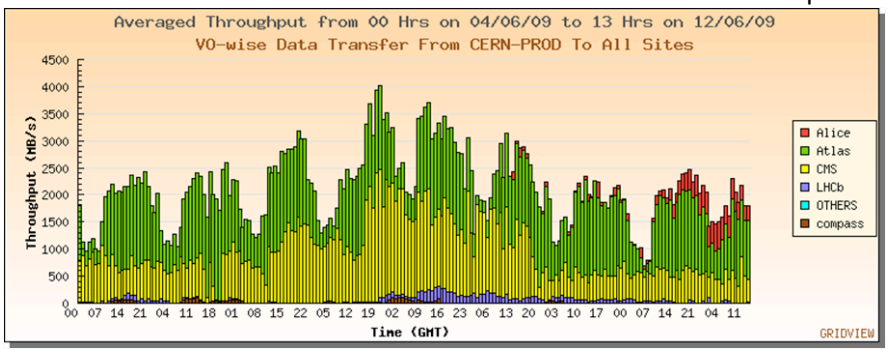

Figure 1: Data transfers from CERN during STEP'09

For CMS the activities included demonstrating the ability to recall data from tape at the Tier 1s at the required rates, testing various pre-staging tools all of which worked successfully. The processing at the Tier 1s ran smoothly at the pledge level of the number of job slots. It was noted that the CPU efficiency was much higher when data was pre-staged to disk. In general the performance of data recall and reprocessing was good, although some sites had downtimes. In terms of data transfers, they tested flows (rates and latencies) from Tier 0 to Tier 1, between Tier 1s, and from Tier 1 to Tier 2. Even transfers to Tier 2s ran smoothly although relied on data being pre-staged at the Tier 1s. For analysis their goal was to demonstrate analysis at a scale using all the pledged resources at the Tier 2s – ensuring that at least 50% was dedicated to analysis, the remainder doing ongoing tasks. They succeeded in doubling the overall use of Tier 2s, and showed that they can easily fill the pledged resources. Moving data to and from Tier 2s was not always optimal, especially for Tier 2s that had not previously participated in such a test.

ATLAS ran the STEP'09 activities at full rate for some 11 days. This included testing data distribution, reprocessing, production, and analysis work flows. In terms of tape recall and processing at Tier 1s ATLAS reported similar success as CMS. In the 11 days of the test, they succeeded in adding 4 PB of data to the ATLAS total stored on the grid. Five of the ATLAS Tier 1 sites achieved more than 5 times the nominal data taking rate for reprocessing, including recalling the data from tape; 4 of the other Tier 1s achieved better than 90% of the target, and one reached only 50%. For analysis they submitted around 1 M jobs, with a good success rate. This was achieved while the Tier 2s were also continuing to run MC production jobs, scaled to keep the Tier 2s' resources full.

LHCb did similar tests, with no problems in achieving the data transfer rates needed. In terms of re-processing they had no problems with staging data from tape, although their reconstruction jobs were hit by a bug in the access to the conditions databases. This bug was fixed during the exercise. They also reached a scalability limit of ~10k simultaneous jobs in DIRAC and are working with IT to devise a means to improve scalability of their system. Nevertheless they achieved significant job submission rates, and involved 115 sites in their analysis tests.

ALICE demonstrated that they could also achieve transfer rates in excess of 300 MB/s which is the rate required for Heavy Ion data taking. They also tested data transfers from Point 2 to the Tier 0 achieving the required 1.25 GB/s to tape at the Tier 0.

The Figure below illustrates the good performance of the Tier 0 tape system, in this example as viewed by CMS.

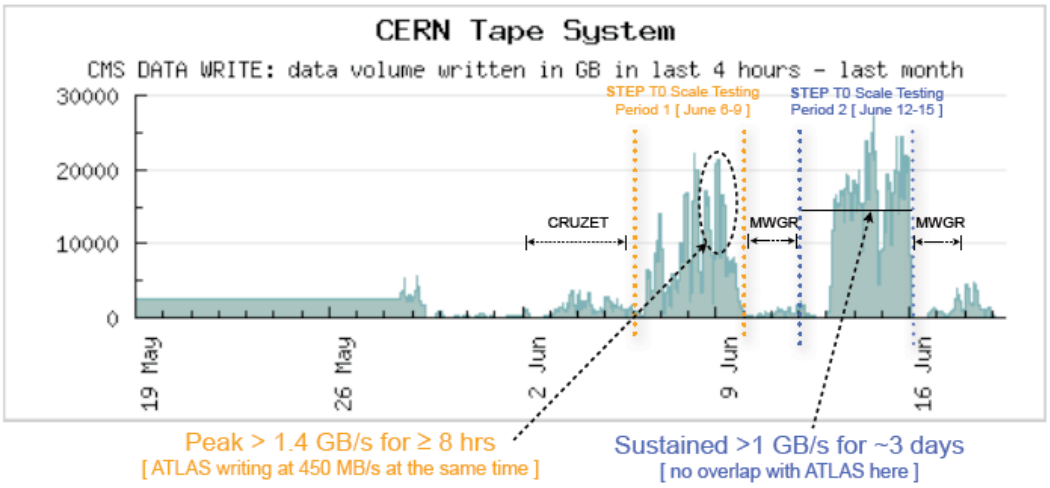

Figure 2: Tape writing performance at Tier 0

In the first part of the test CMS successfully wrote data to tape concurrently with ATLAS without interference. In the second part of the test CMS alone was able to achieve very high rates to tape over sustained periods.

Other details of experiences at the Tier 1s and Tier 2s can be seen in the slides shown in the post-mortem workshop.

Follow up activities

Several specific points arising from the STEP'09 exercise have been noted and are being followed up by the management board. Of the Tier 1 sites where the performance during STEP had not been fully satisfactory (due to scheduled interventions or to problems that arose), specific tests have been made, or are planned at most of the sites by the experiments to validate those sites.

Other general issues that were shown during STEP'09 included the recognition that data access performance (between site Storage Element and the worker node) depends strongly on whether data is copied to the local disk on the WN or read directly from the SE. A working group is investigating how to improve data access flows and hopefully to recommend configurations. It also became apparent with the large number of simultaneous jobs at many sites that the shared file systems (NFS) used for the software areas could not sustain the performance needed. Those sites affected must investigate alternative solutions that scale better, or make local installations of the software directly on the WNs. Finally, there are concerns that at some of the Tier 2s the network infrastructure between the batch systems and storage systems is not adequate to sustain the data rates needed by a full analysis workload.

These and other more detailed specific issues are regularly followed by the project and in some cases milestones have been added to track progress.

Ongoing Service status

Since the end of STEP'09 workloads have continued at a fairly high level, illustrated in the Figure below that shows the level of CPU consumed. Data transfers have not been at such a high rate. Most of the experiments will start taking cosmic data again over the coming weeks in final preparation for the accelerator start up in November.

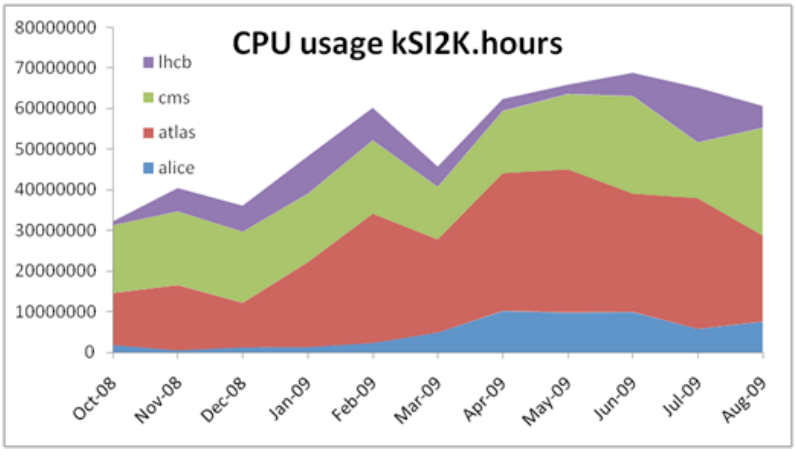

Figure 3: CPU usage

As agreed, significant service interruptions require a documented follow up (Service Incident Report – SIR). The full list (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.

The levels of support required to run the WLCG service, both during the STEP'09 challenge and since during ongoing production work, are now more sustainable and manageable. This is in part due to the improvements and understanding provided by the incident reports, follow ups, and dissemination of the lessons learned.

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

Site	Date	Duration	Service	Impact
CERN	21 Sep 2009	08:00 - 18:00	DB Replication	ATLAS Replication Tier0->Tier1 down
RAL	15 - 17 Sept 2009	2 days	CASTOR	Disk to Disk (D2D) transfers started failing during a planned upgrade to the NS
FZK	7 - 16 Sep 2009	10 days	ATLAS RAC	3D Streams replication blocked then degraded
CERN	5 & 8 Sept 2009	2 * 2 hours	CASTOR LHCb	two Castor Database problems
CERN	26 Aug 2009	18:40 - 23:30	Batch	Public and production queues closed
ASGC	17 Jul 2009	6:00 - 10:00	Power cut	Most services went down and restarted
ATLAS	13 Jul 2009	10:00 - 11:00	Central Catalogs	Degrade of performance
NL-T1	STEP09			Post-mortem of STEP'09 experience
OPN	10 Jun 09	>1 day	LHC OPN	primary circuits to ASGC, CNAF, KIT, NDGF, TRIUMF (incl. backup)
FZK	STEP09	many days	storage	
ATLAS	27 Jun 09	2 days(?)	PVSS2COOL	online reconstruction was stopped
ATLAS	24 Jun 09	8 hours	PanDA and ATLR	Degraded PanDA service, impact on other offline DB services on ATLR
CERN	11 Jun 09	n/a	LHCb conditions access, LFC	scalability problem
CERN	18 Jun 09	2 hours	Batch & CASTOR services	down
IN2P3	10 Jun 09	7 hours	GridFTP	Transfers
CERN	4 Jun 09	n/a	CASTOR LHCb	accidental garbage collection of tape0disk1 files

CERN	3 Jun 09	n/a	CASTOR LHCb	accidental re-enabling of garbage collection in lhcbdata
CERN	1 Jun 09	~4 hours	DB services	unavailable
PIC	23 - 26 May 09	3 days	LFC	instability
PIC	14 May 09	5 hours	cooling	down
SARA	04 May 09	36 hours	MSS	down
IN2P3	3 May 09	44 hours	cooling	down

2. Middleware services

The middleware support processes have been working as intended, with regular updates provided as required for service improvement, security and specific functional issues. Even during the STEP'09 exercise the standard middleware support process was capable of supporting the needs. There are a few specific services and components where important changes have been introduced or are anticipated. These are listed below.

- Compute Element. A new version of the new CREAM CE is now available resolving many of the outstanding issues. This is now considered as ready for wide-scale deployment at EGEE sites in parallel with the existing CEs initially. ALICE have reported good success with this CE in earlier production testing. CREAM avoids some of the scalability problems existing in the present lcg-CEs which are directly based on the old Globus gatekeeper and inherits many problems.
- The WMS version 3.2 now provides a component that allows submission direct to CREAM, removing one of the obstacles to a wider deployment of CREAM. This WMS version can also submit direct to the ARC CE used in the Nordic countries without need for gateways or experiment-specific adaptation.
- The glexec and SCAS components required to support multi-user pilot jobs are now considered ready for more general deployment, although so far deployment take-up by sites has been quite slow.

In terms of Mass Storage systems (MSS) updates are available for both dCache and Castor. For dCache version 1.9.4 introduces ACLs to ensure file protection for tape data, and a migration to a new namespace service ("Chimera") is recommended for large installations to avoid scalability problems later. dCache have also introduced the concept of a "golden" version that will be supported for the first years of data taking. However this version is not yet available, and sites must make individual decisions on upgrades and migration to Chimera before data taking.

Castor version 2.1.9 also consolidates many patches, but importantly provides better support for large scale analysis use cases with xrootd support.

At the end of last year the remaining "missing" functionality of SRMs that had earlier been agreed as part of the requirements process, was put on hold before accelerator start up. These requirements have been reanalysed in the light of experience gained during 2008 and 2009 testing and production use. These needs have been better understood and reprioritised, and will be implemented where appropriate in coming versions.

3. Applications Area

The Application Area has continued to prepare complete configurations of the application software stack. Three new LCG Configurations LCG 56a, 56b and 56c have been released on request of the LHC experiments. These contain the versions of software that are going to be used for the first data run. It also includes validated binaries for the SL5 (32 and 64 bit) platform. The SPI team has continued to develop the LCG nightly build system, including further improvements on multi-core build architectures and replacing the current web infrastructure by a database system. The system is now successfully used by all AA projects. In addition, experiments such as ATLAS and LHCb use the nightly results to build their own software stacks for different configurations providing in this way a very rapid feedback to changes in the underlying software.

The new ROOT production version 5.24 released end of June contains several improvements in the ROOT mathematical libraries. The core libraries have been consolidated with small improvements and bug fixes, thanks to the effort put in developing new test programs. The major new developments have been put in libraries like RooFit, RooStats and TMVA. A new development version 5.24.2 was released end of September and includes, among other changes, an optimization mechanism for the size of the TBuffer baskets, which improves the CPU performance and the used memory for reading/writing event data file for CMS and ATLAS.

For the Simulation project the main achievement has been the delivery of the new public Beta release of Geant4, Geant4 9.3-beta, announced in early June as planned. The new release provides to experiments a preview of some new features and physics configurations to test and attempting to improve the smoothing in energy response in the transition region between the different hadronic models. The migration to gcc-4.3.2 has been completed for most generators in GENSER; also, adoption of 'autotools' for building packages in GENSER has been endorsed; a new HepMC version (2.05) has been released, now providing the possibility to store generated cross section on an event by event basis, and improvements to the I/O streaming. Progress has been made in unifying EvtGen and a new version of the package is expected soon.

New versions of CORAL, COOL and POOL have been released to adapt to the changes in the configurations. The main achievement of Persistency Framework development has been the first release of the CORAL server components. These include for the time being read-only functionalities and passed the first offline validation tests for the ATLAS HLT use case. It will now be tested more extensively in the experiment control room. An enhanced version supporting secure authentication using Grid certificates and VOMS authorization is undergoing some final tests and configuration fixes and will be included in an upcoming release.

4. Site Reliability

The reliabilities for the last 6 months for CERN and the Tier 1 sites are shown in Table 2. The apparent problems for BNL are in fact reporting issues as the BNL site was moved from being associated to the EGEE measurement system to that of OSG.

In addition to the general reliability testing reported in this table, the experiment-specific measurements are now also published monthly together with the general reports. The reliability of the sites as reported by the experiment tests are also reported weekly at the Management Board through maps such as those illustrated in Figure 4. Each box represents 1 day, and the colours represent the availability of the site according to the experiment tests. This example shows the availabilities over the summer period.

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.

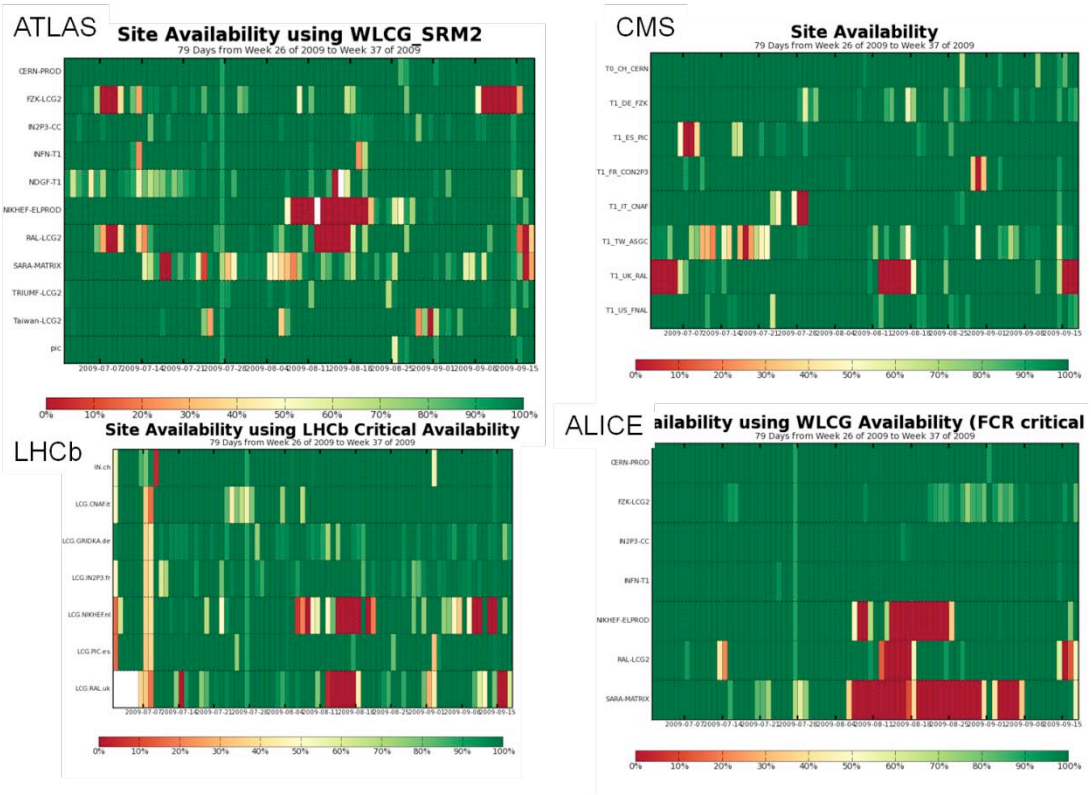

Figure 4: Site availability as measured by the experiments

Table 2: WLCG Site Reliability

Average of the 8 best sites (not always same 8)

Apr-09	May-09	Jun-09	Jul-09	Aug-09	Sep-09
99	99	98	99	99	99

Average of ALL Tier 0 and Tier 1 sites

Apr-09	May-09	Jun-09	Jul-09	Aug-09	Sep-09
98	96	92	99	96	96

Detailed Monthly Site Reliability

Site	Apr-09	May-09	Jun-09	Jul-09	Aug-09	Sep-09
CA-TRIUMF	99	98	99	99	99	99
CERN	99	100	93	99	100	100
DE-KIT (FZK)	97	97	93	99	99	98
ES-PIC	98	96	99	99	98	99
FR-CCIN2P3	99	97	97	99	99	89
IT-INFN-CNAF	98	92	98	100	97	99
NDGF	96	98	96	98	87	97
NL-T1	98	82	99	98	92	89
TW-ASGC	89	96	86	97	100	83
UK-T1-RAL	100	99	99	100	93	93
US-FNAL-CMS	100	100	99	100	100	100
US-T1-BNL	100	100	56	0	92	100
Target	*97*	*97*	*97*	*97*	*97*	*97*
Above Target (+ >90% Target)	10 +2	8 +4	7 +3	11 +0	8 +4	8 +3

Colours: **Green > Target** **Orange > 90% Target** **Red < 90% Target**

5. 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.

Migration of sites to Scientific Linux 5 (or equivalent). This has taken longer than anticipated, largely due to complex dependencies on compiler versions and other external software packages. This has been more of a problem for the experiment software than the middleware, which has been available in the version for the worker nodes for some time. In September, when it was finally thought that there were no more showstoppers for the experiments, the Management Board issued a statement

encouraging sites to migrate their resources to SL5 as rapidly as possible. This migration is ongoing with several of the larger sites having migrated all of their resources.

Reporting of installed capacity. The agreement on the mechanisms for automated reporting of installed CPU and disk capacity was mentioned at the last RRB. Almost all sites are now reporting information, but the validation of that data is still ongoing, and this will take some time. Tools are now available to allow sites and the project management to compare reported information with pledges and accounting.

User level accounting reporting was waiting for completion of the policy to allow data publication. This policy has now been approved; sites are requested to enable the publication of the information into the accounting system.

Metrics. There are several milestones associated with improving metrics, particularly around the area of Mass Storage. Metrics for tape system performance are now available for most Tier 1s in a common format, and have been very useful during STEP'09.

The switch to the new CPU benchmark was reported at the previous RRB. All requirements and pledges are now expressed in these units. The process of sites re-benchmarking their existing capacity is ongoing.

6. Update of Experiment Requirements

Following the discussion at the last RRB where the differences between the experiment requests and the scrutiny group reports were unresolved, a review of the resource requests by the LHCC and scrutiny group was held on July 6. Following that, the LHCC and scrutiny group have worked together with the experiments to arrive at a common understanding of the requirements for 2010. The outcome of this work are the tables of resource requests listed on the WLCG web site: http://lcg.web.cern.ch/LCG/resources.htm. Their results also clearly state that availability of computing resources should not hinder the LHC physics programme.

It should be noted that this process for 2010 has only finally concluded at the very end of September, and that there had been some errors in the numbers stated as the ALICE requests in earlier iterations. The summary table on the WLCG web page and the final C-RSG report at the end of September should be used as the definitive source of information.

With the delay in the LHC, it had earlier been agreed that the 2009 resources need only be available in October 2009 in preparation of the restart of the accelerator. While the C-RSG also commented on the resource requests for 2009, this was already too late in the procurement cycle to affect the procurements and so the full pledged resources for 2009 should be available for first data taking. Note also however, that the experiments have for the most part provided a schedule by quarter of their desired ramp-up of resources and in some cases the sites have taken advantage of this to reschedule their resource deployments accordingly.

The C-RSG also recommended that the 2010 installation deadline exceptionally should be delayed from April 1st until June 1st. This proposal has been agreed by the WLCG Management Board, and will help the procurement process in many cases. In addition the experiments are providing resource requirement profiles broken down by quarters which also helps the sites in planning the installation of equipment. This is particularly important in 2010 as this will happen during data taking.

7. Planning for 2010 and later

Tier 0

The decision on going ahead with the construction of a new Tier 0 centre in Prévessin has been suspended pending a better understanding of costs for container-based solutions and of the long term computing requirements. The delay of the LHC startup has given a little more leeway in providing a long term solution, although there still remain some important issues that must be resolved soon.

There are two distinct, but related issues with the existing CERN Computer Centre. The first is the overall limitation of electrical power available. The second, rather more urgent problem is the available capacity of critical (or backed-up) power. Today this is limited to some 300 kW, all of which is used to provide reliable power for essential equipment (core networking, database services, etc.). Demands for additional reliable power cannot be met by the existing installation. With the rather aggressive replacement of older equipment with newer more power efficient machines, together with savings gained by delaying installations in 2009 following the change in the LHC schedule have meant that the overall power situation has eased slightly for the moment.

It is clear that a better understanding of long term needs for power cannot be improved before there is some experience with real data taking. However it should also be noted that the existing projections do not take into account any additional computing requirements coming from new experiments at CERN.

The strategy now is to understand how to rapidly be able to provision additional capacity before proceeding with a multi-year construction project. Two aspects are being investigated: using container-based capacity, and potential procurement of a hosting agreement in a data centre in the Geneva area. In addition the critical power infrastructure in the existing building will be upgraded to ensure up to 480 kW of diesel-backed load.

The container strategy initially will consider one container to provide additional backed-up power to relieve the problem in that area, and an additional container (or two) to provide capacity for physics services where there is no need for diesel-backed power. A hosting agreement locally will provide a better understanding of the costs involved in such a solution as well as experience in running a remote data centre.

It should be noted that an extended period of using either containers or hosting equipment in a commercial data centre are likely to add significantly to the overall costs if it turns out that eventually a new facility has to be built. Also remote centre operations all carry serious risks that are not understood at this point and that should not be taken lightly.

EGEE to EGI transition

Since the last RRB meeting the preparations for the future European Grid infrastructure have continued. The calls for proposals under the European Commission Framework 7 program have been issued and the EGI communities have been preparing proposals accordingly. For what concerns WLCG there are several important proposals that can be of direct benefit to WLCG in Europe.

The main EGI proposal itself can be expected to provide support for the WLCG Tier 1 and Tier 2 sites, and its scope should include the majority of the core grid services that WLCG relies on. There is a second part of this proposal, responding to a specific funding call in order to support existing multi-national heavy user communities such as WLCG. This part is intended as a transition task expected to last for the first project phase only. This activity covers several of the tasks that the experiments

rely upon including support for the dashboards and Ganga. It also includes experiment-specific tasks for each of the LHC experiments focussed on integrating and supporting the experiment software with grid middleware. This includes for example data management services built upon the general grid data services.

The second area important for WLCG is in response to the funding supporting Virtual User Communities. In the EGI world this means the Specialised Support Centres (SSC) defined as part of the blueprint emerging from the design study. There is proposed to be such an SSC for HEP, providing support for the HEP community and WLCG in particular. There are three main areas of support relevant for WLCG in this activity: integration, operations, and distributed analysis.

The third area is that of middleware. A new project (EMI – European Middleware Initiative) is intended to provide the support for the middleware in production today, not only for gLite but also for the other European middleware stacks – ARC and UNICORE. A second goal of EMI is to integrate and rationalise these separate stacks into a common set of services. At the moment this proposal includes support for all of the middleware services that WLCG today relies on.

At the moment it appears that the services and support structures that WLCG relies upon are covered by one or other of these various proposals. Over the coming weeks we must take care to ensure that the details of the various activities in the proposals are sufficient to support the needs of WLCG and that as the proposals evolve nothing is omitted.

The timescale and transition process are as yet somewhat undefined and are still a cause for concern, particularly as the transition from EGEE to EGI will be during the early months of LHC data taking. The project will continue to monitor these activities and to assist where possible.