



Worldwide LHC Computing Grid Project

Project Status Report

Resource Review Board – 23 October 2007

CERN-RRB-2007-101

Les Robertson

version 2

12 October 2007

This status report covers the period from April to August 2007. 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

During the past six months the WLCG services have been used extensively by the experiments for large scale testing of their computing systems. Each of the experiments plans a final series of tests and “dress rehearsals” over the next nine months to ensure that their distributed computing services are ready for the start of data taking next year and other challenges.

The status report to the April meeting of the CRRB noted that a number of important features in key services had still to be deployed. One of these critical service upgrades remains outstanding: the introduction of a new version of the Storage Resource Manager standard (version 2.2) in the four mass storage systems used by WLCG (Castor, dCache, DPM and STORM). All of the implementations have undergone intensive testing and pilot installations have been installed at a number of sites for testing by the experiments. They were expected to be deployed in production during the third quarter of the year, but this is now scheduled to take place at the Tier-1 and major Tier-2 sites during the fourth quarter, with full deployment to be completed in February 2008.

The growth in the usage of the service over the past year is shown in Figure 1 in terms of the number of jobs run and the cpu used by the LHC experiments. The data covers only sites that are reporting accounting data to the database repository used by WLCG (operated at the Rutherford Laboratory). While this is only a subset of the overall usage it does include CERN, all of the Tier-1s and

more than three quarters of the Tier-2s, together with a number of Tier-3 sites that are connected to the EGEE grid. Full accounting data for cpu and storage usage at CERN and the Tier-1s is available on the web and summarized in CERN-RRB-2007-13.

Accounting for cpu usage has been extended to the Tier-2s from September, and will be presented formally at the next (April 2008) meeting of the CRRB. Figure 2 shows the distribution of the cpu usage reported by the Tier-2 Federations for September. A total of 45 of the 53 federations included in the WLCG MoU are now reporting accounting data, representing 113 different sites. Ten of these sites account for 50% of the cpu usage, and 24 sites for 90% of the total usage. The cpu usage accounted in September is equivalent to 43% of the commitment pledged by the 53 Federations for the month in the WLCG MoU. Figure 2 also shows the distribution of Federations according to the degree to which usage approaches the 2007 commitment. In only 15 cases does the usage exceed 70% of the commitment.




Figure 1 - Jobs Run and CPU used at Sites reporting accounting data

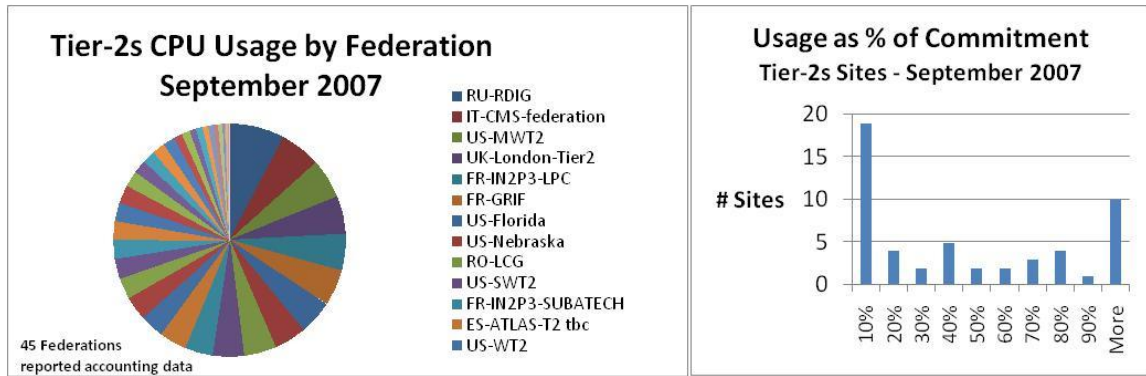


Figure 2 - Distribution of Tier-2 CPU Usage - September 2007

The data transfer services between the different tiers and sites has been the subject of testing by all experiments over the past six months. Figure 3 shows the evolution of CMS data transfer testing over the past twelve months, now including 7 Tier-1s, 49 Tier-2s and 13 Tier-3s. The sustained data rate over recent months is close to that required for operation, although not yet at the full level of complexity.




Figure 3 - CMS Data Transfer between all Sites

Atlas carried out an end to end test of their computing system at the end of August as part of the "M4" cosmic run. This took the events from the data acquisition system through the Tier-0 facility and out to the Tier-1s, Tier-2s and finally to a number of Tier-3s, with processing occurring at each level. The input data rate was up to 250 MB/s. Figure 4 shows the time delay between data being registered at CERN and its successful transfer to Tier-1s. In general this progressed smoothly with problem cases (gaps in the main line) being recovered automatically once the problem was resolved.

At the WLCG workshop in Victoria at the beginning of September it was agreed to schedule a "Combined Computing Readiness Challenge"




Figure 4 - Data Distribution during the ATLAS M4 Test

(CCRC) involving the four experiments and all major sites during the first half of next year, prior to the first collisions. A coordination team has been formed that will prepare a detailed plan, probably with two test periods in February and May 2008. It is undoubtedly complicated to coordinate all the components required for this within one experiment, even more so when trying to schedule such a test across four experiments and around 100 sites, but this will be the normal and continuous operational situation from the middle of 2008.

2. Applications Area

The main emphasis during the last six months in the Application Area has been the consolidation of the software functionality that is in routine use by the LHC experiments. Major improvements have been made in the release process. In particular, the software validation by the experiments has been very much facilitated by introducing the nightly build system.

Simulation and Validation - A new version 9.0 of the Geant4 simulation toolkit was released at the end of June. This version includes performance and algorithmic improvements as well as new functionality. Work on Physics Validation has continued and preliminary results from ATLAS and CMS test-beams consistently show better agreement with the data for hadronic showers. Another important activity has been the Fluka extension of the ATLAS TileCal 2002 analysis, which is aimed to benchmark both Geant4 and Fluka simulations in high-energy calorimetry.

Core Libraries - A number of patch releases of ROOT software have been produced and at the same time new functionality and improvements have been developed with releases in June and August. In particular, the repackaging and modularization of the ROOT core libraries have had a major impact on the size of the ROOT executable module and its start-up time. The merge of the CINT (C interpreter) and Reflex (dictionary) packages is progressing well, and an opportunity to make more fundamental changes in CINT, such as the support for multi-threading, has been taken. New developments have also been made in the graphical user interface, visualization and mathematical packages.

Data Persistency – Improvements have been made in the data persistency framework (POOL) and conditions database (COOL) to improve scalability and performance and these are being tested with realistic conditions data workloads.

MC Generators - All the Monte Carlo generators requested by the LHC experiments have been moved to the agreed new GENSER structure and have been integrated within the software frameworks of the experiments. Several new tests have been implemented and are now used on a regular basis for the validation of the generators.

3. Site Reliability

The results of the site reliability metric for CERN and the Tier-1s for the past 6 months are summarised in Table 1. Fuller data for each site is available from the [LCG Planning Page](#). The Nordic Data Grid Facility (NDGF) and the Open Science Grid (OSG) are developing equivalent sets of tests, adapted to their local needs, which will enable NDGF and the OSG Tier-2 sites to be included in the monitoring and reporting system.

The site target reliability level was 88% to the end of May, and 91% from June. The project target for the eight best sites each month was 88% to the end of June, and 93% from July. The project target was achieved in 4 of the past six months, with an average of 91%, which is just within the six-month target. The evolution of the reliability since May 2006 is shown in Figure 6, showing a clear improvement of more than 10 percentage points over the period for the best eight sites, and a similar, though erratic, improvement for the average of all sites.

Site Reliability - WLCG Tier-1s + CERN								average last six months
		Mar-07	Apr-07	May-07	Jun-07	Jul-07	Aug-07	
site								
CERN-PROD		97%	96%	90%	96%	95%	99%	96%
FZK-LCG2		75%	79%	79%	46%	75%	67%	70%
IN2P3-CC		58%	95%	94%	88%	94%	95%	87%
INFN-T1		76%	93%	87%	67%	82%	70%	79%
RAL-LCG2		80%	87%	87%	87%	98%	99%	90%
SARA-MATRIX		47%	92%	99%	75%	92%	86%	82%
TRIUMF-LCG2		70%	73%	95%	95%	97%	97%	88%
Taiwan-LCG2		95%	92%	98%	80%	83%	83%	89%
USCMS-FNAL-WC1		90%	85%	77%	77%	92%	99%	87%
PIC		96%	95%	77%	79%	96%	94%	90%
BNL-LCG2		6%	89%	98%	94%	75%	71%	72%
average reliability		72%	89%	89%	80%	89%	87%	84%
target site reliability		88%	88%	88%	91%	91%	91%	90%
target 8 best site reliability		88%	88%	88%	91%	93%	93%	90%
8 best sites	availability	85%	91%	93%	87%	93%	94%	91%
	reliability	85%	92%	94%	87%	93%	94%	91%
average	% of target	97%	105%	107%	96%	102%	103%	102%
# sites ≥ target		4	7	6	3	7	6	4
# sites ≥90% target		5	10	9	5	9	8	8
site average colour coding: < 90% of target ≥ 90% of target ≥ target								

Table 1 - Site Reliability Summary - March - August 2007

From October we will also begin to report reliability metrics for Tier-2 sites, using the same test set. For information the results of these tests during this year at 83 of the Tier-2 sites that are connected to the EGEE grid are presented in Figure 5. It is interesting to note that the average of the best 50% of the sites during this period is similar to that of the best eight Tier-1s.




Figure 6 - Site Reliability Metric - CERN + Tier-1s - May 06 to Aug 07

Figure 5 - Site Reliability Metric - Tier-2 Sites - Jan-Aug 07

4. Ramp-up of Resources and Performance

The following figures illustrate the ramp-up required at CERN and the Tier-1 centres to meet the capacity and performance targets for the 2008 run. It is important that these resources are fully available in production by April for the final test sessions a few months ahead of the first beams. This installed capacity at these sites must be increased on average by a factor of 3 or 4 during the coming six months. In several cases the ramp-up required is much steeper as some sites have not yet installed the capacity committed for 2007.




Figure 7 - Resource Ramp-up required for 2008 at CERN and Tier-1 Centres

The situation for CPU at Tier-2 sites is shown in Figure 8, comparing the current cpu usage with the commitments for 2007 and 2008. A factor of four increase is required in the delivered CPU prior to April next year.



Figure 8 - Resource Ramp-Up required at Tier-2 Centres

5. Long-Term Evolution of Requirements and Pledges

During the summer the four experiments revised their planning estimates for resources at CERN, the Tier-1s and Tier-2s, taking account of the machine startup schedule and using the experience of data sizes, programme performance, etc. gained from this year's testing. The detailed numbers for 2008-2012 are given in the resource report to this meeting, CERN-RRB-2007-13,

Figure 9 summarizes the data graphically across the four experiments, showing the growing discrepancy over the next few years between the requirements and the capacity planned by the funding agencies at the Tier-1s and particularly at the Tier-2s. While it is expected that some of the shortfall is due to delays in the national planning processes, there is reason to be concerned by the apparent lack of planned investment at Tier-1 and Tier-2 centres after 2010.

Experiment Requirements and Funding Agency Pledges/Plans - October 2007

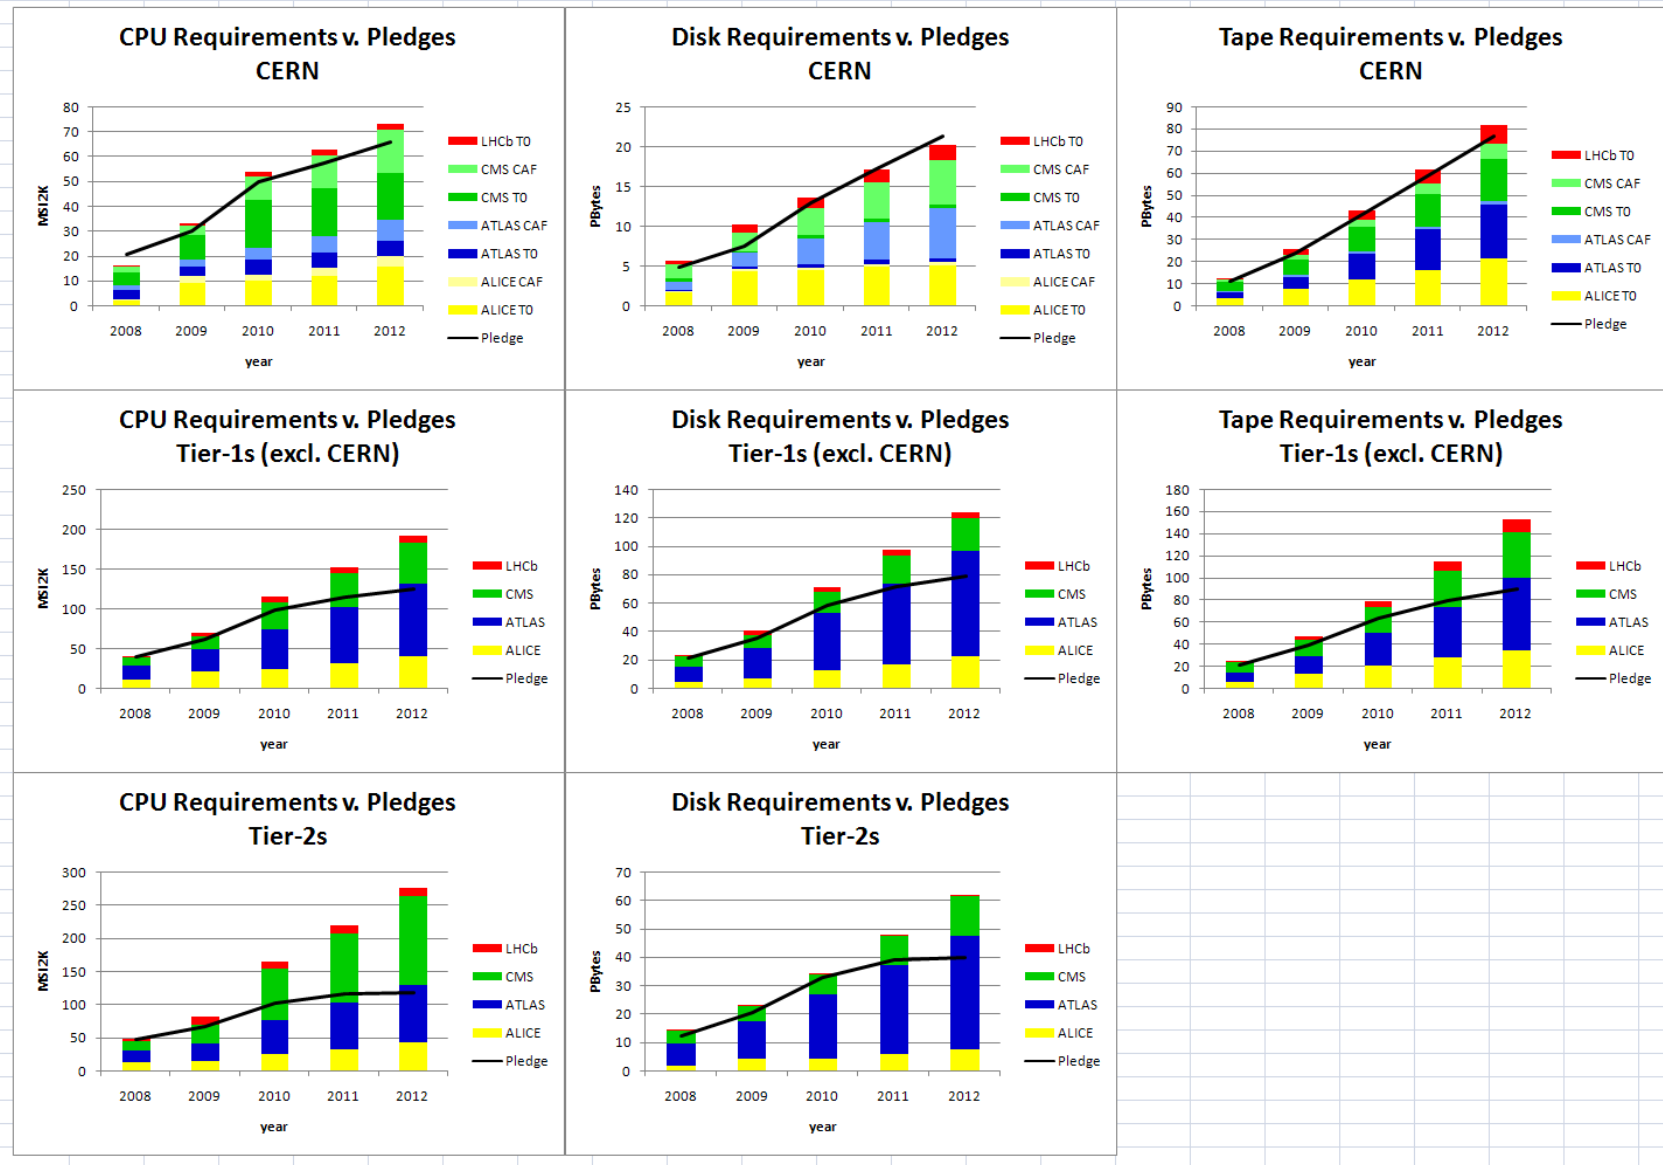


Figure 9 - Requirements and Pledges - 2008-2012