



LHC Computing Grid Project
Quarterly Status and Progress Report
2007 Q1
Executive Summary

Alberto Aimar
17 May 2007

1 Introduction

This document highlights the 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.Hernandez, M.Lamanna, L.Robertson and O.Smirnova). The reviewers then asked for additional information and, in some cases, the reports were modified and re-submitted by the original authors.

For details about the topics summarized in this document please refer to the complete document: Quarterly Status and Progress Report - 2007 Q1” where all individual reports are collected. In addition, all past Quarterly Reports, since 2005Q4, are always available on the [LCG Quarterly Reports](#) web page.

2 Level-1 Milestones

2.1 Completion of Milestones for 2006

DBS-1: Full LCG database service in place

Done – This milestone was still not fully completed at the end of 2006. But since then considerable progress has been observed at all sites that were not ready. All of the Phase 1 sites - ASGC, BNL, CNAF, GridKa, IN2P3 and RAL - were ready in 2006 and now, with TRIUMF and SARA, are taking part in the ATLAS and LHCb conditions database tests. PIC and NDGF that were late in the previous quarter have both progressed: PIC have hired an Oracle DB administrator and setup one two-node cluster that will be their initial production server and NDGF have setup a single-node database server as a temporary solution while acquiring the hardware equipment for the final 3D database cluster.

IS-1: Initial LCG Service in operation

Done – The LCG Service, as result of the work on the SC4 Service Challenge in 2006, has been shown to deliver usable, and used, services throughout 2007Q1. 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 development. In 2007 new milestones and targets will verify that the Experiments use cases are all served adequately and the Experiments have planned thorough testing in preparation for their Full Dress Rehearsals (FDRs) in summer 2007.

2.2 Milestones for 2007

A revision of the high level planning and reporting process has been completed, adapted to the operational phase of the project where all sites must provide a certain number of services and achieve a required level of support. High level milestones are now defined as targets to be met by each of the major sites (CERN, Tier-1 sites). Every site reports on the High Level Milestones in its Quarterly Report. The initial set of High Level Milestones, agreed by the Management Board by February 2007, is listed below:

24x7 Support

- **Feb 2007: 24x7 Support Definition** - Definition of the levels of support and rules to follow, depending on the issue/alarm

LCG Quarterly Report – 2007 Q1 – Executive Summary

- **Apr 2007: 24x7 Support Tested** - Support and operation scenarios tested via realistic alarms and situations
- **Jun 2007: 24x7 Support in Operations** - The sites provide 24x7 support to users as standard operations

VOBoxes Support

- **Apr 2007: VOBoxes SLA Defined** - Sites propose and agree with the VO the level of support (upgrade, backup, restore, etc) of VOBoxes
- **May 2007: VOBoxes SLA Implemented** - VOBoxes service implemented at the site, and tested by the Experiments

Job Priorities

- **Apr 2007: Job Priorities Available at Site** - Mapping of the Job priorities on the batch software of the site completed and information published
- **Jun 2007: Job Priorities of the VOs Implemented at Site** - Configuration and maintenance of the job priorities as defined by the VOs. Job Priorities in use by the VOs.

Accounting Data

- **Mar 2007: Accounting Data published in the APEL Repository** - The site is publishing the accounting data in APEL. Monthly reports extracted from the APEL Repository.

3D Database Service

- **Mar 2007: 3D Oracle Service in Production** - Oracle Service in production, and certified by the Experiments
- **May 2007: 3D Conditions DB in Production** - Conditions DB in operations for ATLAS, CMS, and LHCb. Tested by the Experiments.

SL4 Migration

- **Deployment Date + 30 days: SL4 Operational at Site (for WN and UI nodes)** - This has to happen within 30 days after the release from GD.

Site Reliability Targets

- **Jun 2007: Site Reliability above 91%** - Considering all 11 Tier-1 sites
- **Jun 2007: Best 8 Sites above 93%** - Eight sites should reach a reliability above 93%

- **Dec 2007: Site Reliability above 93%** - Considering all 11 Tier-1 sites
- **Dec 2007: Best 8 Sites above 95%** - Eight sites should reach a reliability above 95%

The LHCC referees have asked for a dashboard style of presentation, showing graphically the progress across the different sites. For the High Level Milestones a display will be maintained showing the progress on each milestone at each site – see Appendix 1 for the dashboard at the end of 2007Q1 and a more recent one at end April 2007. The dashboard is updated every month and is available on the LCG Planning Web.

The cells of the matrix are coloured to indicate the state of progress at each site on each of the milestones: green = achieved; orange = overdue by less than one month; red = overdue by more than one month; white = not yet due.

2.3 Targets for 2007

Another set of targets has been defined for data transfer between sites. These specify for each Tier-1, for each LHC experiment, the data rates to be sustained for a certain minimum period between:

- Tier-0 and the Tier-1 site;
- Tier-1 to/from the other Tier-1 sites;

LCG Quarterly Report – 2007 Q1 – Executive Summary

- Tier-1 to each of the associated Tier-2 sites.

Targets are also defined for simultaneous operation of Tier-0/Tier-1 and Tier-1/Tier-2 traffic. The rates are based on the data rates for 2008 in the Megatable.

All Experiments have expressed their targets for 2007. Those targets will be reviewed in 2007Q2, while the transfer rates will be measured.

3 LCG Service Overview

Summary - The LCG Service is taking shape successfully. Progress continues to be made on better scheduling, and announcing of service interventions, as well as attendance at and effectiveness of the weekly joint Operations meeting. The services that will be in place for the preparations for the Full Dress Rehearsals have been clarified. Among others, these include LFC 1.6.3 which brings support for bulk methods, FTS 2.0, partial support for VOMS roles in job scheduling and 3D Services for conditions data.

File Transfer (FTS) - FTS version 2.0, with improved monitoring and administration tools and support for SRM 2.2 installed for testing at CERN since December. It is expected to be in production at CERN in April, then progressively deployed to the Tier-1 sites during May. At CERN the Tier-0 service has been separated from the general purpose CERN-Tier-2 service, to protect the main Tier-0 to Tier-1s data distribution service.

SRM 2.2 - The situation has improved considerably during this quarter. At the end of March 2007 the basic tests and essentially all relevant Use Case tests can be passed by the three implementations (dCache, CASTOR and DPM). However, there are still stability and operational issues with dCache and CASTOR, and the stress tests only started at the beginning of April. Additional endpoints are scheduled to be set up during April therefore testing by the Experiments should start by the end of May – about two months later than the original schedule.

File Catalogues (LFC) - With respect to the services required for the Experiments' 2007 Dress Rehearsals, LFC 1.6.3 has been released and deployed at CERN and some of the Tier-1s. This version supports the bulk methods required by ATLAS and its production deployment at ATLAS Tier-1 sites is being actively followed up both through the weekly joint operations meeting as well as ATLAS channels.

3D Database Services - As mentioned in Section 2.1, good progress has been made in deploying the hardware required for production distributed database services. All of the Tier-1 sites have now a 3D database service. PIC and NDGF, that were late in the previous quarter, have considerably progressed.

24x7 Support - A proposal for 24x7 on-call service for the main physics services at CERN has been prepared and submitted to the CERN management. The goal is to implement such an on-call service in advance of the 2007 engineering run. Several sites are late in defining and testing their 24x7 model. High level milestones (WLCG-01, -02, -03) have been defined to monitoring the progress at every Tier-1 site. The status of these and other high level milestones is in Appendix 1.

VOMS-based scheduling priorities - Base software components are being delivered by EGEE and have been deployed at most of the Tier-1s for beta testing prior to being included in a gLite release. However, deployment requires local configuration and adaptation to the local batch system, and gLite support is only available for the Torque system in the initial release. Support for other batch systems (LSF, Condor, BQS) has been developed by sites using these batch schedulers. It is expected that the service will be available at EGEE Tier-1 sites by the end of April. LCG is not monitoring a more widespread deployment to Tier-2 sites, OSG, NDGF, etc).

Multi-VO Tests - Building on the experience from the Service Challenges, a programme has been proposed to repeat the multi-VO Tier0-Tier-1 transfer tests of last November, in which concurrent transfers from at least ATLAS and CMS are performed to common sites (possibly also including BNL & FNAL) under realistic conditions and at an agreed fraction of clearly stated nominal data rates (e.g. 50%) to be consistent with CMS' CSA07 goals. 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.

Service Organization - The status of the delivery of the residual services is being closely monitored at the GDB / MB level. Whilst the full roll-out of all of these services is not expected to be completed even

LCG Quarterly Report – 2007 Q1 – Executive Summary

this year, good and steady progress is being made, both in the incremental delivery of new features, as is the case of the DPM/LFC, and in experiment testing / validation, e.g. for the FTS 2.0 service.

SCOD - To address the service issues, it is being discussed that the “Service Coordinator on Duty” be established soon, on a trial basis, before formalising the role and its responsibilities. In addition, a regular high-level “Service Report” should be prepared, as well as attendance at daily and weekly operations meetings. It is considered extremely critical that the current issues in service operation and delivery (e.g. lack of clear and up-to-date information) are corrected rapidly.

Issues During the Quarter - The primary issues of this quarter for the LCG Services were the failure to achieve concurrent multi-VO Tier0-Tier-1 transfers at the agreed percentage of the nominal rates and the delays in the delivery of the residual services (target was the end of March) to be ready for experiment testing and preparations for the summer's Full Dress Rehearsals.

3.1 Global Performance and Reliability

Site Reliability - The site reliability summary for October 2006 through March 2007 is always available on the LCG Web pages. ([Availability and Reliability Data](#)).

Site reliability metrics were established in 2006 and reported monthly for CERN and 10 Tier-1 sites (NDGF still does not participate). The target reliability until June 2007 is 88%, to be achieved by all sites and will be 91% from June 2007 (with 93% for the 8 best sites).

Site	Jan 07	Feb 07	Mar 07	Apr 07
CERN	99	91	97	96
GridKa/FZK	85	90	75	79
IN2P3	96	74	58	95
INFN/CNAF	75	93	76	93
RAL	80	82	80	87
SARA-NIKHEF	93	83	47	92
TRIUMF	79	88	70	73
ASGC	96	97	95	92
FNAL	84	67	90	85
PIC	86	86	96	95
BNL	90	57	6	89
NDGF	n/a	n/a	n/a	n/a

Note: For consistency the (dis)order of the list of sites above is the same as the order used in the [Availability and Reliability Data](#) summary.

The reliability of the 8 best sites in each of the three months of SC4 was 91%, 89% and 85% respectively compared with a target of 90% (and in April was up to 92%).

The Management Board has been focusing its attention on site reliability regularly during the past few months, with each site providing a summary of the problems that have led to each interruption and the solutions applied. The situation is slowly improving and this is an area of some concern in view of the major upgrades and installation that will be performed in 2007 and 2008.

Note: Some March values are particularly low because of some problems with the network and with the test system itself. In some cases the failure of the tests is not depending on the site but on the general infrastructure (network, etc) or in other situations a test fails at a particular site but the site is still able to

LCG Quarterly Report – 2007 Q1 – Executive Summary

run jobs for one or more experiment (e.g. BNL performed well for ATLAS even if did not pass the SAM tests for most of March).

VO/Sites Specific Reliability - The strategy that is being developed in the next quarters it to define tests more relevant to the VO trying to reproduce their operational environment. In addition the Experiments now specify their own tests and integrate them into SAM. This will allow the reliability of the site to be calculated in the Experiment's environment and will provide a more realistic value of the reliability they obtain from their grid applications on every LCG site.

Job Reliability Data - At the end of the 2007Q1 quarter a first set of job reliability metrics has been prepared, using specific job sets: from the experiments: (1) CMS jobs submitted via CRAB to EGEE Resource Brokers; (2) ALICE job agents and (3) LHCb pilot jobs. Work is progressing in April and May to provide job reliability data for ATLAS and job submitted via Ganga.

At present only jobs submitted through an EGEE Resource Broker are included, as insufficient logging information is available for jobs submitted by other methods. The first set of results for February and March are still being discussed by the Management Board that will also monitor the job reliability reports now produced monthly.

The work developed in the previous quarters 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. This analysis of job logs is now used for debugging and understanding why jobs fail, or need several attempts to succeed, on a given site. This kind of investigation requires someone within the experiment to work with a site to investigate the job reliability graphs and why in some cases a job succeeds only after several attempts.

Monitoring the LCG Services - Three monitoring working groups have been proposed, agreed and started. They will 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. More information is available on this web page: [WLCG Monitoring Working Groups](#).

3.2 Sites Procurement and Resources

Procurement and Mid-Term Requirements - 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 and 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 a viable but risky approach therefore is being monitored and issues reported to the Management Board. A process to collect and verify the Mid-Term Planning Table (Appendix 2) that tracks on a quarterly basis the sites - sites resources vs. experiments requirements - in order to make sure that the sites have sufficient resources for the Experiments' activities.

Appendix 2 shows the values pledged and available at each site compared with the requirements from each experiment (for CPU resources, Disk and Tape capacities).

Resource Requirements - Longer-term Experiments' requirements and site installations for network, disk and tape resources at the Tier-0, Tier-1 and Tier-2 sites continue to be collected in one single Resource Table, for convenience now 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. In parallel the Experiments continue to refine it as the understanding of their computing models develops.

4 Summary of Progress and Issues

4.1 Scientific Linux version 4 (SL4) Migration

There have been some technical difficulties in preparing the full SL4 ports, but now all gLite components are built using the ETICS build system and on SL4 in 32 and 64 bits mode. Currently several components are being tested and thoroughly re-certified, also as often they are newer versions of the same components and not identical to those currently running on SL3.

gLite Components on SL4 - The Worker Node (WN) package under SL4 is available on the Pre-Production Service (PPS) at CERN, and will be ready for production deployment in mid April. The User Interface (UI) builds but some packaging problems need to be resolved. The Workload Management System (WMS), CE and the Data Management components have been resolved to get complete builds. The WMS testing effort moved to INFN to allow CERN to focus on the CE testing. The port of the remaining components is progressing, but release dates cannot be given at present.

In order to judge the quality of the gLite WMS and CE, detailed requirements documents were written based on the performance of the current production components and taking into account the scaling requirements of the Experiments for 2007. The WMS has now reached a level of performance and reliability that is acceptable and will be put back into the usual certification chain. It is anticipated that the gLite WMS will likely be in the production release in early June 2007 and can soon replace the current Resource Broker (lcg-RB). The new CE still has efficiency problems - failing some 9% of jobs. Certification and release dates are not yet scheduled.

SL3 Future Support - The deployment of an SL4 version of the gLite distribution is urgent. Most of the computer systems being installed by regional centres since the last quarter of 2006 are not certified for older versions of Linux. Support for the version of Red Hat Linux on which SL3 is based stops in October 2007 (even if Red Hat will provide security patches until 2010). A backup solution has been available since the end of the year for Worker Nodes, the User Interface and the LCG Compute Element, using code compiled under SL3 executing on SL4. There are some potential risks with this solution, although no problems have been found in practice, and it does not fulfill the Experiments' requirement to have a native SL4 environment.

4.2 FTS Deployment and Services

FTS Deployment - FTS 2.0, which implements the improvements agreed at last October's workshop at SARA, has been deployed on a pilot service at CERN and will also be deployed in pre-production at a number of Tier-1 sites. At least one minor release prior to the start of the Dress Rehearsals in July is to be foreseen.

FTS Services - In addition to the FTS 2.0 release described above, the FTS service at CERN has been split into one serving the WLCG Tier-1 sites and another serving the Tier-2 sites. This allows for better separation of the different tasks, as well as improved tuning of service parameters according to the specific requirements. The old Tier-2 channels will be switched off by the end of April 2007.

Note that the service for Tier-0 to Tier-1 traffic is operated by CERN, but each Tier-1 has to operate the service for transfers between itself and the Tier-2 sites, and for data received from another Tier-1. This is a major responsibility for the Tier-1s that is being tested by each of the Experiments during the first half of 2007.

4.3 SRM Interface and Implementations

During 2007Q1 the SRM v2.2 project has made very significant progress. The last ambiguities and inconsistencies in the specification were essentially resolved during the WLCG workshop in January 2007, after lots of discussion through e-mail and phone conferences. At the same time the proposed testing plan was agreed by the developers. However, already for the basic tests dCache and CASTOR missed the milestone because of other major commitments of the developers.

SRM Implementations - Each of the implementations has come to comply with all the important use cases. In that respect the v2.2 specification and the S2 test suite have had to be adjusted a few times to

LCG Quarterly Report – 2007 Q1 – Executive Summary

allow for more diversity in the behaviour of the back-ends while still maintaining a sufficient amount of common functionality available to clients. DPM has production releases available. dCache 1.8 beta is already deployed at five sites. CASTOR SRM v2.2 is also deployed at CNAF. A new release of the CASTOR core is awaited for the crucial fast srmPutDone implementation. Some of the sites in INFN will use StoRM and some of the US sites will use BeStMan

SRM Testing - A lot of work has been done on the S2 test suite: bug fixes, implementation of timeouts, addition of a large number of use case tests, support for GFAL and lcg-util tests. The test suite has been adapted for stress tests to run on multiple machines in parallel. The documentation has been improved, as well as the web pages showing the results, given the ever increasing number of test endpoints. A lot of time has been spent on chasing the causes of test failures. The dCache developers at FNAL lost a week debugging what turned out to be a problem with ESnet packet losses. The LBNL test framework has been expanded and automated. Also the FTS developers have created a test framework with results published on web pages.

GSSD Working Group - A working group, the Grid Storage System Deployment Working Group ([GSSD](#)), with members from Experiments, sites and mass storage systems development has been established to understand and guide the deployment of the storage classes using SRM v2.2. Three storage classes have been agreed: tape1disk0 - the classic HEP storage class of tape with a front-end disk cache; tape0disk1 - disk only storage with the space managed by the experiment; tape1disk1 - backed-up disk also managed by the experiment. It is not yet fully understood how these can best be implemented by the sites and how they will be used and managed by the Experiments.

SRM V1.1 to V2.2 Transition - The GSSD SRM v1-v2 transition study group completed their report: after generic certification of v2.2 implementations Experiments agreed to make v2.2 work for them and stick to it, rather than switching back to v1.1 whenever problems are encountered. LHCb and CMS provided input on the storage class instances they need, data rates, access patterns etc. ATLAS has started that effort, ALICE not yet. The monitoring study group made an inventory of what dCache and DPM can provide through native tools. A common subset may form the basis of an implementation-independent monitoring system.

Issues to Follow - Problems seen with the main storage systems, including file corruption under certain conditions with dCache, as well as stability, functionality and performance problems with CASTOR, suggest that priority should be given to these critical issues, as opposed to new functionality - which can be expected to introduce a further period of instability - again addressing the immediate needs of the 2007 engineering run. It is too soon to say that the original target of production operation of dCache and CASTOR in July at some of the major sites will not be met, but it is clear that the Experiments should plan to depend on all the SRM v2.2 features during 2007.

4.4 Sites Reliability Monitoring Tools

SAM and GridView – The two systems are in production with no major changes apart from the integration of these two tools. 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 in use since May 2006. Availability metrics have been published for CERN and nine Tier 1 sites since then.

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

4.5 Experiments Preparation

Tier-1 - Tier-2 Associations – The activity which 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.

LCG Quarterly Report – 2007 Q1 – Executive Summary

This work resulted in an extensive and very important spreadsheet (called informally the Megatable) describing all such requirements and data transfers. 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.

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.

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 among the high level milestones for 2007.

DAQ-Tier-0-Tier-1 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 Tier-0 data flow.

5 Current Concerns

CASTOR - During last year all of the CASTOR milestones at CERN associated with experiment tests of the Tier-0 and data distribution were met, although it was recognised that there were shortcomings in the ability of CASTOR to handle the workloads that could be expected from simultaneous Tier-0 and general analysis activity. The start of the ATLAS tests in January showed up the current deficiencies and the service has not been able to fulfil the ATLAS requirements. On the other hand the performance requirements for the CMS and ALICE tests during the first quarter have been met.

A task force has been initiated to work on these problems in the context of the ATLAS Tier0 and data export tests, bringing in additional experienced people to help. Until the current development plan has been completed and put into production we cannot be sure whether this will fully resolve the performance problems, therefore the task force will also look further ahead. Until these problems are resolved this will remain the highest priority development activity for CASTOR.

Mass Storage Services - During last year there has been important progress in exercising the basic set of services for distributing data from CERN to the Tier-1s and, although to a lesser extent, from Tier-1s to Tier-2 sites. The MC productions have also provided experience with the reverse data flow into the Tier-1s. By the end of the year the local mass storage services to be used for LHC start-up were established at most Tier-1s and had been used during the period of experiment testing in the second half of the year. The exceptions were NDGF, which is now in the process of deploying a distributed implementation of dCache, and PIC, which still had to make a final decision on which mass storage system to use. Many Tier-2 sites have also established their mass storage services using either dCache or, where a tape layer is not required, DPM.

This year the Experiments' full dress rehearsals will also test the ability of the mass storage services to handle simultaneously the full set of use cases at the Tier-1s. It is probable that this will uncover additional performance and reliability issues, as well as stressing the storage operations and management processes at the sites and within the Experiments. The situation is complicated as in many cases the mass storage system is only one element in a long chain of components: the experiment data management systems, experiment data operations, the file transfer service, the wide area network, site operations, and the mass storage service at the remote site.

Debugging reliability and performance problems is very difficult in this complex distributed environment and the operations costs to both sites and Experiments must not be underestimated. It is important that we keep things as simple and stable as possible, avoiding trying to inject new functionality before we have integrated the full set of basic components.

Sites Reliability - Site reliability has been measured using a straight-forward set of tests since May 2006, reported to each meeting of the Overview Board. As can be seen from the report for the present meeting ([Section 9](#)) it is proving very hard to get individual site reliability as seen from the grid up to the current target of 88%, with the average still around 80%. The end of year target of 95% does not look to be

LCG Quarterly Report – 2007 Q1 – Executive Summary

achievable for all sites. With only 90% availability as seen from the grid, Experiments will have to live with an environment in which one or more Tier-1 centres will usually appear to be unavailable. While the reliability as observed within a site is certainly higher than that viewed from the grid, the difficulty in obtaining higher grid reliability figures with the current limited set of tests suggests that there are operational and monitoring difficulties at sites.

Multi-VOs Operations - Building on the experience from the Service Challenges, a programme has been proposed to repeat the multi-VO Tier0-Tier-1 transfer tests of last November, in which concurrent transfers from at least ATLAS and CMS are performed to common sites (possibly also including BNL & FNAL) under realistic conditions and at an agreed fraction of clearly stated nominal data rates (e.g. 50% to be consistent with CMS' CSA07 goals). In this respect, it will be necessary to agree on the sizes of the ATLAS RAW/ESD/AOD, as well as the number of copies of each that are exported and the fraction that is delivered to each ATLAS Tier-1 site.

The schedule of these tests will be coordinated through the LCG Experiment Coordination Meeting (LCG ECM) and operational issues followed up on via the weekly joint operations meeting. 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.

DAQ-Tier-0 Integration - The integration of the ALICE DAQ system with the Tier-0 has been tested with a throughput of 1 GByte/sec. It is planned to begin integration tests with LHCb in the coming weeks. Testing with ATLAS and CMS will not take place before September. Although the internal operation of the Tier-0 has been exercised for both of these Experiments this leaves little time for problems arising from unexpected usage patterns to be resolved.

End-user Analysis - There is an assumption by most sites that end-user analysis will largely be through batch job submission, where there is some experience in managing the workload. An evolution towards intensive interactive or notebook based analysis will be more difficult to manage. We still do not have a good characterization of analysis use of the grid or of individual computing centres. This is particularly problematic for CERN and the Tier-1s, where the analysis workload could interact strongly with the more production oriented activities through end-user initiated access to storage systems and network resources.

6 Specific Areas and Projects

6.1 Applications Area

During the first quarter of 2007 the Application Area has been consolidating the functionality released at the end of the year. 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 has been put into production. It allows the build of all LCG Application Area projects on most of the LCG platforms. Windows is about to be finished. Apart from the LCG projects also the Gaudi software framework used by ATLAS and LHCb has been added to the builds. The nightly builds are used by the Experiments for validation of changes and new features.

The 7th ROOT workshop took place from 26-28 March at CERN with a participation of about 80 people. The complete status and representative use cases of the ROOT software were presented in 38 talks and 7 posters.

All persistency framework builds are now built using CMT and regularly tested via the nightly build and test system. CORAL improvements, such as thread safety and access to stored procedures, have been delivered. The recent COOL releases include a new API for user payload specification, a port to the AMD64 architecture and new 'locking' and 'dynamic replication' functionalities as well as examples on how to use the CORAL LFC Replica Service.

Important deliverables has been achieved 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

LCG Quarterly Report – 2007 Q1 – Executive Summary

of hadronic shower shapes) have been published. In particular, the joint investigations performed by the Physics Validation and Geant4 teams for understanding the key components responsible for the development of the hadronic shower shapes, has already produced good positive results in the preliminary developments made available in the last Geant4 development releases.

The Generator Services subproject after the reorganisation initiated few months ago is now stable and progressing well, the new GENSER release (2.0.0) which follows the new structure and configuration policy defined in December 2006, is now available and used by the Experiments.

6.2 Distributed Database Deployment – 3D

During the last quarter the 3D project has seen significant progress from both the Tier-1 sites and the experiment tests. As of today all 10 Tier-1 sites have set-up a database service, which is included in the experiment tests. Not all sites have so far fully satisfied the experiment request in terms of database server nodes, but a complete set of database replicas will now move into production mode with the ATLAS and LHCb conditions challenges in April. In particular PIC, NDGF, NIKHEF and TRIUMF made a significant effort to catch up with the other Tier-1 sites, which is acknowledged by the experiment teams.

On the experiment side the work has continued to validate the access to the database resources at Tier-1 sites including some of the more recent software components in CORAL/COOL to select a one of the replicas. The work is now moving focus from optimization of the pure Tier-0 to Tier-1 replication to an evaluation of the Tier-1 database performance for grid clients. Scalability tests with the experiment software and workload are now starting, which will allow confirming (or if necessary modifying) the database resource request. These tests are expected to produce the main input to a database resource review during the summer to define the Tier-1 setup for the LHC start-up in autumn.

During the last 3D workshops in January and March the current experiment requests have been confirmed until summer and policy documents for the database security and backup/recovery have been agreed. 3D plans to perform a joint Tier-1 database recovery test on the timescale of mid-May to make sure that the backup procedures and the streams synchronization after a potential problem are valid and can be executed at all sites.

The dates by which the main services - initially LFC and conditions - continue to slip with a new target for the ATLAS conditions of April 15th. This is due to a new release date for the corresponding ATLAS software. Some issues with Oracle streams behaviour in the case of site outages have been seen and have been escalated to Oracle. For the conditions application in particular, a rapid resolution of these problems is considered essential by ATLAS. A number of stop-gap solutions and work-arounds are being discussed, but none appear to satisfy the basic requirements, hence the urgency for a solution from Oracle to these problems. The date by which the ATLAS TAG application requires a production service needs to be clarified. In particular, should this application form part of the 2007 Dress Rehearsals, this must be established without delay. The overall resource requirements will in future form part of the overall resource requirement tables maintained by WLCG.

6.3 ARDA

The milestone of producing a new plan is delayed. Discussion with the Experiments have taken place, but have not been formalised in a new plan. In a nutshell, All 2007 activities to continue along the lines of 2006, with the notable exception of ATLAS, where the scope of the collaboration has been expanded to support their distributed data management activities and on the dashboard project, which is now involving all 4 Experiments (in 2006 only ATLAS and CMS).

Dashboard - The activity is continuing (evolution of the toolkit and operation of the system for all 4 Experiments). Starting mid February all 4 experiment dashboards are in production. The system is evolving in order to make conveniently available more information, to interpolate missing information lost before entering the dashboard data base. They are also preparing to better monitor jobs submitted via CondorG or executed via pilot-job mechanism (i.e. several users' jobs executed by a single pilot).

Ganga - The activity is continuing focusing on the evolution of the system, tutorial and support of users. Since the beginning of the year, Ganga users are monitored, collecting information about user, Ganga version and TCP/IP domain.

LCG Quarterly Report – 2007 Q1 – Executive Summary

During this quarter the Ganga team, not only the CERN members, has animated a significant series of tutorials (e.g. Edinburgh, Milan, Lyon, Munich, and Taipei). The very intense tutorial effort is having visible effects: during this quarter over 400 different users have tried the system at least once. Around 50 regular users are using Ganga on daily basis. Around 2/3 of the users belong to ATLAS and 1/3 to LHCb; and a sizeable ~10% from non-HEP communities are using the system.

Middleware - During the quarter, several EGEE/LCG components have been tested: LFC (new interface allowing bulk operations); WMS (reliability studies); FTS2.0 (in the framework of the ATLAS Tier0 activity); SRM (contribution to the general testing); xrootd/storage manager integration (ALICE).

Job Reliability Reports - Started to publish (within the LCG MB) efficiency numbers for the Tier0/Tier-1s as seen by selected applications. In addition a number of web pages to ease the sites to understand eventual failures and pin down the origin has been discussed and is being put in production.

ATLAS DDM - The so-called ARDA monitor of the ATLAS DDM using dashboard technology has been adopted as an official monitoring tool and is in use in production and in the Tier0 activities.

ALICE Analysis - ALICE Distributed Analysis is growing and in the user community using the developed analysis infrastructure has grown to 25 regular users. 500,000 user jobs have been successfully executed.

7 Experiments

7.1 ALICE

AliRoot Framework - The development of the AliRoot software has entered the last phase leading to the version to be used for the processing of the data from the first LHC run scheduled for the end of this year. Work is still going on to reduce memory and CPU consumption, to make the code more robust and to reduce the size of data (raw and ESD). The common AOD data structure has been prototyped. The analysis framework has been released and adopted by users. The access to calibration data in the OCDB is fully operational and is being optimized further. The implementation of the detector geometry as installed is in progress including cables and services. The validation of the Fluka transport is progressing. The integration of the GEANT 4 package into AliRoot is completed and validation is starting.

Calibration and Condition DB - The detector visualisation is fully integrated in ROOT and is available for almost all detectors. Alignment and calibration is implemented in the reconstruction algorithms for almost all detectors but the integration of the survey data is mostly missing. The framework to calculate the calibration parameters online (DAQ, DCS, HLT) and to feed them into the Offline Condition Data Base is fully operational. The implementation of the detector specific algorithms within this framework has started and several have been tested.

ALICE Grid Tools - The development of a new version of the AliEn file catalogue with faster access and improved scalability is being finalized and will be released in April. The interface to the Nordic ARC GRID is still under development. The development of the interface to the US OSG GRID is pending the approval of the ALICE-USA computing plan. The ALICE storage architecture based on xrootd is fully operational. Interfaces to the LCG storage solutions are in advanced tests and validation stage (dCache) or close to be released for testing (CASTOR2 and DPM). Further validation of xrootd interface in the LCG SE implementation CASTOR2 and DPM is on-going.

Physics Data Challenge - The currently running 2007 Physics Data Challenge is a continuation of the PDC 06. Its main goals are to produce events required by the Physics Working Groups and to test and validate the new releases of the Grid and application software. In the past 11 months, 1200 jobs on average were running over the entire period. During this challenge new sites will be added to the ALICE Grid, the new components of the LCG Grid and of the application software will be validated with special emphasis on standard LCG storage solutions. All sites which have pledged resources to ALICE take part to the challenge. However, a substantial amount of resources still need to be set-up in Poland, NDGF, KISTI, China, and US.

7.2 ATLAS

ATLAS is continuing their Data Streaming tests that were extended and will go through the spring 2007. As mention in the 3D section, ATLAS also successfully used the 3D database in several Tier-1 sites and the problem with Oracle streams replication has been solved.

The Tier-0 tests including SRM and data distribution to Tier-1 and Tier-2 sites was not completed in 2007Q1 because of the problems encountered with CASTOR. The main problem has been in the data management of which CASTOR is an essential part but it is not CASTOR alone.

ATLAS realizes the full impact of all the services they use: DDM, SRM, CASTOR, FTS, GridFTP, TCP/IP, and then the network and then the way back up into dCache/CASTOR/DPM on the other end. The interplay of those systems is difficult and it is only now they are beginning to test the full chain. All individual components have been tested before but mostly only with the layer just above to drive it.

The overall plan was to do Tier-0 tests and data exports in March, then run together with CMS and ALICE in April and do a second final Tier-0 test and data export in May and start the final dress rehearsal in June. To understand the new stager and the new software an estimated minimum of at least 4 weeks is needed, in the best case the position at the end of May will be that planned for March meaning a 3 month delay.

Outstanding issues are site, mainly Storage Element, instabilities that affect production activities and data transfer tests. The user access to the data on the Grid using DDM tools also needs attention in the coming quarters.

7.3 CMS

Quarterly report not received.

7.4 LHCb

For the alignment and calibration the challenge has been delayed, all the LHCb sub-detectors will have their code in place in June 2007. During 2007Q1 replication of conditions data across available Tier-1 using 3D was successful.

The software release for the Physics Book Production was completed and its production commenced. It provides MC production samples. There were issues with access to data in order to perform reconstruction and stripping of data. Analysis is yet to really get underway.

Data transfers from DAQ to the Tier-0 are about to start and will be organised by the LHCb Online project. It will need to reach 70 MB/s, aiming at the target of 200 MB/s.

Appendix 1: WLCG High Level Milestones Status (31 March 2007)

31.03.2007		WLCG High Level Milestones - 2007												
ID	Date	Milestone	ASGC	CC IN2P3	CERN	FZK GridKa	INFN CNAF	NDGF	PIC	RAL	SARA NIKHEF	TRIUMF	BNL	FNAL
WLCG-07-01	Feb 2007	24x7 Support Definition Definition of the levels of support and rules to follow, depending on the issue/alarm												
WLCG-07-02	Apr 2007	24x7 Support Tested Support and operation scenarios tested via realistic alarms and situations												
WLCG-07-03	Jun 2007	24x7 Support in Operations The sites provide 24x7 support to users as standard operations												
WLCG-07-04	Apr 2007	VOBoxes SLA Defined Sites propose and agree with the VO the level of support (upgrade, backup, restore, etc) of VOBoxes												
WLCG-07-05	May 2007	VOBoxes SLA Implemented VOBoxes service implemented at the site, and tested by the Experiments												
WLCG-07-06	Apr 2007	Job Priorities Available at Site Mapping of the Job priorities on the batch software of the site completed and information published												
WLCG-07-07	Jun 2007	Job Priorities of the VOs Implemented at Site Configuration and maintenance of the jobpriorities as defined by the VOs. Job Priorities in use by the VOs.												
WLCG-07-08	Mar 2007	Accounting Data published in the APEL Repository The site is publishing the accounting data in APEL. Monthly reports extracted from the APEL Repository.												

LCG Quarterly Report – 2007 Q1 – Executive Summary

WLCG-07-09	Mar 2007	3D Oracle Service in Production Oracle Service in production, and certified by the Experiments												n/a
WLCG-07-10	May 2007	3D Conditions DB in Production Conditions DB in operations for ATLAS, CMS, and LHCb. Tested by the Experiments.												
WLCG-07-11	Depl. Date + 30d	SL4 Operational at Site (for WN and UI nodes) This has to happen within 30 days after the release from GD.												
WLCG-07-12	Jun 2007	Site Reliability above 91% Considering each 11 Tier-1 sites												
WLCG-07-13	Jun 2007	Best 8 Sites above 93% Eight sites should reach a reliability above 93%												
WLCG-07-14	Dec 2007	Site Reliability above 93% Considering all 11 Tier-1 sites												
WLCG-07-15	Dec 2007	Best 8 Sites above 95% Eight sites should reach a reliability above 95%												

WLCG High Level Milestones Status (30 April 2007)

30.4.2007		WLCG High Level Milestones - 2007												
ID	Date	Milestone	ASGC	CC IN2P3	CERN	FZK GridKa	INFN CNAF	NDGF	PIC	RAL	SARA NIKHEF	TRIUMF	BNL	FNAL
WLCG-07-01	Feb 2007	24x7 Support Definition Definition of the levels of support and rules to follow, depending on the issue/alarm	Red	Green	Green	Red	Green	Red	Red	Red	Red	Green	Green	Green
WLCG-07-02	Apr 2007	24x7 Support Tested Support and operation scenarios tested via realistic alarms and situations	Orange	Green	Green	Orange	Orange	Orange	Orange	Orange	Orange	Orange	Green	Green
WLCG-07-03	Jun 2007	24x7 Support in Operations The sites provides 24x7 support to users as standard operations												
WLCG-07-04	Apr 2007	VOBoxes SLA Defined Sites propose and agree with the VO the level of support (upgrade, backup, restore, etc) of VOBoxes	Orange	Orange	Green	Orange	Orange	Orange	Orange	Orange	Orange	Orange	Orange	Green
WLCG-07-05	May 2007	VOBoxes SLA Implemented VOBoxes service implemented at the site, and tested by the Experiments												
WLCG-07-06	Apr 2007	Job Priorities Available at Site Mapping of the Job priorities on the batch software of the site completed and information published	Orange	Orange	Green	Orange	Green	Orange	Green	Green	Green	Orange	Orange	Orange
WLCG-07-07	Jun 2007	Job Priorities of the VOs Implemented at Site Configuration and maintenance of the jobs priorities as defined by the VOs. Job Priorities in use by the VOs.												

LCG Quarterly Report – 2007 Q1 – Executive Summary

WLCG-07-08	Mar 2007	Accounting Data published in the APEL Repository The site is publishing the accounting data in APEL. Monthly reports extracted from the APEL Repository.													
WLCG-07-09	Mar 2007	3D Oracle Service in Production Oracle Service in production, and certified by the Experiments													n/a
WLCG-07-10	May 2007	3D Conditions DB in Production Conditions DB in operations for ATLAS, CMS, and LHCb. Tested by the Experiments.													
WLCG-07-11	Depl. Date + 30d	SL4 Operational at Site (for WN and UI nodes) This has to happen within 30 days after the release from GD.													
WLCG-07-12	Jun 2007	Site Reliability above 91% Considering each 11 Tier-1 sites													
WLCG-07-13	Jun 2007	Best 8 Sites above 93% Eight sites should reach a reliability above 93%													
WLCG-07-14	Dec 2007	Site Reliability above 93% Considering all 11 Tier-1 sites													
WLCG-07-15	Dec 2007	Best 8 Sites above 95% Eight sites should reach a reliability above 95%													

Appendix 2: Tier-0 and Tier-1 Capacity Available and Required

1Q2007	Tier 1 Capacity: Available vs. Required (Scheduled)									Scheduled Capacity Required by LHC Experiments											
WLCG Site	CPU KSi2K			Disk TB			Tape TB			ALICE			ATLAS			CMS			LHCb		
	2006/7 pledge	Avail	Req.	2006/7 pledge	Avail	Req.	2006/7 pledge	Avail	Req.	CPU	Disk	Tape	CPU	Disk	Tape	CPU	Disk	Tape *	CPU	Disk	Tape
ASGC	950	950	313	400	400	85	500	500	122				133	15	29	180	70	93			
CC-IN2P3	1170	1171	1231	520	280	307.2	535	500	501	400	161	319	234	30	65	180	70	93	417	46.2	24
FZK/GridKa	1030	1030	1267	280	280	311.2	393	395	695	670	167	532	182	28	58	180	70	88	235	46.2	17
INFN/CNAF	1800	1800	1202	850	450	384.2	850	450	449	433	161	213	130	17	31	420	160	189	219	46.2	16
NDGF	520	505	577	160	202	133	240	225	233	482	120	213	95	13	20						
PIC	250	600	390	140	70	131.2	158	240	123				95	15	27	180	70	81	115	46.2	15
RAL	980	775	545	450	171	136.2	664	286	136				130	20	40	180	70	77	235	46.2	19
SARA-NIKHEF	306	742	742	170	129	128.2	143	187	187	131	54	106	225	28	60				386	46.2	21
TRIUMF	150	250	92	12	12	14	12	12	22				92	14	22						
US-ATLAS BNL	1120	1300	416	520	420	50	300	300	110				416	50	110						
US-CMS FNAL	728	2000	720	100	700	280	250	350	343							720	280	343			
US-ALICE		180	131		45	48		35	106	131	48	106									
TOTALS	9004	11303	7626	3602	3159	2008.2	4045	3480	3027	2247	711	1489	1732	230	462	2040	790	964	1607	277	112