



Computing Resource Review Board – 24 April 2007 Status of the LCG Project

This status report covers the period from October 2006 to March 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.

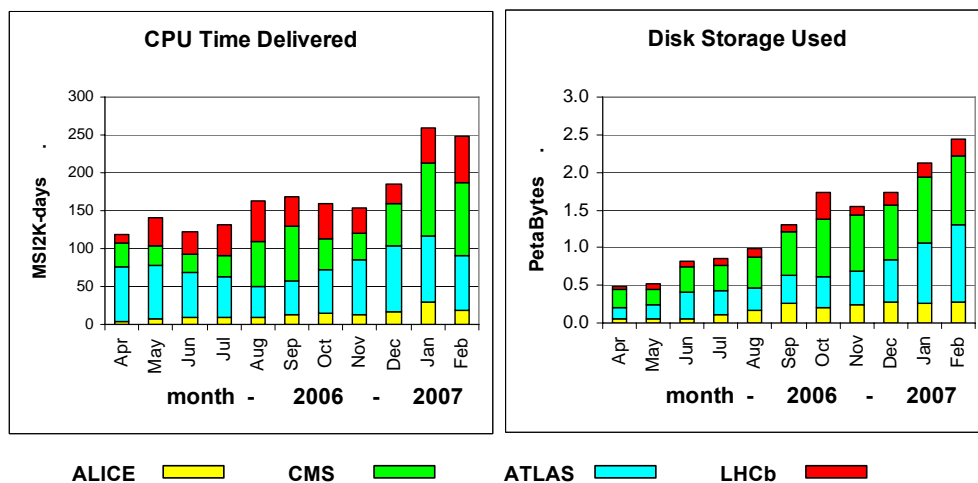
1. The LCG Service

By the end of Service Challenge 4 (SC4) in October 2006 peak data transfer rates between the participating sites had been achieved that satisfy the full nominal requirements for 2008, with some sites achieving double or even triple those required. While these rates were achieved under test conditions, considerable progress was also made in both data rates and stability using experiment-driven transfers which more accurately reflect real-life use cases. For example CMS was able to achieve more than the 150MB/s target set for SC4, and ALICE was also able to demonstrate peak rates above their heavy ion requirement. While this showed that the basic software and networking services are sufficient to handle the data throughput required in 2008, a number of shortcomings in the software and storage management operations were identified that are being tackled in the first half of 2007.

The service has continued to run uninterrupted as it transitioned from the final pilot phase (SC4) to a permanently available service. A first test of the new service was that it ran over the Christmas holidays, when many sites had very much reduced operations support, without any significant interruption and with relatively high job throughput rates.

The usage of the service has grown over the past year as shown for CERN and the Tier-1 sites in the diagrams below, with cpu usage increasing by a factor of about 2.5 and disk usage by a factor of 5. The drop around November 2007 was due to a large part of the CERN equipment being taken for ATLAS data acquisition tests.

LCG Service - CPU and disk usage at CERN and the Tier-1s



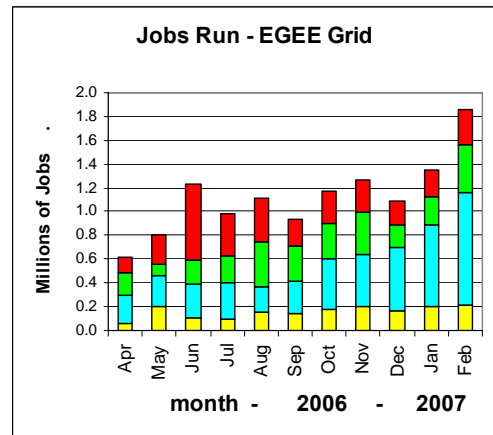
It is planned to begin formal reporting of accounting data for Tier-2 sites during this year and the first results for the full set of Tier-2 sites will be reported at the October meeting. However, an

indication of the growth in usage by the LHC experiments of the grid in general can be seen from the graphic showing the number of jobs run on the EGEE grid at all sites that are reporting their accounting data to the central repository. The current job rate of 1.8M jobs per month will have to increase by a factor of about 5 when the accelerator is fully operational in the second half of 2008.

A system to measure job reliability is under development and is being used within ATLAS and CMS as part of the experiment dashboards for some of their workloads. A set of metrics is being developed to track reliability as seen by the end-user for specific sets of jobs.

From July 2007 all of the experiments plan to run full tests of their computing systems, integrating the complete set of applications software components with grid and site services.

The preparations for these are going on now, uncovering several problems which have to be resolved during the next few months, particularly in the area of data and storage management.



2. Applications Area

New versions of the Applications Area (AA) packages, in particular ROOT, GEANT4, CORAL and COOL, were provided at the end of the year and are currently being integrated by the experiments as the base versions that will be used for the first accelerator run. Several LCG configurations have been made available to LHC experiments and are being used in their latest releases. The nightly build system for the LCG software stack has been put into production. It allows the build of all LCG AA projects on most of the supported LCG platforms. The system is used by the experiments for validation and integration of changes and new features.

Important deliverables have been completed for the different Simulation sub-projects. Two LCG notes (one on the results of the 3rd simple benchmark analysis, and one on the studies in Geant4 related to the simulation of hadronic shower shapes) have been published. In particular, the joint investigations performed by the Physics Validation and Geant4 teams to understand the key components responsible for the development of the hadronic shower shapes, has already produced good positive results from the preliminary developments made available in the last Geant4 development releases. The Generator Services sub-project after a re-organisation initiated in the second half of 2006 is now stable and progressing well. The new GENSER release follows the structure and configuration policy defined in December 2006 in agreement with the Monte Carlo generator developers and is now available and used by the experiments.

Improvements to the Common Relation Abstraction Layer (CORAL) have been delivered such as authentication functionality based on the LHC File Catalog (LFC), thread safety and access to stored procedures. The Conditions Database (COOL) version 2.0.0 is almost ready for release. It includes a new API for user payload specification, a port to the AMD64 architecture and new locking and 'dynamic replication' functionalities. This new version is currently being tested for integration by ATLAS and LHCb since it requires some changes in the DB schema and API.

ROOT developments include many new functionalities and improvements in most areas. The work of re-engineering parts of ROOT to integrate with the LCG dictionary (Reflex) is progressing steadily. The 7th ROOT workshop took place from 26-28 March at CERN with a participation of about 80 people. Development of the PROOF interactive parallel ROOT facility has focused on the ALICE analysis use cases. The developments included an extensive monitoring facility to track performance on the ALICE analysis cluster at CERN. ALICE measured a very satisfying speedup of the analysis with high cluster usage efficiency.

3. Site Reliability

The results of the site reliability metric for CERN and the Tier-1s for the past 6 months are summarised in the following table. Fuller data for each site is available from the [LCG Planning Page](#). The Nordic Data Grid Facility (NDGF) does not yet take part in the system. The Open Science Grid (OSG) in the United States has begun a project to develop a set of tests adapted to their local needs, which will replace the current set of tests for OSG sites, including FNAL and BNL. In the case of BNL, the lengthy period of unavailability in February and March was due to a system upgrade that led to the failure of tests using a service that is not required by ATLAS, the only experiment to use the site.

Site Reliability - WLCG Tier-1s + CERN																	
The table shows <i>availability</i> for May through August, and <i>reliability</i> from September on BNL included in average from November																	
Reliability = Availability/Scheduled_Availability (Scheduled_Availability=(1-Scheduled_Down_Time); tests are not run while scheduled down)																	
	CERN- PROD	FZK- LCG2	IN2P3- CC	INFN- T1	RAL- LCG2	SARA- MATRIX	TRIUMF- LCG2	Taiwan- LCG2	USCMS- FNAL-WC1	PIC	BNL- LCG2	average reliabilities	8 best sites average			# sites ≥ target	# sites ≥90% target
												availability	reliability	(% target)			
Oct-2006	96%	54%	85%	85%	77%	74%	80%	67%	55%	84%	27%	76%	80%	81%	92%	1	5
Nov-2006	90%	85%	62%	94%	87%	77%	87%	95%	77%	79%	56%	81%	86%	87%	99%	3	7
Dec-2006	93%	63%	22%	77%	86%	82%	91%	95%	79%	90%	70%	77%	86%	87%	99%	4	6
Jan-2007	99%	85%	96%	75%	80%	93%	79%	96%	84%	86%	90%	87%	91%	91%	103%	5	9
Feb-2007	91%	90%	74%	93%	82%	83%	88%	97%	67%	86%	57%	82%	88%	89%	101%	5	8
Mar-2007	97%	75%	58%	76%	80%	47%	70%	95%	90%	96%	6%	72%	85%	85%	97%	4	5
average last three months	96%	83%	77%	82%	80%	75%	79%	96%	80%	89%	54%	81%	--	88%	100%	3	8

The target for the Site Reliability metric set for the period until June 2007, 88% as the average of the eight best sites, has been met on average for the past three months. However, only five sites have achieved the target more than once during the 6-month period of this report. During the last three months three sites have achieved the target on average (two sites at the time of the October meeting), and a further four sites (three in October) have come within 10% of the target, while the average for all eleven sites is 81% (73% in October). The situation has clearly improved since October, but there is some way to go to achieve the 93% target for the eight best sites set for the second half of the year.

In order to help sites improve their reliability a series of working groups has been started to share experience in and to propose standard tools for monitoring site and grid services.

4. Critical Services that have not been fully deployed

The main focus for the rest of this year will be on improving stability of the software and services, and gaining the experience needed for sustained reliable operation. The second priority is expanding the capacity of the services while increasing their performance to reach the levels necessary for the first physics runs. In general the development and deployment of new functionality have lower priority during the commissioning period, but there are several key developments which are scheduled for deployment during the first half of the year.

Distributed Database Deployment - 3D Service

Eight of the participating sites (ASGC, BNL, CNAF, GridKA, IN2P3, NIKHEF/SARA, RAL, TRIUMF) are now taking part in the ATLAS and LHCb conditions database tests using the distributed Oracle database service. The two remaining sites, PIC and NDGF, are in the process of joining and it is expected that all of the sites will be available for experiment production by the beginning of May, even if some of the sites will not have the full planned capacity.

SRM v2.2

As reported at the last meeting of the C-RRB an agreement on a new version of the Storage Resource Manager standard (version 2.2) was reached at a workshop at Fermilab in May 2006. The developments required for the three mass storage systems used by WLCG (Castor, dCache and DPM) were completed by the end of the year, but problems discovered during testing have

delayed the deployment of the Castor and dCache implementations. This was foreseen for the first quarter of 2007, but they are now expected to be available in production at a small number of sites by the end of the second quarter, and more widely deployed only during the third quarter.

File Transfer Service (FTS)

A new version of the FTS system, with improved monitoring and administration tools and support for SRM 2.2 was installed for testing at CERN in January and is on schedule to be in production at CERN at the end of April, then progressively deployed to Tier-1 sites during May.

VOMS-based scheduling priorities

This concerns the deployment of a system whereby the group and role of a user as defined by the experiment is used by the local site scheduling system in assigning priorities and sharing resources. Base software components have been delivered by EGEE and deployed at most of the Tier-1s for beta testing prior to being included in a formal release. Each site has to adapt the system to the details of the local configuration and batch system, but it is expected that the service will be available at EGEE Tier-1 sites by the end of April.

Scientific Linux version 4 (SL4) Migration

Most of the computer systems being installed by regional centres since the last quarter of 2006 are not certified for the version of the Linux operating system that is supported by the middleware package used by EGEE sites, and support for security updates for this version is scheduled to stop in October 2007. The availability of a middleware distribution using a later version of the operating system, Scientific Linux version 4 (SL4), is therefore urgent. There have been some technical difficulties in preparing the full SL4 port, but it is expected that this will be available during May.

5. Level-1 Milestones

A revision of the high level planning and reporting process has been completed, adapted to the operational phase of the project. High level milestones are now defined as targets to be met by each of the major sites (CERN, Tier-1s). Individual milestones are also defined for each site, activity area (applications, ARDA, deployment, ..) and experiment. The [full set of milestones](#) can be found on the LCG Web. The status of Level-1 milestones due since the last Resource Review Board meeting is summarised below. The site-specific status is colour-coded as follows:

achieved	not achieved < 1 month late	not achieved > 1 month late
----------	--------------------------------	--------------------------------

Level 1 Milestones due October 2006 to March 2007					
ID	Date	Milestone			
DBS-1	30.09.06	Full LCG database service in place Completed by the end of the year for Frontier/Squid (CMS Tier-1s) and for Oracle/Streams at 6 of the 10 Tier-1s concerned (ASCC, BNL, CNAF, GridKA, IN2P3 and RAL), all of which have been tested at the replication rates estimated to be required by ATLAS and LHCb. A new milestone (WLCG-07-09) has been defined for the final step in getting the Oracle/Streams into production including the remaining Oracle Tier-1s (NIKHEF, NDGF, PIC, TRIUMF).			
IS-1	30.09.06	Initial LHC Service in Operation Began at the end of SC-4. The service will be used for extended testing of the computing systems of the four experiments, for simulation and for processing of cosmic-ray data.			
WLCG-07-01	28.02.07	24 X 7 support defined at CERN and Tier-1s Definitions of the level of support provided outside of normal working hours, including processes for monitoring and reporting problems, according to the urgency.			
		ASGC	IN2P3	CERN	FZK
		PIC	RAL	NIKHEF	TRIUMF
				INFN	NDGF
				BNL	FNAL

Level 1 Milestones due October 2006 to March 2007							
ID	Date	Milestone					
WLCG-07-08	31.03.07	Accounting Data published into the APEL repository					
		CERN and Tier-1 sites automatically publish their CPU accounting data to the central repository at the Grid Operations Centre (GOC) provided by Rutherford Lab. This data will be used as the sole source for the accounting reports from May 2007.					
		ASGC	IN2P3	CERN	FZK	INFN	NDGF
		PIC	RAL	NIKHEF	TRIUMF	BNL	FNAL
WLCG-07-09	31-03-07	3D Oracle/Streams Service in Production					
		Oracle services in production and certified by the experiments.					
		ASGC	IN2P3	CERN	FZK	INFN	NDGF
		PIC	RAL	NIKHEF	TRIUMF	BNL	
		NIKHEF has been certified by ATLAS, and is awaiting certification by LHCb. TRIUMF is awaiting certification by ATLAS.					