CERN/LHCC 2006-031
LHCC-G-119
28 September 2006

LARGE HADRON COLLIDER COMMITTEE

LCG PROJECT COMPREHENSIVE REVIEW
September 2006

1. EXECUTIVE SUMMARY

The fourth annual LHCC Comprehensive Review of the LCG Project took place on 25-26 September 2006. The LHCC referees addressed the following areas: Management and Global Collaboration, Middleware, Applications Area, CERN Fabric, Distributed Fabric, Grid Deployment & Operations, Service Challenges, and Project Planning & Communications. The LHCC acknowledges the considerable amount of work that has gone into the preparation of the LCG Project Comprehensive Review.

The LHC Computing Grid (LCG) Project was created by the CERN Council in September 2001 with the aim of prototyping and deploying the computing environment for the LHC experiments. The formal launch of the project was at a workshop held in March 2002. Since that time, the LCG has demonstrated progress towards the realization of the computing requirements of the experiments in time for LHC operation in 2007.

The LCG Project is a collaboration of the LHC experiments, the Regional Computing Centres, CERN and the physics institutes with the aim of preparing and deploying the computing environment that will be used by the LHC experiments to analyze the LHC data. The project includes support for applications and the development and operation of a computing service.

The LCG Project is divided into two phases. Phase I (2002-2005) had the objective of building a service prototype, based on existing Grid middleware, of running a production Grid service and producing the Technical Design Report for the final system. Phase II (2006-2008) is building and commissioning the initial LHC computing environment. The LCG is not a Grid development project and it relies on other Grid projects for the middleware development and support.

The LHCC considers that the LCG Project has shown significant progress since the last Comprehensive Review. In particular, the interoperability between the various Grids, particularly the Enabling Grids for e-Science in Europe (EGEE) and the Open Science Grid (OSG), has improved considerably, the Worldwide LCG (WLCG) has successfully deployed the gLite 3.0 middleware, the CERN computer centre is on schedule to be ready to handle the requirements for the Tier-0 and Calibration Analysis Facility (CAF) for LHC start-up in 2007, the data transfer rates achieved in the Service Challenges, although not meeting fully the original targets, are compatible with the latest estimated requirements from the experiments, and the planning, reporting and reviewing system has improved significantly.

Moreover, the Collaborations are updating their resource requirements in view of the revised LHC schedule for the years 2007-2010. Initial indications show that the modification in the schedule reduces the gap between the required and available resources for these years, but the total level of funding requested will be required to complete the LCG Project.

However, the Committee did note some concerns. The overall stability of the service needs to be improved, and a reliable, unattended operation has not yet been achieved, particularly for the Tier-1 and Tier-2 sites. A complete test of the entire chain from the DAQ to the physics analysis is still lacking, although many pieces have been tested. Attention should be given to the storage management systems – the CASTOR2 performance and stability needs to be improved both for the Tier-0 and Tier-1 sites, the new version of the Storage Resources Manager interface needs to be implemented in all of the mass storage systems.

The conclusions and concerns of the LHCC are given below. They will help the Committee to follow up outstanding issues and to monitor future progress of this project in forthcoming sessions of the LHCC prior to the next LCG Project Comprehensive Review one year hence.

2. OVERVIEW

- Good progress was reported on the interoperability between the various Grids, with much effort being put into authentication, job submission and mass storage access across the Grid borders. The future role of the Nordic Data Grid Facility within the WLCG needs to be clarified. The Collaborations are updating their resource requirements in view of the revised LHC schedule for the years 2007-2010. Initial indications show that the shortfall in resources is reduced for these years, but the total level of funding requested will be required to complete the LCG Project.
- WLCG has successfully deployed gLite 3.0 and is improving the usage of the EGEE and OSG Grids. The experiments presented a coordinated plan of the necessary functions and services of the middleware that have still to be deployed. Adequate manpower should be devoted to fixing of bugs and to consolidating the existing code.
- Good progress was reported on the Applications Area, with no major concerns having been identified. Action must be taken soon to ensure the required level of manpower remains available beyond 2007. The experiments are encouraged to make decisions on the use of PROOF.
- The CERN computer centre infrastructure is on track to be ready to handle the known Tier-0 and CAF requirements for LHC start-up in 2007. The funding and manpower situation of the CERN Fabric has improved considerably.
- The Distributed Fabric is well suited to the challenges it faces for LHC production and analysis. Tier-1 sites should improve communication with the experiments and local Tier-2 sites, and focus on stability. Middleware software developers must concentrate on stability over functionality. Attention should be given to implement the Storage Resource Manager interface to the storage management systems. The performance and stability of CASTOR2 needs to be improved for both the Tier-0 and Tier-1 sites. The 3D Phase 2 Tier-1 sites should move rapidly to deployment. Experiments should provide realistic estimates of their database requirements.
- Considerable progress was reported on the deployment of the Grid services and operations, highlighted by the rapidly increasing usage of the EGEE and OSG Grids. The Committee recommends that emphasis is put on improving further the stability and reliability of EGEE and OSG services and that the experiments get further involved in the first line support for users in their organization.
- There has been a significant amount of work and progress in the services during the experience gained through the Service Challenges. The rates achieved, although not meeting fully the original targets, were compatible with the latest estimated nominal requirements from the experiments, and are considered to be an important result. The experiments have successfully used the Grid for many tasks, albeit with some problems that need to be corrected.
- The planning, reporting and reviewing system of the LCG has improved significantly in Phase 2 of the WLCG Project.

3. MANAGEMENT AND GLOBAL COLLABORATION

The organizational structure of the WLCG is essentially unchanged since the last LCG Comprehensive Review. The activity area `Distributed Analysis & Grid Support' has emerged, and includes now the ARDA (A Realization of Distributed Analysis for the LHC) activity with a focus on analysis.

The WLCG infrastructure is based on the individual infrastructure of the science Grids EGEE and OSG. Interoperability had been a major concern at the last LCG Comprehensive Review, but in the mean time, much work has been invested to support authentication, job submission and mass storage access across the Grid borders. While this enables the Collaborations to successfully combine both Grids in their production models, this integration is by no means effortless. The role of Nordic Data Grid Facility (NDGF) within WLCG appears to have significantly diminished compared to a year ago, and there are presently doubts regarding the services and service levels that the Nordic Tier centre will provide. The situation should be clarified between WLCG and NDGF. Since the Nordic Tier-1 centre is expected to play a significant role for the computing of several LHC Collaborations, it is important that interoperability issues with the Advanced Resource Connector (ARC) middleware are addressed by WLCG and not deferred as internal problems of these experiments.

Monitoring of availability has made progress and is now routinely available for the Tier-1 centres. Monitoring and classification of job failures turned out to be a much harder task. While presently much manual work is still required for job failure analysis, it is important to push vigorously further for an automated system that will deliver key information for improving the performance and reliability of the Grid centres.

Accounting has made significant progress. For about four months running, comprehensive accounting information on the use of CPU and storage resources has been collected routinely from CERN and the Tier-1 centres. Monthly use of the resources is relatively low, about half of the installed resources, which is not unexpected given the present use patterns, but indicates that any observed performance bottlenecks are usually not due to resource limitations.

After the recent update of the LHC running schedule for 2007-2010, the experiments have been requested to update their resource requirement estimates given in the Computing Technical Design Reports. While official numbers are not yet available, preliminary figures indicate a significant reduction of the funding shortfall for WLCG material at CERN for Phase 2. However, the total level of funding requested will be required to complete the LCG Project. The balance of WLCG personnel beyond spring 2008 depends critically on the establishment of a successor project for EGEE-II. WLCG should aim for a consolidation in terms of a structural project, and this is equally important for the remote centres.

Good progress was reported on the interoperability between the various Grids, with much effort being put into authentication, job submission and mass storage access across the Grid borders. The future role of NDGF within the WLCG needs to be clarified. The Collaborations are updating their resource requirements in view of the revised LHC schedule for the years 2007-2010. Initial indications show that the shortfall in resources is reduced for these years, but the total level of funding requested will be required to complete the LCG Project.

4. MIDDLEWARE

The LHCC congratulates the WLCG for the successful deployment of gLite 3.0 and the rapidly increasing performance of the EGEE and OSG Grids. The Committee recommends that emphasis is put on improving further the stability, reliability and interoperability of EGEE and OSG services, and that the experiments become more involved in the first line support for users in their organization.

The experiments presented a coordinated plan on components and services considered critical but still not fully deployed in the current implementation of the middleware release. Historically, and partially due to persistent delays on the availability of middleware packages, the experiments have developed independent solutions to similar problems, relying nonetheless on common core middleware components.

The experiments, endorsed by the Committee, emphasized the need to achieve stability and robustness on common tools related to file transfer as well as resources and data management. It is essential that adequate manpower be devoted to fixing of bugs and to consolidating the existing code.

Available middleware tools for monitoring the status of Grid nodes, usage accounting, and job priority assignment do not yet meet the needs of the experiments. The Committee notes that these issues, together with activities which aim to guarantee their long-term reliability, constitute main items within the EGEE and OSG short-term plans.

WLCG has successfully deployed gLite 3.0 and is improving the usage of the EGEE and OSG Grids. The experiments presented a coordinated plan of the necessary functions and services of the middleware that have still to be deployed. Adequate manpower should be devoted to fixing of bugs and to consolidating the existing code.

5. APPLICATIONS AREA

The Applications Area develops and maintains that part of the physics applications software and associated infrastructure that is shared among the LHC experiments. The scope includes common applications software infrastructure, frameworks, libraries, and tools, together with common applications such as simulation and analysis toolkits, Grid interfaces to the experiments and assistance with the integration and adaptation of physics applications software in the Grid environment.

The LHCC appreciates the amount of work done in the context of the Application Area projects and the achievements reached. The progress is, however, not uniform in the various projects and subprojects, which is due to the intrinsic fragmentary nature of parts of the projects and to limitations of manpower such as in the Software Process Infrastructure (SPI) project, causing difficulties in managing the activities.

The LHCC invites the Simulation project to overcome the present difficulties in interfacing some Monte Carlo generators to the LCG simulation infrastructure. The LHCC takes note with satisfaction of the achievements, in particular in the sectors of consolidation and fast access to data, and of the ROOT interactive tool for analysis, which is properly managed and has appropriate manpower resources. The merger of the ROOT and SEAL (Shared Environment for Applications at LHC) is continuing successfully. The development of PROOF, a prototype extending ROOT to run on a distributed, heterogeneous system, is progressing satisfactorily. The experiments are expressing interest in using PROOF. Clear decisions by the experiments are urged since the use of PROOF has major implications in the design of the computing centres.

The constant effort to keep a strong and effective link with the experiments and the individual users via the direct input of the experiments at the Applications Area management level, via regular dedicated meetings and with appropriate interface tools (SAVANNAH service for bug tracking and task management) is certainly beneficial to the effectiveness of the Applications Area projects and the LHCC encourages progress in this direction.

The present manpower level is globally very near to the needs and some reassignments of manpower can cure the limitations affecting individual projects. A possible manpower crisis can take place in 2008, due to retirements and contract terminations. Appropriate action has to be taken during next year to keep, also after 2007, the manpower level adequate for the project completion and for the software maintenance phase beyond 2007.

Good progress was reported on the Applications Area, with no major concerns having been identified. Action must be taken soon to ensure the required level of manpower remains available beyond 2007. The experiments are encouraged to make decisions on the use of PROOF.

6. CERN FABRIC

The CERN computer centre infrastructure is on track to be ready for the LHC startup. Impressive progress has been made regarding the power, cooling, and building areas. The funding situation has improved significantly, though it is still slightly short of the pledged resources.

The technology available on the market, with minor shortcomings in the disk storage area, is sufficiently advanced and scalable to handle the known Tier-0 and the CAF requirements. The storage systems´ performance has improved significantly. Yet, Service Challenge 4 (SC4) failed to reach the nominal data transfer rates from Tier-0 to Tier-1 due to stability problems at the Tier-1 centres.

The computing, storage, and network infrastructure of the CERN computing centre will need to serve two LCG applications - the Tier-0 and the CAF. The Tier-0 requirements are sufficiently clear and have been verified in an ATLAS full scale test which exceeded the nominal data rates. Further tests will have to include phases of unattended running, integration of DAQ, and Tier-0 to Tier-1 data transfers as well as full scale tests of the four experiments in parallel.

In contrast to the Tier-0, the application footprint of the CAF is not yet well defined. Choosing analysis tools like PROOF may have an impact on the hardware selection for the CAF. Hence, such a decision should be made in time for the Fabric managers to react with appropriate purchases and procurements, i.e. 6-9 months before full CAF functionality is actually expected.

Tape libraries, network components, and wide area network bandwidth have reached a level which already fulfils the LHC requirements and are able to scale beyond. Commercially available solutions fit well into the CERN Fabric. Driven by the recent release of Intel´s dual core Xeon CPUs, both the CPU price/performance ratio and the performance/Watt ratio have leapt forward, thus reducing the number of boxes to be installed at CERN and the infrastructure requirements. This is one factor in the more optimistic budget forecast. However, this is most likely a one-off event. Further improvements will be moderate increases of CPU clock frequency and number of cores per chip.

The disk systems available on the market now easily fulfill the capacity requirements of the CAF and Tier-0. However, while the commodity market delivers ever larger and cheaper SATA disk drives, these are not easily integrated into systems that fulfill the access bandwidth requirements. The Tier-0 and CAF storage systems will have to be designed with both capacity and throughput in mind, at a cost level that is above that of

systems that only offer high capacity. Again, optimizing the storage hardware requires timely input from the physics groups.

The funding situation of the CERN Fabric has improved significantly. Major factors for this are the reduced, albeit preliminary, requirements of the experiments for 2007 following the modified LHC operation plans. A second significant contribution comes from the improved CPUs that are now entering the market. Overall, the Tier-0 and CAF are still slightly under-funded with respect to the claims made, and budget cuts by funding agencies for 2007 might make the situation worse. A side effect of the reduced 2007 requirements is an even steeper ramp up of hardware resources in 2008.

The manpower situation of the CERN Fabric seems tight, but manageable. Care should be taken regarding employment contracts ending in 2008, such as those in the EGEE-2 project. A loss of manpower in this critical year could significantly interrupt operations in all Tier centres and development in software projects.

The CERN computer centre infrastructure is on track to be ready to handle the known Tier-0 and CAF requirements for LHC start-up in 2007. The funding and manpower situation of the CERN Fabric has improved considerably.

7. DISTRIBUTED FABRIC

The Distributed Fabric provides resources at national Tier-1 and regional/institutional Tier-2 sites for production and analysis of simulated and reconstructed data. This requires the infrastructure to store, process, and distribute bulk data. Distributed databases are needed for detector conditions, calibration, geometry, event tags, and file metadata catalogues.

Tier-1 and Tier-2 sites have taken enormous strides over the last year. Production operation has been clearly demonstrated by the SC4, in which all Tier-1 and many Tier-2 sites participated. The issues now lie in the area of scaling, stability, and support - all of which highlight the need for improved communication between sites and with the experiments.

The recent change in the LHC start-up schedule will have a direct bearing on resource procurement at Tier-1 and Tier-2 sites. This is particularly an issue for Tier-1 sites in Europe, where large orders must go through a European Union tender process, which can take 6 to 9 months. Tier-1 sites must be fully integrated into the experiments' planning and decision processes to allow them to understand the consequences of this, and any future, requirement changes.

While all Grid computing centers are preparing for 24x7 operations, the expectation of having expert debugging capabilities at this level cannot reasonably be fulfilled. The stability of the middleware software is crucial, both to reduce problems and the necessity of frequent upgrades. Therefore, Grid services have to be made more fault tolerant as they mature. The LHCC stresses that developers should concentrate on reliability and providing well-tested releases over functionality, and considers that conventional load balancing and clustering mechanisms as already used for Internet services might improve both stability and scalability also in the Grid context. This is particularly important for the Tier-1 sites, where any downtime will not only impact the aggregate CPU available but, during running, will necessitate buffering new data at the CERN Tier-0, and increased load to catch up once the problem is resolved.

Having Tier-1 support staff on-call to resolve problems significantly improves site stability. This is especially important to resolve issues between Tier-0 and Tier-1 sites, where experts on both sides must be available to work together. Many Tier-1 sites have implemented such a system, and others are strongly encouraged to follow suit wherever possible. However, not all problems can be resolved by those on-call - this can only reduce, not eliminate, outages. Scheduled outages should be carefully coordinated with the experiments to reduce the chances of multiple Tier-1 sites being down at the same time.

The Tier-1 sites are managed by dedicated support staff, working in computer centres. In order to support the needs of the experiments and feed back the experimental requirements, it is vital to bridge any "culture gap" between the experiments and Tier-1 centres. The Committee recommends that Tier-1 centres appoint one or more experiment liaison officers, who should attend experiment computing meetings. Similarly, experiments should have a Tier-1 liaison officer to meet regularly with each Tier-1 site's support staff. This person could also be a first line of support within the experiment on Tier-1 issues.

A number of Tier-2 centres are federations of geographically and administratively distinct sites. The sites are not federated at the middleware level. Federation can nevertheless have the advantage of sharing support load, expertise, and some services - provided there is close cooperation between sites, perhaps with a dedicated Tier-2 coordinator.

Collaboration between Tier-1 and Tier-2 sites is essential to maintain the data transfers. The Tier-1 sites are also encouraged to provide advice, and perhaps some services, to the Tier-2 sites. On this point, the Committee notes with concern the paucity of Tier-1 sites in relation to Tier-2 sites for CMS in Europe.

With CASTOR2, dCache, and DPM, the WLCG project has embraced three mass storage systems with very different complexity. While CASTOR2 expects and includes a hierarchical storage manager (HSM) and is in use at Tier-0 and some Tier-1 sites, dCache can interface to several existing HSMs as well as manage “disk only” pools, while DPM is solely a disk pool manager mostly used at smaller Tier-2 sites.

All three projects are currently implementing within a highly constrained and short time-scale the Storage Resource Manager SRM v2.2 interface released in June 2006 with a clear focus on WLCG requirements. While CASTOR2 is on target, both dCache and DPM will most likely not have SRM v2.2 fully implemented for the scheduled large scale deployment starting 1 November 2006, with possible consequences for extended tests planned for early 2007. The Committee notes that parts of the SRM specification are very theoretical and that it is not clear how e.g. a transition of a large number of files from “TAPE0” to “TAPE1” is to be implemented. While SRM v3 will be developed further for a non-LCG community, the LHCC recommends that new features should be integrated with great care in the startup phase of WLCG and with stability of services in mind.

CASTOR2 has been successfully deployed at the Tier-0, although its performance still needs to be improved for LHC operations. Rolling it out at several Tier-1 sites went less smoothly and created a significant support load for the project. The external site support issues will need further attention. While the manpower headcount within CASTOR2 is back to its nominal level, several experienced members had to leave the project and left a significant gap, at least temporarily. The CASTOR review in June 2006 was generally positive towards the project, but included a warning that there would be “Many years … of periods of operational distress”. The recent addition of the “durable storage” feature raises concern because negligent use of such manually managed pools might block the system if they run full. At this stage, the Committee recommends that after the SRM v2.2 implementation, further development should focus on the known performance and stability issues, while adding features should have to step back behind that. The LHCC cannot conclude that there is sufficient (or sufficiently experienced) project manpower in CASTOR2 now.

In the dCache project, the manpower situation has improved visibly. User support and additional communication channels have been made available, and installation procedures have been streamlined and simplified. The workload for implementing SRM v2.2 is particularly high for dCache since it is the only system supporting both tape-backed and disk-only storage pools and has to interface several different HSMs. While dCache performance and stability are not a direct issue for the CASTOR2 based CERN Fabric, they are important for the Tier-0 to Tier-1 data transfer rates which have not yet met the design expectations.

The DPM user community has grown significantly, and DPM is actively being deployed at a large number of Tier-2 sites. The project core team is small and currently understaffed since an open position has not been filled. Regarding the increasing visibility of the Tier-2 sites as LHC approaches data-taking, the LHCC feels that long-term support for DPM should be assured.

Consideration should be given to certain issues in the mass storage systems. The CASTOR2 and dCache system development and rollout strategies at the Tier centres will need appropriate attention over the next year. In addition, two thirds of LCG Tier-2 sites currently rely on DPM, which is considerably simpler to install and manage than the dCache and CASTOR alternatives. The lack of manpower in the DPM project for support and necessary upgrades (for SRM v2.2 currently) is thus a concern.

The distributed database infrastructure (3D) is now ready for production at CERN and six Phase 1 Tier-1 sites. The remaining four Phase 2 Tier-1 sites will not meet the October 2006 milestone. This is a concern.

The database service and replication performance has proved sufficient for the experiments' applications tests so-far performed. However, the eventual performance cannot be confirmed until database update rates and access patterns are better known. The replication performance between Tier-0 and Tier-1 sites is less than had been expected (though probably still adequate), but more optimization should be possible.

The Distributed Fabric is well suited to the challenges it faces for LHC production and analysis. Tier-1 sites should improve communication with the experiments and local Tier-2 sites, and focus on stability. Middleware software developers must concentrate on stability over functionality. Attention should be given to implement the Storage Resource Manager interface to the storage management systems. The performance and stability of CASTOR2 needs to be improved for both the Tier-0 and Tier-1 sites. The 3D Phase 2 Tier-1 sites should move rapidly to deployment. Experiments should provide realistic estimates of their database requirements.

8. GRID DEPLOYMENT AND OPERATIONS

Under considerable time pressure, gLite 3.0 has been successfully deployed on practically all EGEE sites. A large number of bugs and problems have been uncovered and fixed, leading to a stable service presently. In the near future the code will be cleaned up without major new releases and the service will be consolidated. The test bed will be enhanced to better resemble full production conditions. Despite the proper announcement of the gLite 3.0 deployment, the communication with the deployment sites turned out to be sub-optimal and should be improved.

The EGEE organization, representing at this time more than 28000 processors and 26 PB of storage in over 190 sites in 40 countries, has matured with the Regional Operations Centres. Performance measures have been established, both for basic services as well as for services for the LHC experiments. This site monitoring will be further improved. At peak times more than 1.5 million jobs per month have been processed.

The LHCC took note that the OSG project received funding for the next five years. The OSG organization, representing now more than 15000 CPUs and 10 PB of storage spread over 96 centres, supports more than 27 different organizations, with ATLAS and CMS as the largest users presently, using approximately 20000 CPU hours per day.

For both EGEE and OSG, improving stability and reliability is a key issue.

Communication and problem tracking have considerably improved. The EGEE and OSG organization of problem tracking is largely similar. The EGEE organization is moving to a model with first line support inside the experiments and this move is encouraged by the LHCC. The OSG is already organized in this way.

Considerable progress was reported on the deployment of the Grid services and operations, highlighted by the rapidly increasing usage of the EGEE and OSG Grids. The Committee recommends that emphasis is put on improving further the stability and reliability of EGEE and OSG services and that the experiments get further involved in the first line support for users in their organization.

9. SERVICE CHALLENGES

The purpose of the Service Challenges (SC) is to test the performance and reliability of the Grid service. Following the last exercise, SC4, the LHC Grid is entering its commissioning phase until July 2007. This phase requires the continuous provision of services.

There has been considerable work and improvements in the services due to the experience over the last year through SC3 and SC4. Progress has been very positive although there are some areas which continue to cause problems. One of the main goals of SC4 was to test data output rates from the Tier-0 to the various Tier-1 sites, an essential path both for LHC data storage and analysis. The rates achieved, although not meeting fully the original targets, were compatible with the latest estimated nominal requirements from the experiments, and are considered to be a significant result.

The experiments have successfully used the Grid for many tasks, including simulation production, reconstruction and analysis. A significant number of events have been processed, although in most cases the experiments are operating at rates around a factor of ten below nominal. The experiments plan to ramp up to the nominal rates over the next 18 months. All four experiments reported that they had problems with the services with a diverse set of causes. Quantitative information is difficult to collect at present, but the most frequent reasons reported were instabilities of the Tier-1 sites, hardware failures and the interfaces to the data storage systems, and in particular CASTOR. In addition, site differences, such as firewalls, required significant effort to overcome.

Many of the items had been noted as issues at the previous LCG Comprehensive Review. The response to such issues has been to improve communication and error reporting between all sites and this is the most effective route to reducing them. However, the amorphous nature of the problems implies there will be no major breakthroughs but only a gradual progress should be expected.

Over the upcoming service commissioning period, these tests will become more realistic, with other data traffic between all the Tier centres being generated at the same time. It is important to achieve a rate well above the estimated requirements so as to clear any backlog which might occur. In particular, the full transfer matrix, including Tier-1-to-Tier-2, Tier-2-to-Tier-1, Tier-1-to-Tier-1 and Tier-1-to-Tier-0, should be included in the tests.

The maintenance of the other services throughout the commissioning period will remain a challenge. Several residual items need to be released but their implementation should not be allowed to disrupt the ongoing services. This will require careful planning which is complicated by the lack of a definite schedule for some of these items.

There has been a significant amount of work and progress in the services during the experience gained through the Service Challenges. The rates achieved, although not meeting fully the original targets, were compatible with the latest estimated nominal requirements from the experiments, and are considered to be an important result. The experiments have successfully used the Grid for many tasks, albeit with some problems that need to be corrected.

10. PROJECT PLANNING AND COMMUNICATION

The planning, reporting and reviewing system has been significantly improved in Phase 2 of the WLCG Project. Each site and project prepares a detailed milestone plan and reports quarterly its achievements and actual capacity. The site quarterly status reports are elaborated and summarized in a quarterly status and executive summary handed to the Overview Board. All the reporting is done through predefined forms to make them uniform and quantitative. The Committee commends the improvement in the planning and reporting system. The workload for the projects and sites generated by the reports are manageable. Sites should be encouraged to provide detailed planning through this system.

To ensure proper communication flow, three kinds of coordination meetings are held weekly: Experiment, Services, and Operations Coordination. These are important forums to coordinate short-term activities. The participants to the meetings should be selected carefully to ensure rapid and authoritative discussion. In particular, the experiments representative at the Experiments Coordination Meetings should have the full endorsement of the experiment management. In addition, LCG Bulletins containing the "things to know" are published bi-weekly to ensure everybody in the system is aware of what is going on. The LHCC commends the improvements in the communication system and encourages the project to make all efforts to keep running effectively.

*The planning, reporting and reviewing system of the LCG has improved significantly in Phase 2 of the WLCG Pro*ject.