



Computing Resource Review Board – 24 October 2006

Status of the LCG Project

This status report covers the period from April 2006 to September 2006. 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. Service Challenges

Service Challenge 4 (SC4) began as scheduled at the beginning of June, as an evolution of the production services in place at that time using LCG resources connected to the EGEE and OSG grids. In the case of the EGEE sites, this coincided with a new release of the basic middleware – gLite 3.0. There had been concerns over the state of readiness of this release but the deployment went smoothly and the middleware has proven to have a good level of reliability.

Data distribution tests carried out prior to the start of SC4 in April between CERN and the Tier-1s have shown that the service can now sustain over an extended period a transfer rate of about 80% of the 1.6 GBytes/sec rate that will be required in 2008 (see Figure 1). ATLAS and CMS have further exercised the data transfer service with their computing model testing during SC4 under more realistic operating conditions and have shown sustained operation at the level of a PetaByte per month (3 PB of data between Tier-1s and Tier-2s for CMS over a 90-day period, 1.25 PB of data between CERN and Tier1s for ATLAS in a two-month period).

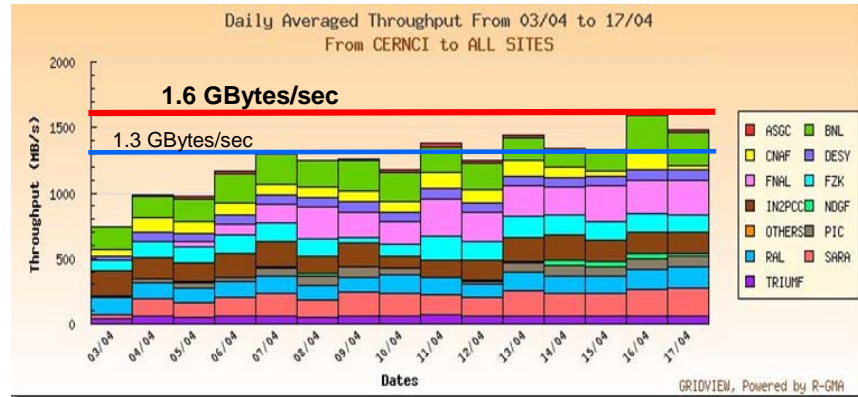


Figure 1 - Daily average throughput - CERN-Tier-1s

SC4 operates across the two major grid infrastructures: EGEE and OSG. Over the past year the number of jobs executed each month has grown significantly. On the EGEE grid an average of 50K jobs were run each day during the experiment test periods in June and August, compared with 12K jobs per day in June 2005 (see Figure 2). Up to 15K jobs per day run on OSG,

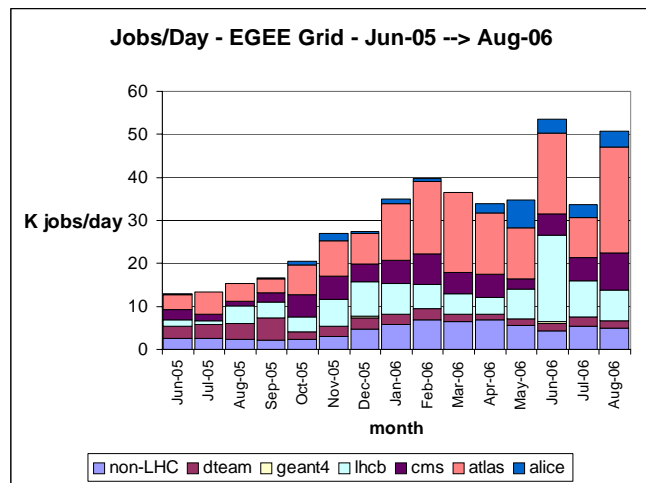


Figure 2 - Evolution of jobs/day on EGEE Grid June 05 through August 06

and these two grids together regularly sustain 17K simultaneously executing jobs.

The experiment activities during SC4 have tested production readiness significantly more than in previous test periods. Many problems have been identified and resolved, and work still has to be done to bring all sites up to target performance and reliability levels. Overall, however, it is clear that there is now a baseline service covering almost all of the major LCG sites that is being used effectively by the experiments.

2. Applications Area

The projects in the Applications area continue to support the experiments in their preparation of the software releases that are being used in the various data challenges and productions this autumn. Several iterations of the software packages (ROOT, CORAL, POOL, COOL, etc.) have been made available in various configurations to allow experiments to integrate the new functionality, provide feedback and be ready with production quality releases by this summer. Special emphasis has been put in the optimization of this iterative process. For that, a new set of procedures for testing and building the software is being put in place to optimize the time that it takes for the experiments to integrate the changes and bug fixes in libraries provided by the Application Area.

In the ROOT project strong development is taking place in the integration of the C++ interpreter (CINT) with the C++ reflection system (Reflex). It is planned to release the new version of the system this quarter. The mathematical libraries have been consolidated and additions have been added concerning Fast Fourier Transforms and Multivariate Analysis. Many developments are currently ongoing in the PROOF system as the result of the serious testing being done by ALICE in the context of their prototyping of the CERN Analysis Facility (CAF).

The POOL/CORAL project has been consolidating the generic relational database management system (RDBMS) interface for Oracle, MySQL, SQLite and FroNTier. New functionality has been developed for improving the overall reliability of user applications with database back ends. In addition, the COOL project (conditions database) has been improving the versioning capabilities by the use of tags and hierarchical tags.

The Simulation project is putting considerable effort on the study of hadronic shower shapes, to understand the discrepancies observed between simulation and test-beam data. Also comparisons between Geant4 and FLUKA simulation packages are being made with the help of the set of tools that has been developed to facilitate this task.

A new version of Geant4 has been released. It includes a new precise elastic process for protons and neutrons. It is particularly relevant to improving the accuracy of energy deposition in scintillators. In addition, the new version includes an efficient method to detect overlaps in a user's geometry and updated particle definitions to match with the Particle Data Group's PDG-2005.

3. Site Reliability

The algorithm for computing site availability was agreed by the Management Board at the beginning of April and measurement started for CERN and the Tier-1 sites from the beginning of May. In July it was decided to measure also *reliability*, defined as the time that the site passes a set of standard tests as a percentage of the time that it is scheduled to be available. *Availability* is defined as the percentage of the time that the site passes the tests, taking no account of scheduled down time. The results from May through August are summarised in Table 1. Fuller data for each site is available from the [LCG Planning Page](#). BNL and NDGF do not yet take part in the system.



Site Reliability - WLCG Tier-1s + CERN

The table shows *availability* for May through August, and *reliability* from September on

	CERN- PROD	FZK- LCG2	IN2P3- CC	INFN- T1	RAL- LCG2	SARA- MATRIX	TRIUMF- LCG2	Taiwan- LCG2	USCMS- FNAL-WC1	PIC	average - all sites	8 best sites		target
												availability	reliability	
May	89%	85%	83%	89%	68%	58%	77%	87%	68%	61%	77%	81%	-	88%
June	92%	15%	89%	62%	76%	49%	88%	75%	64%	88%	70%	79%	-	88%
July	90%	54%	87%	31%	73%	84%	80%	98%	20%	87%	70%	82%	83%	88%
August	95%	69%	94%	69%	59%	83%	87%	97%	4%	88%	74%	85%	86%	88%
September	89%	57%	83%	88%	69%	92%	66%	85%	26%	90%	74%	80%	82%	88%
average	91%	57%	87%	68%	69%	74%	79%	89%	36%	83%	73%	80%	80%	88%

target (90% of MoU)	88%	# sites meeting target	2	8 best site avge	91%	full period
90% of target	79%	# sites > than 90% of target	5	as % of target	93%	latest month

Table 1 - Summary of CERN + Tier-1 Reliability - May-September 2006

The target availability for SC4 is 88%, to be achieved by at least 8 sites. Only two of the ten participating sites achieved the target on average during the four month period, with a further three coming within 10% of the target. The average for all 10 sites is 73%. Taking the 8 best sites in each month (different sites in each month) the average hovers around 80%, and in some individual cases the situation has deteriorated rather than improved.

All of the Tier-1 sites are now producing their plans for 24 X 7 operation next year. In general this will be achieved by an on-call service initiated through automated monitoring systems. In some cases these services will depend on experts being available on a "best efforts" basis while other sites will have formal arrangements to have experts on stand-by duty. It is important that we gain experience with these services at an early stage in order to train stand-by staff and establish the effectiveness of the best-efforts services.

4. Critical Services that have not been fully deployed

The main focus for the next nine months will be 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, with the exception of four services that have not yet been fully deployed but have been agreed as essential to the operation of WLCG.

Distributed Database Deployment – 3D

Two technologies are used: synchronised copies of a relational database system (Oracle) at Tier-1 sites participating in the ATLAS and LHCb experiments; a central Oracle installation with web caching of current queries (Frontier/SQUID) at Tier-1 and Tier-2 sites participating in CMS. For the former case, the database systems at the first set of sites (ASGC, BNL, CNAF, GridKA, IN2P3, RAL) are now synchronized with experiment clusters at CERN for direct experiment tests with conditions data. The remaining sites have now nominated database contacts but there will be delays in getting these sites into production.

The Frontier/SQUID systems have been set up and tested at all of the CMS Tier-1s and all but three of the CMS Tier-2s. Multi-client stress tests have been done by CERN to validate the Frontier back-end installation at the Tier-0.

SRM version 2.2

An important step in completing the base storage services for LHC startup is the evolution from the current version of the Storage Resource Manager standard to a newer version (2.2) which has additional features of importance to the experiments. Defining the details of this took longer than anticipated at the time of the last C-RRB meeting, but an agreement between the developers, experiments and sites was reached at a workshop in Fermilab in May, and development of the support for the three mass storage systems used by WLCG (Castor, dCache

and DPM) is now well in hand. However the implementations will not be ready for deployment before the first quarter of 2007.

Full deployment of the File Transfer Services

As explained above considerable experience has been gained running File Transfer Services (FTS) between CERN and the Tier-1s, and between Tier-1s and some of the Tier-2s. A comprehensive set of data has been compiled, using information provided by the experiments, defining for each Tier-1 site the Tier-2 sites with which it will have to exchange data, along with data rates and storage system implications. This information will be used to define the test programme to be used to commission over the next few months all of these data paths.

Job Scheduling using Virtual Organisation groups and roles

Many computing centres are now using well developed resource schedulers to provide fair sharing of the installed computational capacity with predictable job turnaround times, taking account not only of the relative priorities of the experiments but also sharing between different groups within each experiment. Such facilities are even more important with the large number of collaborators on the LHC experiments sharing resources spread across more than one hundred computing centres. Facilities for doing this at the level of the grid are not well-developed, but an initial implementation supporting a small number of groups and roles within each experiment is being tested and is scheduled to be put into production during the first quarter of 2007.

5. Level-1 Milestones

The status of Level-1 milestones due since the last Resource Review Board meeting is summarised in Table 2. Full milestone tables are available via the LCG [Planning Page](#).

Level 1 Milestones due March to September 2006		
ID	Date	Milestone
SC4-1	15Mar06	gLite 3.0 middleware and services available on the pre-production service for beta testing by the experiments. <ul style="list-style-type: none"> The first version of the gLite 3.0 distribution was installed on the Pre-Production Service in mid- March, but there were many problems with the new components and only limited testing by experiments was possible before the end of April. <i>(The software for baseline services was operational at all Tier-1s and more than 20 Tier-2 sites by the end of SC-3. All sites upgraded with the LCG 2.7 distribution early in the year.)</i>
OPN-2	31Mar06	Tier-0/1 high-performance network operational at CERN and 6 Tier-1s, at least 3 via GEANT. <ul style="list-style-type: none"> The high-performance network was operational at CERN and 6 Tier-1 sites. FNAL, SARA, IN2P3, TRIUMF, BNL and CNAF were connected but, as anticipated at the March Overview Board meeting, only two sites (CNAF and FZK) were connected via GEANT.
CAS-1	15Mar06	Castor2 Readiness Review <ul style="list-style-type: none"> The review took place on 6-8 June.
SC3-4	31Mar06	All services on all Tier-1 sites monitored <ul style="list-style-type: none"> The monitoring of the services was in place by the end of April with a new version of the framework (Service Availability Monitor - SAM). This includes a tool to calculate site availability as agreed by the MB (see SC3-5). These figures are published on a monthly basis for CERN and the Tier-1s from the beginning of May.
SC3-5	31Mar06	Proposal on availability levels specified in Annex 3 of the WLCG MoU. Proposal agreed by the MB on 4 April 2006.

Level 1 Milestones due March to September 2006		
ID	Date	Milestone
SC4-3		Service Challenge 4 Set-up: Set-up complete and basic service demonstrated, capable of running experiment-supplied packaged test jobs, data distribution tested.
SC4-3.1	30Apr06	Testing completed on pre-production service <ul style="list-style-type: none"> • Only limited testing was possible before the end of April (see SC4-1), but the Management Board decided to stick to the deployment schedule and the gLite 3.0 distribution was made available for installation at CERN and on Tier-1s at the beginning of May.
SC4-3.2	30Apr06	SC4 disk-disk throughput tests complete <ul style="list-style-type: none"> • Initial T0->T1 disk-disk tests took place, achieving the target aggregate throughput of 1.6 GB/sec for only one day. The average throughput for the last week was 1.3 GB/sec.
SC4-3.3	01Jun06	Middleware and services deployed and in operation at participating sites. Start of SC4 service phase. <ul style="list-style-type: none"> • Services deployed on 1 June at CERN+7 Tier-1s (exceptions IN2P3, FZK, FNAL, NDGF). All sites had installed by the end of June.
DRC-3	30Apr06	1.0 GB/s data recording demonstration at CERN: Data generator -> disk -> tape sustaining 1.0 GB/s for one week using the CASTOR 2 mass storage system and the new tape equipment. <ul style="list-style-type: none"> • Completed March 18.
SC4-4	31May06	Service Challenge 4: Start of stable service phase <ul style="list-style-type: none"> • Service opened on 1 June, with CERN and Tier-1 sites.

Table 2 - Level 1 Milestones due since last C-RRB