



LHC Computing Grid Project

Quarterly Status and Progress Report

2006 Q4

Executive Summary

Alberto Aimar
15 February 2007

1 Introduction

This document highlights major achievements and issues of the quarter, also providing an overview of the individual Quarterly Reports that were submitted by sites, experiments and projects.

All reports were reviewed by a Review Team (A.Aimar, F.Carminati, M.Gronager, L.Robertson and J.Templon). The reviewers then asked for additional information and, in some cases, the reports were modified and re-submitted by the original authors.

For further details about the points summarized in this document, please refer to the complete document: "Quarterly Status and Progress Report - 2006 Q4" which contains all the individual reports, twenty three in total. All Quarterly Reports since 2005Q4 are here: <https://cern.ch/twiki/bin/view/LCG/QuarterlyReports>

2 Level-1 Milestones

2.1 Status of 2006 Milestones

DBS-1: Full LCG database service in place

Status – This milestone is still not fully completed at the end of 2006Q4. It has been completed for all sites that installed the Frontier/Squid framework (CMS Tier-1 sites) and for also for six of the ten Tier-1s concerned with the Oracle/Streams solution.

New milestones for 2007, and a closer monitoring of their progress, have been defined for the remaining Tier-1 sites (NDGF, PIC, SARA-NIKHEF, and TRIUMF) that did not achieve their milestone and that are not yet providing a Distributed Database Service

In particular NDGF and PIC have major delays to face, as described in the 3D section of this executive summary.

IS-1: Initial LCG Service in operation

Status – The "LCG Service", as result of the work on the SC4 Service Challenge, has been shown to deliver usable, and used, services. Started at end of SC4 t should be capable of handling the full nominal data rate between CERN and Tier-1 sites. 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. During the first six months of 2007 each site will build up to the full throughput needed for LHC operation, which is twice the nominal data rate.

The current status of the LCG Service provides a solid basis for consolidation and further hardening. In 2007 new milestones and targets will verify that the experiments use cases are all served adequately.

2.2 Definition of Targets and Milestones for 2007

In order to be ready for data taking in November 2007, the LHC experiments are preparing the commissioning their software by summer 2007. After that time, in order to minimize risk, only improvements and debugging will be performed, but there will not be any major release of new software applications or libraries.

The same constraint clearly applies also to the LCG Middleware and, more in general to the LCG Services. Therefore major software changes and sites installations still to perform have to be performed in Q1 and Q2 2007. In order to monitor the progress of these final installations a set of Milestones for 2007 are being

LCG Quarterly Report – 2006 Q4 – Executive Summary

defined. In addition, being the LCG Service running permanently, it is now possible to measure and monitor its performance (in terms of reliability, jobs executions, etc) and also define the performance targets that must be achieved in order to fulfil the Experiments' requirements in 2007 and 2008.

In 2006Q4 the "definition of Targets and Milestones for 2007" has started. Targets and milestones are being set on activities that the experiments are planning for 2007. This work on planning requires gathering as many details as possible on the activities and plans of the Experiments for 2007. It is important that all targets are set and measured in conjunction with experiment real work with the services under realistic usage conditions. The current goal is to define the 2007 targets by February.

3 Service Challenge Progress

All participating sites have shown peak rates that satisfy the full nominal pp values and/or the full nominal Heavy Ion rates. Peak values of double or even triple those required have been seen, as well as the proven ability to recover from backlogs. The main service problems are still related to the storage systems at both ends, as well as to the additional complexity of coupling additional and needed services.

From SC4 to LCG Service – The service has continued to run uninterrupted, with the final phases of SC4 as well as continued ongoing productions from all the experiments. In particular the service continued over the Christmas holidays without significant interruption, and with continued use by the experiments. In particular ATLAS achieved very high rates of jobs. The individual data transfer targets for each of the Tier-1s was achieved but the target of long-term stable transfers at this rate to tape at the Tier-1 sites was not fully demonstrated.

The job reliability measurement system is under development and is being used within ATLAS and CMS as part of the experiment dashboards for some of their workloads, but there was no overall measure of job reliability during SC4.

Important additional components are being added to the LCG service, in order to be available to the experiments by July 2007. These services include production Distributed Database Services, SRM v2.2-based services, as well as support for VOMS groups and roles.

Network Transfers – Considerable progress was made in both data rate and stability, using experiment-driven transfers. Such transfers much more accurately reflect real-life use cases, as opposed to the pre-staged DTEAM driven transfers. CMS were able to achieve good rates to all of their sites, achieving more than the 150MB/s target set for SC4 (the full nominal rate for CMS is 270MB/s, without accelerator efficiency factored in). ALICE were also able to demonstrate peak rates above the Heavy Ion requirement, although service stability meant that these values were not achieved over longer periods of time. Combined transfers hinted at some coupling between the activities of the different experiments, and a programme to systematically debug such issues was launched, initially with IN2P3 and subsequently extended also to FZK and SARA.

The DTEAM resources at CERN were gradually reduced during the SC4 service phase and finally removed completely. DTEAM now uses the shared public CASTOR2 instance. This, coupled with the issues seen from multi-VO tests, meant that daily averages of full nominal rates were not achieved. However, it was understood as part of the LHCC Comprehensive Review that such targets were almost certainly overambitious and that more realistic values, taking the latest Megatable numbers together with some realistic accelerator efficiency factor for 2007/8, should be established. Progress on FTS monitoring helped to significantly improve the success rates of transfers. Nevertheless, it is understood that the service is likely to need significant 'manual' effort, at least in 2007 and 2008.

Multi-VO and Sites Tests - The multi-VO tests with IN2P3, coupled with an upgrade to dCache 1.7 plus fixes, have shown a marked improvement in service stability. Multi-VO tests will continue in Q1 2007 and be repeated in Q2 (using SRM v2.2-based services), so as to be ready for experiment 'dress rehearsals' and other such exercises by May/June 2007.

The active participation of Tier-1 sites at an early stage is required, as experience shows that debugging services is a lengthy and iterative process. Whilst the stability of the services involved needs to be further improved, it is important not to forget that these data transfer activities must take place in conjunction with

experiment production work - it is not sufficient to demonstrate that individual components work successfully, but rather the whole chain under realistic usage conditions. This is a priority for 2007.

4 Summary of Progress and Issues

4.1 Global performance and Reliability

Performance and Reliability Targets - The individual data transfer targets for each of the Tier-1s was achieved at some time during SC4. The aggregate target of 1.6 GB/s from CERN to the Tier-1s was met for a short period but the global performance targets set for SC4 have not been fully met.

Site reliability metrics were established in April and measured throughout the period of SC4 for CERN and 9 Tier-1s (BNL and NDGF did not participate). The reliability of the 8 best sites in each of the three months of SC4 was 83%, 86% and 82% respectively, compared with a target of 88%. However, the target of long-term stable transfers at this rate to tape at the Tier1s was not demonstrated. The target reliability was achieved, though only just, during November and December, the first two months of the production LCG service.

Monitoring the LCG Services - Three monitoring working groups have been proposed, agreed and started. These will try and address the various aspects of monitoring with the goals of improving the overall reliability of the LCG services, and providing more complete information to the site administrators, service managers, and users. They are not development projects but they focus in organizing how current tools and procedure should be defined.

Job Reliability - The job reliability measurement system is under development and is being used within ATLAS and CMS as part of the experiment dashboards for some of their workloads. There was no overall measure of job reliability during SC4.

The work on job log processing to identify the reasons for job failure has made good progress providing reports to ATLAS, CMS and to the daily grid operations service. The dashboard work is now being adapted also for ALICE and LHCb. To work efficiently it is important that there is a formal structure linking the experiment to the site; this requires someone within the experiment to work with a site to investigate the reliability graphs.

4.2 Procurement and Resources

Procurement - Following the reschedule of the LHC programme for 2007 and 2008, several sites have re-defined their procurement schedule for next few months. Some sites are reducing their disk expansion purchases, and wait for the availability of lower prices or for the introduction of new technologies. In most cases the requirements during 2007 and 2008 are lower than before, offering an opportunity to the funding agencies to fulfil these requirements with the funding that is available.

This is being monitored, and a process to collect and verify “site resources vs experiments requirements” has been put in place in order to make sure that the sites have sufficient resources for the experiments’ activities.

Resource Requirements - Experiments’ requirements and sites installations for network, disk and tape resources at the Tier-0, Tier-1 and Tier-2 sites are all collected in one single Resource Table, for convenience usually called the Megatable. This is the reference for all current procurement and preparation of resources for 2007.

The first version of the Megatable was made available at the end of 2006. This shows the implications of the computing models and the resources planned at sites on relationships between the sites, data transfer performance, and allocation of disk space to the different storage classes at each site. The table is currently being reviewed by the sites and by the network experts. At the same time the experiments continue to refine it as the understanding of their computing models develops.

Here is the Megatable, as on the 24 January 2007: <http://cern.ch/LCG/documents/Megatable240107.xls>

LCG Quarterly Report – 2006 Q4 – Executive Summary

Milestones for LCG Service - The quantitative measures in the Level-1 milestone for successful completion of SC4 cover only reliability and Tier-0/Tier-1 data distribution rates. Additional quantitative targets will now be set for the LCG Service, such as job submission rate, number of simultaneous jobs, catalogue access rates, etc. Milestones should also be defined to follow experiment preparations for the beginning of their data challenges and to measure their success in using the services in 2007.

4.3 LCG Middleware Releases and gLite 3.0

GLite Deployment - gLite-3.0 was deployed on schedule for the start of SC4. Subsequent releases (gLite-3.0.1, 3.0.2) have been produced and deployed to address many issues and bug fixes. The model of updates for gLite has now evolved to facilitate updates of individual components or services, and since gLite-3.0.2 there have been 11 updates (covering >100 patches). The updates are now very regular, and sites upgrade as needed, unless a patch is signalled as critical.

GLite Components - The new Resource Broker (RB) has been shown to be capable of handling about 25-30K job submissions per day from a single server. This is acceptable for CMS and probably for ATLAS, but needs to be further improved over the next year. However the code is still not stable enough from a service manager or users point of view to be deployable in full production. Work continues with the RB developers, ATLAS, CMS and the deployment team to debug the problems.

A first version of FTS v2 (compatible with SRM v2.2 and 1.1) is being tested; the common rfiio library is now used in DPM and is being implemented by Castor to remove the current problem with two incompatible versions of the libraries.

A version of the xrootd implementation in DPM has been made available to ALICE for testing and evaluation. Work on the xrootd interface for Castor, being developed by SLAC, has not progressed. In both of these cases it is too early to make any formal commitments for long term support.

The VOMS services are now in full production, the old LDAP-based services have been closed in October as scheduled. Security vulnerability in Torque which had been publicised by the Torque community was patched and deployed rapidly to a large number of LCG sites without unduly disrupting the service.

Progress on testing the gLite CE is halted waiting for an updated version of Condor (VDT) that avoids the bug limiting the number of jobs that can run on a site to 100.

The gLite WMS is still not stable enough from a service manager or users point of view to be deployable in full production and is being investigated by the developers, together with ATLAS, CMS and the deployment team.

GLite Releases - The release process is now much cleaner with the majority of changes being provided as updates rather than full releases. Since gLite-3.0.2 there have been 11 updates (covering >100 patches). The updates are now very regular, and sites upgrade as needed, unless a patch is reported as critical.

There have been continuing delays in the testing and certification of new releases of the EGEE products. This was studied in depth at an EGEE meeting in Bologna during the week of 15 January and a radical procedure has been put in place to try to resolve the problems. The intention is to make a thorough examination of the code base with a view to removing unnecessary dependencies and cleaning up sections of the code that causes build difficulties. A management group has been identified led by the EGEE Technical Director to provide day to day management of the process. While this attempt to solve a long-standing problem is certainly welcomed by LCG we will require sensitivity to the needs of the LHC experiments during this critical period.

SL4 Porting - Progress of the port of the gLite distribution to the SL(C)4 version of Linux has been delayed. This will have implications for the deployment of the experiments' application packages, in particular requiring the maintenance of both SL3 and SL4 versions during at least the first half of the year.

The WN port (components built under SL3 and packaged for SL4) is available and being tested on the PPS and will be available for deployment shortly. The full port including 64bit versions awaits updates from VDT and does not yet build fully. This issue has been raised with EGEE-JRA1 as a major problem.

4.4 SRM Interface and Implementations

The progress on SRMv2.2 is considerably slower than anticipated. Test end-points for Castor, DPM, and dCache are available, but are not yet neither stable nor providing the full set of functionality agreed in the LCG SRM 2.2 MoU. It is nevertheless hoped that we will be able to stabilize the situation and complete the basic test cycle within one to two months.

While the SRM 2.2 implementations will be certified each site has to understand how to map the storage classes of each experiment on to its Mass Storage Service. A working group was formed to help with this and provided some examples and guidelines. Now experiments and sites must collaborate closely on getting this going before summer 2007.

If the SRM 2.2 is not deployed in time for commissioning the LCG services, some sites will have to fall back to the more limited functionality of SRM 1.1.

4.5 Sites Reliability Monitoring

SAM and GridView – The two systems are in production and there are not major changes except that these two tools are now being integrated. The decision is to use the SAM system to collect the test data meanwhile GridView will visualize such data, providing a richer and customizable display via a web-based GUI (<http://gridview.cern.ch>).

Some of the SAM tests have shown problems and will need adaptation for certain situations, and the full set of service tests is not yet there. In particular experiments-specific tests are needed in order to provide a realistic verification of the services that are needed by each experiment.

The SAM system is now in full production and used since May 2006. Availability metrics have been published for CERN and nine Tier 1 sites since that time.

OSG decided not to use the SAM tests from CERN, but to develop equivalent tests using their own tools.

4.6 Accounting and Security

Accounting- During 2006Q4 a considerable progress has been made on User and Storage Accounting. Data is now collected by the APEL system and most sites are providing data or working to provide it (INFN, OSG, etc). The solution for the user privacy issues for user-level accounting has been proposed and implemented. But for the time being the accounting data gathered is not published because privacy and security policies must be officially approved.

OSG decided not to use the APEL tests from EGEE, but to develop equivalent accounting metrics using their own accounting tools and directly inserting their data in the APEL accounting repository.

Privacy and Security - Users privacy is a very important and delicate issue. Therefore all user-level accounting information is encrypted and all users identities masked to all viewers, except to the authorities allowed to view it (the user herself, site admins, VO admins, etc). The major Security and Privacy policies have been defined by the Joint Security and Policy Group (JSPG) and the final documents were circulated in 2006Q4. The goal is to approve them during the Grid Deployment Board meetings in 2007Q1, probably March.

One major security vulnerability in Torque, which had been publicized by the Torque community, was patched and deployed rapidly to a large number of LCG sites without disrupting the service.

4.7 Experiments Preparation

Tier-1 - Tier-2 Associations – The activity started in summer 2006 to integrate data from all four experiments on the expected relationships between Tier-1 and Tier-2 sites is now completed. It includes both data rates and storage requirements of the experiments on each Tier-1 and Tier-2 site.

This work resulted in an extensive and very important spreadsheet (called informally “the Megatable”) describing all such requirements. This information enables sites to configure and test the necessary FTS services, verify that there is adequate networking performance, and initiate operations and networking with their partner sites.

LCG Quarterly Report – 2006 Q4 – Executive Summary

As mentioned before the Megatable is also the reference of the targets for all transfers between Tier-0, Tier-1 and Tier-2 sites for the preparation of the 2007 Plans.

DAQ-T0-T1 Data Flow - Although this is technically outside the scope of the project it is clear that these connections may cause performance or other problems for the Tier-0 operation. There is still concern over the schedule for connecting and testing the DAQ systems to the Tier-0, but for 2007 all experiments have scheduled tests of their DAQ to T0 data flow. More in detail:

- ALICE: DAQ to T0 has been tested with the full data rate, the offline data flow was not included
- ATLAS: tested the chain emulated-DAQ to T0 to T1 with nominal speed and the full offline chain, including data base access T0 to T1 was only at the 50 % nominal speed level and no real DAQ access. The next tests including the missing features are planned for February and June 2007.
- CMS: first test of the T0toT1 part including the full T0 software setup at the 25-50% nominal speed level integration of DAQ and moving to 50-75% speed is planned , but no details yet
- LHCb: DAQtoT0 tests planned for March 2007 , offline integration not clear yet

24 X 7 Support - The 24x7 support plans are being implemented, with different levels of completion, at all Tier-0 and Tier-1 sites and the support scenarios will be tested in spring 2007. By summer 2007, when experiments will start commissioning their software, all sites must be able to provide permanent and reliable support and maintenance to the grid users.

VO Boxes Service Level - The VO Boxes needed by the LHC experiments are installed at all Tier-1 sites and are functioning adequately. No work was done between sites and experiments in order to agree on how VO Box configurations and data are stored and backed up (by the sites or by the VOs). Defining a clear “SLA on VO Boxes Support” is one of the milestones for 2007.

5 Specific Areas and Projects

5.1 Applications Area

New versions of the Applications Area packages, in particular ROOT, GEANT4, CORAL and COOL, are currently being integrated by the experiments and will basically be the versions, besides possible bug fixes, that are going to be used for the start of the LHC experiments. Production releases for Geant4 and ROOT have been made available at the end of the year.

GEANT4 Release - Version 8.2 of Geant4 includes improvements in the standard electromagnetic physics such as multiple scattering, which provide improved results for large angle scattering; better particle transport near geometry boundaries and less cut dependence for sampling calorimeters. In the hadronic physics, coherent elastic scattering processes have been reviewed and data tables are now computed on-flight. The physics lists are now built as part of the kernel libraries by default.

ROOT Release - ROOT version 5.14/00 includes many new functionalities and improvements in basically all areas. The release notes in <http://root.cern.ch/root/Version514.news.html> give all the details. The work of re-engineering CINT to use Reflex is progressing steadily. A new version of CINT, known as CINT 7.0, that uses Reflex to store the information regarding

CORAL Release - The Relation Abstract Layer (CORAL) have had the first release of authentication functionality based of LFC and the Python API. The Conditions Database (COOL) version 2.0.0 is almost ready for release and includes an API for the Record Specification and the port to AMD64 architecture. This new version is being currently tested for integration by ATLAS and LHCb since requires some changes in the DB schema and API.

SEAL Components Deprecated - The migration of to the SEAL component model has been completed in the POOL CVS repository for most packages. The same applies for the migration to the CORAL AttributeList. However, no corresponding releases have been made available yet and the finalization of the development has been cancelled as a consequence of the recent decision to move the SEAL functionality into CORAL and to deprecate the use of SEAL in POOL and CORAL.

MC Event Generators - The detailed program of work for the MC Generator services sub-project has been discussed and reviewed by all the stakeholders in a planning meeting end of October. Rapid progress

LCG Quarterly Report – 2006 Q4 – Executive Summary

is currently being made to release all the required generators in a new structure more suited for the needs of the experiments.

Nightly Build System - A nightly build system has been put in place following the recommendation of the AA internal review. The main goal for this system is to be able to tests/validate new versions of the AA software by the experiments before they are released, thus speeding up the overall development and release cycle. The introduction of the nightly build system has been synchronized with the migration from the SCRAM configuration and built system to the CMT one, which was also supported by the internal review. The system starts being functional and has been very useful for the experiments.

PROOF - Development in the last quarter was focused on providing the needed features for the ALICE analysis use cases. The developments included an extensive monitoring facility to track performance on the ALICE CAF cluster and to provide the infrastructure to successfully run the ALICE analysis. ALICE measured a very satisfying speedup of the analysis with high cluster usage efficiency. At the end of 2006 first interactions with CMS were started and we expect a lot to happen in the coming quarters. New plans are in the area of a PROOF-lite, which is aimed at using PROOF on single multi-core machines (simplification of daemon architecture) and in helping CMS getting their environment PROOF enabled.

5.2 Distributed Database Deployment – 3D

All phase 1 Tier-1 sites are now in operation and use by the experiments for the conditions databases is beginning. The system is also being used by LHCb for replication of their LFC catalogue.

All phase 2 sites (NDGF, NIKHEF/SARA, PIC and TRIUMF) are all actively working on the commissioning of the requested database setups, but not all are available yet for use by the experiments. At TRIUMF and SARA database clusters have been setup and are currently validated by the experiments. At NDGF and PIC the cluster setup is still being performed and both sites risk to not be ready in time for their availability milestone end of March. All other sites have been tested by ATLAS and LHCb and did reach the estimated replication rate for conditions data.

The 3D streams monitoring setup has been put into production since November last year and does alert experiment and site responsables about operational problems. It also provides graphical displays of the over all status of the database installations at all sites, which is now being integrated into the experiment dashboards.

The deployment of the initial Frontier/Squid system and the associated monitoring for CMS is complete and has been successfully tested during the CSA06 activity with some 28 Tier 1 and Tier 2 sites.

Updates on the experiment resource requirements and agreement on the database backup and security policies which have been prepared by Tier 0 and discussed with experiments and sites are expected for the upcoming database workshop (26 Jan).

5.3 ARDA

Ganga - The status of the tools (now in version 4.2) is satisfactory and includes support for experiment-specific software (ATHENA and GAUDI) plus splitting and merging facilities (transform a task on a data set in a set of jobs and merge the results).

In ATLAS the interest and the visibility of the tool has improved (independent users are using the tool especially in the UK; in some case new components are contributed). Feedback is sizeable and satisfactory. The ATLAS system PANDA is the standard tool for job handling in the OSG and we might use Ganga as a front-end if needed. The main risk factor is the actual demonstration that enough data can be made available via the ATLAS DDM (Distributed Data Management).

In LHCb the status of the tool is well-established (main mode of operation: Ganga as a front-end of DIRAC).

Dashboard - All "views" and functionality we foresaw are available for ATLAS and CMS. In this quarter LHCb showed interest and immediately after also ALICE did the same. The work plan has been modified in the following way: we used LHCb has an opportunity to rationalise some part of the code (preparing for the departure of one developer) and to improve the installation. This system is now in production and the

LCG Quarterly Report – 2006 Q4 – Executive Summary

ALICE one should follow soon due to the code/installation improvements. We plan to provide regular monthly report for each VO and each Tier1 starting early next year.

ALICE - Support of the production is continuing. The monitor effort (MonAlisa related) is an important ingredient of their system benefiting also the dashboard and support activities. The system is composed by a large number of VO boxes (one per Tier1 and Tier2 site) and could be run successfully, but some necessary automation must be added.

ATLAS - Support of their production and analysis activities has continued on several fronts. The LEXOR-based production has considerably improved (because of better operational procedures and improvements in the software): this experience will be part of the WMS evaluation (see CMS section, next paragraph). A very important new activity is starting to further support ATLAS on the data management area. By refocusing existing effort, a team is being set up to support ATLAS in the deployment of the DDM system: the programme of work is being prepared and an ATLAS DDM review took place in December.

CMS - Support of CSA06 was a central part of the activity. The feedback from CMS was very positive. We are planning (also together with ATLAS) to repeat/improve the validation performed in the last quarter as preparation of CSA06. The expectation is to move to the 3.1.x WMS and improve on especially on reliability in close contacts with the developers and the operation teams.

LHCb - Support of their production continued to support the successful DC06 activity. The LHCb dashboard opens new possibilities of collaboration on monitor activities (so far done in the framework of DIRAC outside our team).

6 Experiments

6.1 ALICE

AliRoot Simulation: The standard transport package remains GEANT3 because there are some difficulties to put Fluka, which has much better physics, in production. Work has started to implement the day 0 geometry.

AliRoot Reconstruction: The parameters are accessed from the central Offline Condition Data Base. The detector algorithms are being refined. Progress in reducing the memory consumption and the processing time has been achieved.

AliRoot Analysis and Visualization: The new analysis framework has been released and is under test by users. A prototype for the AOD has been released. The visualization framework is fully implemented and is under evaluation by the users.

Production on the Grid: Job submission is working smoothly and has been exercised continuously during PDC06, including during the Christmas holidays. The Grid operation has been automated and now it is possible the entire ALICE production even with a small team.

Grid SE: The interface xrootd-dCache is in early testing phase, and the interface xrootd-CASTOR2, DPM is still under development. Operational problems encountered with FTS during the exercise of pushing data from CERN have been identified and discussed with the FTS development and operation teams. In general, the Tier-0 to Tier-1 transfer exercise was extremely helpful in tuning our high level services to work with FTS. The nominal target rate has been reached and sustained for a short period of time. This exercise will restart and run continuously during 2007.

Grid Middleware: The evolution of the ALICE services and the use of production LCG services are steady. ALICE tested the gLite resource broker at CERN and found it to be satisfactory. The interface to ARC is progressing very slowly and the development of the interface to OSG is stalled.

Development and tests of PROOF on the CAF are going on. The implementation of quotas, load balancing and disk management has started.

Grid resources: only about 50% of the resources pledged by the sites were usable or available for PDC06. On the long term, whereas the anticipated situation with the resources seems adequate in 2008 in later years the missing resources are a danger to the execution of the ALICE physics program.

6.2 ATLAS

The work on optimisation of data access led the collaboration to choose a new Event Data Model for reconstruction output, which allows direct ROOT access to the AOD data. Its implementation is in progress and should be completed for release 13 in March.

The ATLAS Software Release 13 will also include:

- Time-dependent calibration infrastructure and support for Calibration Data Challenge
- Parameterized shower simulation in production
- Full Trigger chain
- Optimized reconstruction algorithms
- Schema evolution support in place for all data
- Port to SLC4 and 64-bit architectures

Simulation production operation continued with Release 12 at a rate of 2M events/week in November-December (with a peak of 4M events the last week of December). CSC Tier-0 tests showed that the system can sustain the nominal number of jobs and data transfer rates. Tests in Q1-2007 will include more realistic calibration and reconstruction code and exercising error conditions.

Data distribution tests operated concurrently, over Tier-0, Tier-1 and Tier-2 sites, with other activities and experiments in autumn 2006 showed that only a fraction of the nominal data transfer rate can be achieved currently. Focussed tests to optimise the site and transfer channel configurations are in progress.

Open issues are mainly due to instabilities of the sites, mainly Storage Element problems, that affect productions and also, in different ways, data transfer tests. Also the user access to the data on the Grid using DDM tools will soon need attention.

6.3 CMS

During the quarter CMS activities concentrated on the Computing, Software and Analysis (CSA06) challenge. All defined goals were met; some technical metrics were exceeded by large factors. More details are available here: <https://twiki.cern.ch/twiki/bin/view/CMS/CSA06>

The major achievements of CSA06 are here summarized:

- T0 prompt reconstruction of RECO, AOD, and AICaReco with Frontier access performed at 100% efficiency for 207M events.
- Data were exported to T1 at rates of 150 MB/s and higher.
- Offline database access at all participating remote sites via Frontier/Squid successfully tested.
- Skim production at T1s was demonstrated, results were transferred to T2s.
- Re-reconstruction was performed at 6 T1 centres.
- Alignment/Calibration/Physics analyses were successfully demonstrated.
- Job load reached 55K jobs/day.
- Data transfers performed between T1 centres and non-associated T1-T2 transfers executed.

A comprehensive report on CSA06 including lessons-learned is in preparation and will be submitted to the LHCC in Feb. 2007.

6.4 LHCb

The experiment did not provide a large amount of information in the Quarterly Report (details at pg.65 of the “Quarterly Status and Progress Report - 2006 Q4” document).

Four sites (RAL, CERN, PIC and IN2p3) are generally behaving smoothly. The outstanding problems with dCache between NIKHEF and SARA seem resolved and tests are ongoing. GridKa tests are ongoing after the recent upgrade of dCache. A temporary fix has been implemented at CNAF whilst awaiting the deployment of a fixed CASTOR. Accessing data from the MSS systems continues to be problem at several sites. The stripping phase will commence once smooth operations of the reconstruction is achieved.