

Performance of the UK Grid for Particle Physics

D.I. Britton (*Imperial College*), S. Burke (*CCLRC*), A.J. Cass (*CERN*),
P.E.L. Clarke (*University of Edinburgh*), J.C. Coles (*CCLRC*),
A.T. Doyle (*University of Glasgow*) a.doyle@physics.gla.ac.uk,
N.I. Geddes (*CCLRC*), J.C. Gordon (*CCLRC*), R.W.L. Jones (*Lancaster University*),
D.P. Kelsey (*CCLRC*), S.L. Lloyd (*QMUL*), R.P. Middleton (*CCLRC*),
D. Newbold (*University of Bristol*), S.E. Pearce (*QMUL*)
on behalf of the GridPP Collaboration.

Abstract

GridPP operates a Grid for the UK particle physics community that is fully integrated with the Enabling Grids for E-science (EGEE) and LHC Computing Grid (LCG) projects. GridPP provides CPU and storage resources at 19 sites across the UK, runs the UK-Ireland Regional Operations Centre for EGEE, provides Grid-wide configuration, monitoring and accounting information via the Grid Operations Centre, and provides support directly for its own system managers and users.

This paper discusses GridPP-wide performance and accounting methods. Over the last three years, a prototype Grid has been developed and put into production with computational resources that have increased by a factor of 100. In the last year more than a million jobs were processed and accounted on the system. In 2006 performance is being improved with emphasis placed on the improvement of data management and file handling across the Grid.

1. Introduction

A close relationship has existed between particle physics and computing for the last quarter of a century. Driven by economic, political, and performance issues Particle physicists have moved from the gold-standard of service and performance provided by mainframes, through smaller institutional based single machines, to modest sized cluster-based solutions. The Grid, a global and heterogeneous aggregation of hardware clusters, is the latest step along this path, which strives to minimise the computing cost by the use of commodity hardware; provide scalability to a size beyond that of mainframes; and deliver a quality of service sufficient for the task primarily by relying on redundancy and fault tolerance to balance the intrinsic unreliability of individual components. The Grid model matches the globally diverse nature of the particle physics experiment collaborations, providing

politically and financially acceptable solutions to an otherwise intractable computing problem.

At the LHC, the raw data are reconstructed and calibrated in a CPU-intensive process at the Tier-0 and Tier-1 centres, before being catalogued and archived as Event Summary Datasets (ESD). The data are further refined and rarefied to produce Analysis Object Datasets (AOD) and Tagged samples enabling data navigation. All these datasets may be used subsequently for data analysis and metadata is also required to be compiled and catalogued. The raw data are complimented by a comparable quantity of simulated data that are generated predominantly at regional Tier-2 sites before being processed in a similar manner to the raw data in order to understand detector performance, calibration, backgrounds, and analysis techniques. The computing requirements are enormous: in 2008, the first full year of data taking, CPU capacity of 140 million SPECint2000 (140,000 3GHz processors), 60 PB of disk storage and 50 PB of mass storage will be needed globally. The hierarchy of Tier centres represents an optimisation of the resources mapped to the functionality and level of service required for different parts of this problem. On the one hand this recognises that there are economies of scale to be gained in the management and operations of computing resources, particularly commodity hardware where there is only basic vendor support; on the other hand it acknowledges that not all parts of the problem need the same services or quality of service and that substantial benefits in cost and scale can also be gained by embracing an architecture where institutes, regions, or even countries, can plug-and-play. This, then, is the optimisation afforded by the Grid approach.

Since September 2001, GridPP has striven to develop and deploy a highly functional Grid across the UK as part of the LHC Computing Grid (LCG) [1]. Working with European EDG and latterly EGEE projects [2], GridPP helped develop middleware adopted by LCG. This, together with contributions

from the US-based Globus [3] and Condor [4] projects, has formed the LCG releases which have been deployed throughout the UK on a Grid consisting presently of more than 4500 CPUs and 0.65 PB of storage. The UK HEP Grid is anchored by the Tier-1 centre at the Rutherford Appleton Laboratory (RAL) and four distributed Tier-2s known as ScotGrid, NorthGrid, SouthGrid and the London Tier-2 [5]. There are 16 UK sites which form an integral part of the joint LHC/EGEE computing Grid with 40,000 CPUs and access to 10 PB of storage, stretching from the Far-East to North America. There are well over 2000 registered users. Managing such a large, distributed infrastructure is a non-trivial problem, and the projects have been working to develop tools and procedures to provide a production-quality service.

GridPP contributes to middleware development in a number of areas, mainly through the EGEE project. An interface to the APEL accounting system (Accounting Processor for Event Logs: an implementation of Grid accounting which parses log files to extract and then publish job information) has also been provided and is being tested. The development of the R-GMA monitoring system, discussed below, has continued, with improvements to the stability of the code and robustness of the system deployed on the production Grid. A major re-factored release of R-GMA was made for gLite-1.5. Similarly, GridSite was updated where it provides containerised services for hosting VO (Virtual Organisation) boxes (machines specific to individual virtual organisations that run VO-specific services such as data management: an approach which, in principle, is a security concern) and support for hybrid HTTPS/HTTP file transfers (referred to as "GridHTTP") to the httpc tool used by EGEE. GridSiteWiki has been developed, which allows Grid Certificate access to a wiki, preventing unauthorised access, and which is in regular use by GridPP. The cornerstone of establishing a grid is a well-defined security policy and its implementation: GridPP leads the development of that security policy within EGEE, identifying 63 vulnerability issues at the end of 2005. Monitoring and enhancements of the networking, workload management system (WMS) and data management systems have been performed in response to deployment requirements, with various tools developed e.g. GridMon for network performance monitoring, Sun Grid Engine integration for the WMS, and MonAMI a low-level monitoring daemon integrated with various data management systems.

2. Performance Review

The current phase of GridPP moves the UK HEP Grid from a prototype to a production platform. Whilst

progress can be monitored by milestones and metrics, success can ultimately only be established by the widespread and successful use of substantial resources by the community. Collecting information about Grid use is, in itself, a Grid challenge. GridPP sites form the majority of the EGEE UK and Ireland region (UKI), with RAL as the Regional Operations Centre (ROC). RAL also runs the Grid Operations Centre (GOC) which maintains a database of information about all sites and provides a number of monitoring and accounting tools that provide insight and information.

At the basic level the Site Functional Test (SFT) uses a small test job that runs at each site and determines the availability of the main Grid functions. Similarly, the Grid Status Monitor (GStat) retrieves information published by each site about its status. Their use and subsequent triggering of follow-up action, by way of the Core Infrastructure Centre (CIC) on Duty staff raising tickets against sites to resolve observed problems, has greatly helped improve the usability of grid resources. The Site Functional Test results have been captured and archived since October 2005 by the CIC portal developers. The total number of critical tests failed by GridPP sites is illustrated in figure 1.




Figure 1: Accumulated site functional test failures for GridPP sites in Q4 2005.

Several problems recorded as site failures may in fact be the result of intermittent problems elsewhere. For example, approximately 80% of the observed problems at UK sites came with the replica management test. However this test has a number of external failure points such as problems accessing the replica catalogue, BDII configuration errors and CERN storage replication problems. The job submission test also fails if the site queue is full of jobs and a response from is not received within a specified period.

The primary record of the downtime is recorded by the system administrator directly into the GOC database. The figures are extracted for regional reports that are reviewed weekly. Figure 2 shows the scheduled downtime for GridPP sites for Q4 2005. Like SFT failures, downtime is now logged in a database over which the CIC portal can be used to run queries.




Figure 2: Accumulated hours of scheduled downtime by site for Q4 2005.

The performance data may not be complete and can include additional site-initiated entries. Despite this it is clear that the same issues will affect all sites and yet there is an uneven distribution of recorded failures and downtime. This data was used to improve performance at all sites in the subsequent operation period.

The transitory nature of some of the problems is such that the above figures should be used to typify the approach taken to maintain a Production Grid service. Figure 3 shows the average CPU availability by region for April 2006 derived from sites passing or failing the SFTs. Although this particular data set is not complete (and an improved metric is being released in July) it can be seen that within Europe the UKI region (second entry from the left) made a significant contribution with 90% of the total of just of 4000 CPUs being available on average. With the introduction of the gLite 3.0 middleware and improved performance monitoring in June 2006, we plan to improve the overall reliability.




Figure 3: Available CPUs for April 2006. CPUs are deemed available (green) when the site passes the SFTs or unavailable (red) if the site fails.

In addition to the GOC database containing Grid-wide information, statistics are also recorded at the RAL Tier-1 centre using Ganglia to monitor CPU load, memory usage and queue data from the batch system. Figure 4 shows the usage by VO for 2005. Full capacity is roughly the top of the graph so the Tier-1 facility was running around 90% of capacity for the latter half of the year, though about half of this was non-Grid use by the BaBar experiment.

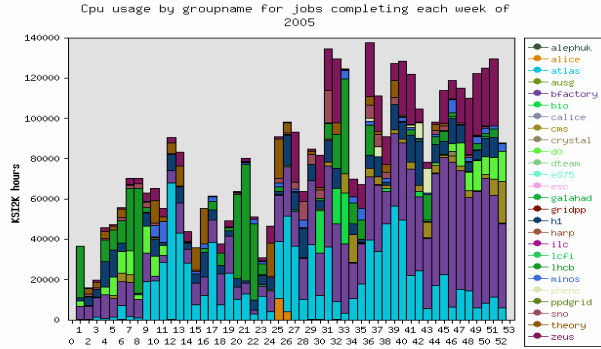


Figure 4: Tier-1 CPU use for 2005.

In order to understand the apparently low CPU utilisation in the first half of 2005, a detailed analysis of batch job efficiency was carried out where the efficiency is the ratio of CPU time to elapsed time. A highly CPU intensive batch job can achieve 95-98% utilisation of a CPU, an I/O intensive job is more likely to be around 85-95% utilisation of a CPU, and jobs waiting for busy resources can vary from 0-100% efficient. As can be seen in figure 5, the overall efficiency was rather low during the second quarter of 2005 until the applications and their data-access patterns were better understood. When CPU time is corrected by job efficiency (to give job elapsed time), it is apparent from figure 6 that the farm ran with greater than 70% occupancy for most of the year, rising to 100% in December.




Figure 5: 2005 CPU efficiency (CPU/Wall-time).




Figure 6: Tier-1 Calculated occupancy (purple curve is scaled for observed efficiency).

The global efficiency has continued to improve in 2006 with many experiments maintaining efficiencies well over 90%. The large difference in efficiency by VO shown in figure 7 illustrates that future efficiency increases will depend upon VO-specific improvements in their use of the resources: efforts are underway to improve individual applications based on this information.

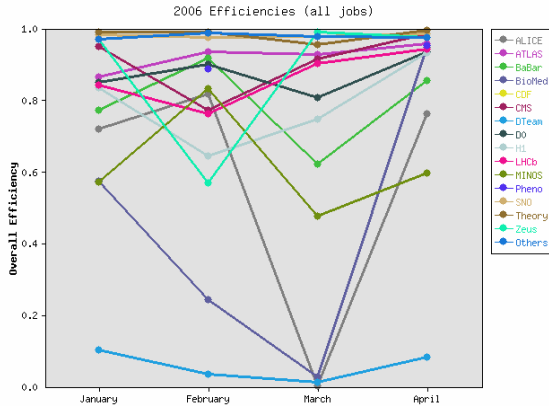


Figure 7: Job Efficiency (CPU-time/Wall-time) for 2006, by Virtual Organisation.

Whilst the statistics above address the scale, performance and efficiency of the Grid, reliability is also critical. The EGEE-JRA2 project addresses the issue of Quality Assurance and figure 8 shows the success rate of jobs run in the UK for the period May to Dec 05. Information by VO and average wait times are also available.




Figure 8: Number of jobs run (blue) versus successful jobs (orange) for UK sites in 2005.

3. Accounting Review

The hardware deployed and managed through GridPP is guided by a Memorandum of Understanding (MOU) signed by PPARC with CERN which defines the intended provision of hardware for the LHC experiments. In addition, GridPP has internal MOUs

with the Tier-2 institutes which outline the hardware provision intended. However, actual purchases are optimised to reflect the anticipated needs during the near-term future so that, overall, the hardware resources can be maximised. In 2005 the Tier-1 hardware purchase was delayed and the hardware at many of the Tier-2 sites, particularly disk, ramped up more slowly than originally planned. Tables 1 and 2 at the end of this paper show the CPU and Storage installed at the Tier-1 and Tier-2s over the last five quarters, compared with the original plans contained in the MOUs. The Tier-1 has provisioned 60% of the CPU and 70% of the storage (which includes Tape) originally envisaged. The Tier-2s have been somewhat slower to ramp-up and although significant CPU was added at the start of 2006 taking the overall provision tot 75%, the storage is still much lower than planned at 34%. All these numbers need to be understood in the perspective of the actual usage, contained in Tables 3 and 4.

The usage tables show that the resources allocated to the LCG Grid, i.e. declared via the LCG/EGEE mechanisms and monitored via the SFTs, with storage via an SRM interface. The tables also show the fraction of this allocation that was used or, more precisely, the fraction of use that was recorded by the Grid Operations Centre for CPU and the GridPP SRM Storage accounting for disk and tape. There are a number of caveats associated with the accounting system; most notably that it currently does not account usage at Cambridge and London LeSC due to their use of Condor and Sun Grid Engine respectively. A preliminary new version of APEL with Condor support has now been released to Cambridge for testing. Nevertheless, despite these known inefficiencies in the accounting, it is apparent that there was little pressure on the Tier-2 resources in 2005. Part of this is explained by a lack of confidence from the experiments in the quality of the Tier-2 storage. GridPP is working to build explicit relationships between individual sites and experiments in order to create a better understanding of needs and services.

Overall the Tier-1 is considered by GridPP to have delivered for all the experiments at the required target levels in 2005. The UK Tier-1 centre delivered 29% of the CPU of the LCG Tier-0 and Tier-1 centres. The Tier-2s are also considered to have delivered more than the required capacity. Currently, the Tier-2s account for twice the delivered CPU and 1.5 times the storage at the Tier-1. One of the challenges for 2006 is to achieve a more precise view of Grid usage, in the context of the accounting developments discussed below.

3.1. CPU

CPU accounting is relatively well developed. The collection of accounting usage records is done through R-GMA, an implementation of the Grid Monitoring Architecture (GMA). GMA models the information and monitoring system infrastructure of a grid as a set of consumers (which request information), producers (which provide information) and a registry, which mediates the communication between the producers and consumers.

In EGEE accounting, each site publishes its own accounting data using an R-GMA primary producer and its locally assigned R-GMA server. To collect the data from all participating sites, data is streamed to a centralised database via a secondary producer. The central database is located at the Grid Operations Centre (GOC) that provides a web front end, generating a summary of resource usage across the EGEE grid network. APEL is a log processing application which is used to interpret gatekeeper and batch system logs to produce accounting records. It currently supports PBS and LSF batch systems but is extensible. An accounting record is composed of (among others) the grid user, the job id of the submitted job and the resources used when executing the job. This information is typically dispersed between several different log file types such as those produced by the gatekeeper or batch system. For resource usage, a query is issued to the site's information index (GIIS) to look up the CPU performance for the computing nodes where the job was executed. APEL collects this information together and manages it within a database. APEL then joins the data to produce a list of final accounting records with all necessary details.

APEL is used to publish the generated accounting records into R-GMA where they are collated at the GOC using an R-GMA secondary producer, as shown in figure 8. APEL provides support for republishing the complete local copy of accounting records to R-GMA (in cases when the GOC was offline). It also provides a mechanism for reliable delivery using a basic integrity check to compute the number of records that were last published compared with the actual count stored on the GOC. Each accounting record is unique and there is only one record per grid job. The records may be consolidated in different ways to provide high-level views of accounting data, such as the total CPU time consumed by each VO. Figure 9 illustrates how the accounting system is used. The overall contributions from various resource providers of the global accounts for EGEE from July 2005 to June 2006 were 25 billion S12k-hours. The pie chart shows the breakdown by

resource provider illustrating that the UKI (GridPP) resources account for 1/3 of that CPU resource.

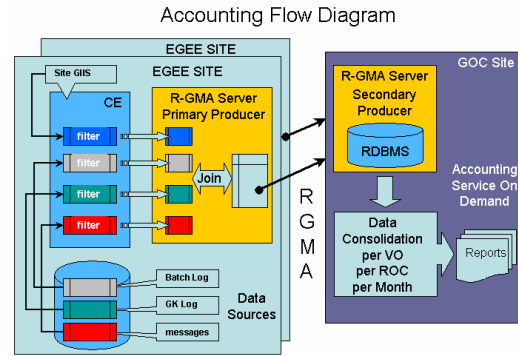


Figure 8: Accounting flow diagram providing an overview of the APEL data collection process.




Figure 9: Pie chart of CPU time by region within EGEE from July 2005 to June 2006. GridPP resources account for 33% of the total resource.

The breakdown of CPU time by VO within the UK shows that ATLAS (37%), LHCb (31%) and BaBar (10%) are the major users. The BioMed VO is the largest of the non-particle physics VOs, accounting for 4% of the UK CPU resources.

3.2. Storage

Grid file access is based upon the Storage Resource Management (SRM) interface specification. This protocol is designed to allow access to large scale storage systems on the grid, allowing clients to retrieve and store files, control their lifetimes (in volatile stores) as well as reserve filesystem for uploads etc. It is expected that each site on the grid will offer an SRM compliant storage element (SE) providing an interface to the low-level file store.

GridPP supports small-scale (DPM), medium-scale (dCache) and, in future at the Tier-1, large-scale (CASTOR) implementations. Each provide disk pool management with an SRM interface, with dCache and CASTOR offering tape support as part of a mass storage system. The SRM is a web service handling the incoming file requests which access the Unix file

systems on which data resides. When grid jobs, which need access to certain files, are submitted to a site a copy of the file exists on the local SE (it may have been there already when the job was submitted, or may have been pre-staged). The SRM does not offer random access facilities to its files, so jobs will take a copy of the file from the local SRM and place it on their local disk, enabling efficient access to the data.

Figure 10 illustrates the current status, with growth of the storage from 50TB in Aug 05 to 200TB in May 06. The problems in consistently monitoring and maintaining access to that distributed file store are evident in the day-to-day variations of the resource.

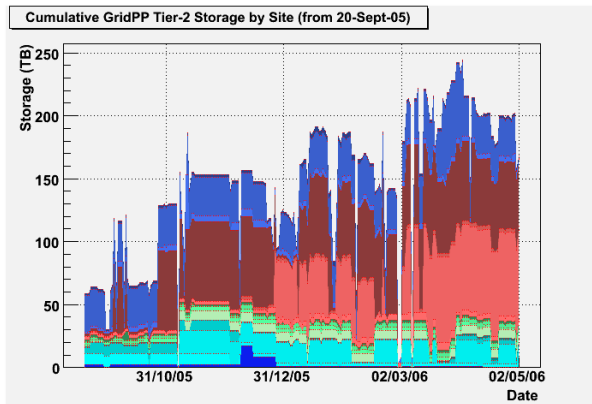


Figure 10: Histogram of Tier-2 storage capacity from September 2005 to May 2006 by region within GridPP.

In recognition of the inherent complexity and the problems of scaling Tier-2 resources to the required file storage capacity, GridPP has implemented global monitoring of the file store accessed via the SRM. It should be noted that the storage monitoring data is extracted directly from the information system, and is generally dedicated storage that is not part of a Worker Node. The data is generally a replica and not backed up to tape, but sites do need to provide a level of redundancy in cases of disk failure, so they are typically configuring their disks to use RAID 5 or 6, meaning that they must store recovery information in addition to the actual data. This reduces the capacity down to the usable level, available and accounted via the SRM.

3.3. File Transfers

The LCG has planned, and is executing, a series of world-wide Services Challenges designed to stress test the infrastructure and establish, incrementally, the levels of service and throughput needed for the LHC computing challenge. In the autumn of 2005, and through into the New Year, GridPP participated in Service Challenge 3. One aspect of this challenge was

to establish Tier-0 to Tier-1 transfer rates of 150 Mbytes/sec (disk to disk) and 50 Mbytes/sec (disk to tape). Although the initial tests in July only achieved about half these targets (and with poor stability), by January 2006 the target rates had been established. Figure 9 shows a snapshot of the Tier-0 to Tier-1 disk-disk transfers for all Tier-1s in January 2006.

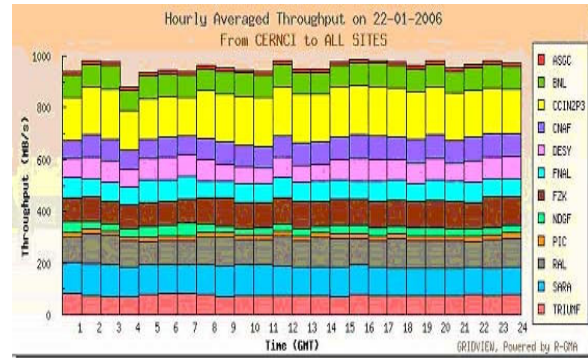


Figure 11: SC3 Tier-0 to Tier-1 throughput summary.

Currently GridPP are engaged in LCG Service Challenge 4 with goals that include ramping up the Tier-0 to Tier-1 transfer rates to full nominal rates (to tape); to identify and validate all other production data flows (Tier-x to Tier-y); to increase Tier-2 participation from 20 sites worldwide in April 2006 to 40 by September; to broaden focus from production to analysis (where there are many more users); and to streamline Operations & User Support building on existing efforts. At the time of writing, the UK Tier-1 was sustaining disk-disk transfer rates to CERN of up to 160 Mbytes/sec and Tape-Disk rates of 40 Mbytes/sec. Figure 10 shows a snapshot of disk-disk throughput for a week in April 2006 and it can be seen that the total concurrent flow from the Tier-0 is close to the target of 1600 Mbytes/sec. The UK Tier-1 at RAL took part in the test, receiving data from CERN at close to 200 Mbytes/sec.




Figure 10: SC4 Tier-0 to Tier-1 disk to disk transfer summary.

In addition GridPP is conducting transfer tests between the RAL Tier-1 and each of the UK Tier-2 sites using the File Transfer Service (FTS) developed as part of the gLite middleware stack. The target rate is a minimum of 250 Mbytes/sec for just reading or writing and 200 Mbytes/sec for simultaneous read and write. The eventual aim is to demonstrate that this can be sustained over long periods over the shared UK academic network, SuperJANET4. Initial tests of separate read and writes have now been completed with 11 of the 19 sites exceeding the targets in at least one direction and 7 exceeding them in both, as shown in table 5. The highest speeds obtained were over a lightpath (UKLight) from Lancaster to RAL where a transfer rate of over 900 Mbits/sec was sustained for more than 90 hours and 1Gbit/sec was exceeded for significant periods, as illustrated in figure 11. These rates are close to the sustained target rates between the Tier-0, -1s and -2s which are O(100 Mbytes/s) in 2008. We thus conclude that the connectivity across the academic networks are sufficient. The challenge will be to sustain these rates throughout the period of data analysis with combined data flows that will be particularly large at the Tier-0 and Tier-1 centres.




Figure 11: SC4 Tier-1 to Tier-2 file transfer summary.

4. Outlook

We now have a working Grid with resources that have grown by a factor of 100 over the last three years. In the next three years the Grid has to increase in scale by another factor of 10 and make large strides in functionality, robustness, and usability. In particular, the current Grid is largely used for fairly coordinated work by a relatively small number of approximately 250 active users in the UK. The future Grid must provide a platform, not only for coordinated production and reconstruction, but also for much more responsive (sometime called “chaotic”) use by a much larger community intent on individual analyses. The LHCb collaboration has taken some initial steps and their DIRAC workload management system enables these shorter analysis jobs to be submitted with higher internal priority compared to background production jobs.

Some of the required functionality that is currently missing includes the ability to chain jobs through the Resource Broker and the ability to pin data for subsequent access. Without these two facilities, the Grid will become grossly inefficient as jobs occupy CPU waiting for data to be staged. The whole area of data movement and data management is underdeveloped at the Grid level and individual experiments have currently resorted to proprietary solutions that have spawned the need for experiment-specific persistent services (so called VO boxes) at individual sites, which introduce security and scalability concerns. Although these are now envisaged to be limited to Tier-0 and Tier-1 sites, the better way forward would be to incorporate common services in the upper-level middleware of the Grid and experiment-specific services within the experiment software itself. Similarly, generic metadata handling remains a challenge and is a potential area of concern as the move is made to live analysis where perhaps hundreds of users are simultaneously making repeated and complex relational queries on databases that grow (in the case of Tag data) by about 4TB a year. Finally, debugging on the Grid is notoriously difficult even for experts and reproducing the conditions of a failure, logging and error reporting all need to be raised to the level expected from PBS/LSF/NQS.

In the next 18 months, GridPP needs to successfully complete the defined GridPP2 work programme; deploy the new middleware releases based on gLite-3; continue to participate in Service Challenge 4 and move into the service phase; the hardware promised in the Memorandum of Understanding with CERN, now signed by PPARC, must be provided; and the user communities must be developed and served. In parallel, PPARC has announced a call for the future support of the Grid and a proposal for an extension of GridPP2 to March 2008; a subsequent three-year GridPP3 project is being developed.

5. References

- [1] <http://lcg.web.cern.ch/LCG/>
- [2] <http://www.eu-egee.org/>
- [3] <http://www.globus.org/>, I. Foster, C. Kesselman and S. Tuecke, *International Journal of Supercomputer Applications*, **15** 3 (2001).
- [4] <http://www.cs.wisc.edu/condor/>, D. Thain, T. Tannenbaum and M. Livny, *Concurrency and Computation: Practice and Experience* **17** 323-356 (2005).
- [5] <http://www.gridpp.ac.uk/>, P J W Faulkner et al, *J. Phys. G: Nucl. Part. Phys.* **32** N1-N20 (2006).

CPU Capacity [KSI2K]							
	Delivered (i.e. Installed)					Promised	Current
	2005-Q1	2005-Q2	2005-Q3	2005-Q4	2006-Q1	MOU 2005	Ratio
London	910	910	935	1049	1049	1351	78%
ScotGrid	32	186	237	273	354	340	104%
SouthGrid	354	483	492	508	516	667	77%
NorthGrid	205	750	750	776	1783	2602	69%
Total Tier-2	1502	2329	2413	2607	3703	4960	75%
RAL Tier-1	830	830	830	830	830	1282	65%

Table 1: Delivered CPU [kSI2k] by quarter from 05Q1 to 06Q1 compared to MOU planning.

Storage Capacity [TB]							
	Delivered (i.e. Installed)					Promised	Current
	2005-Q1	2005-Q2	2005-Q3	2005-Q4	2006-Q1	MOU 2005	Ratio
London	32	13	19	37	38	102	37%
ScotGrid	9	23	36	45	45	90	49%
SouthGrid	39	40	40	48	48	46	105%
NorthGrid	14	48	84	86	132	543	24%
Total Tier-2	93	124	178	216	263	781	34%
RAL Tier-1	180	325	410	440	440	629	70%

Table 2: Delivered storage [TB] by quarter from 05Q1 to 06Q1 compared to MOU planning.

CPU Usage [KSI2K]										
	Available (i.e. Allocated to LCG)					Fraction of Allocation Used (accounted)				
	2005-Q1	2005-Q2	2005-Q3	2005-Q4	2006-Q1	2005-Q1	2005-Q2	2005-Q3	2005-Q4	2006-Q1
London	557	542	592	884	884	3.2%	13.8%	39.2%	10.3%	22.7%
ScotGrid	31	78	237	237	182	0.1%	0.8%	4.9%	6.4%	31.9%
SouthGrid	151	205	207	243	243	4.4%	18.9%	47.3%	31.0%	45.3%
NorthGrid	311	745	765	772	1777	1.1%	9.2%	10.6%	10.1%	18.6%
Total Tier-2	1050	1569	1800	2136	3086	2.6%	11.5%	22.9%	12.2%	22.7%
RAL Tier-1	444	444	444	444	444	49.8%	69.8%	67.8%	26.3%	77.0%

Table 3: CPU resources made available and accounted by LCG from 05Q1 to 06Q1.

Disk Usage [TB]										
	Available (i.e. Allocated to LCG)					Fraction of Allocation Used (accounted)				
	2005-Q1	2005-Q2	2005-Q3	2005-Q4	2006-Q1	2005-Q1	2005-Q2	2005-Q3	2005-Q4	2006-Q1
London	1.3	10.7	18.6	27.5	22.4	18.9%	7.6%	6.4%	7.4%	80.3%
ScotGrid	3.1	3.8	36.9	36.5	37.1	25.8%	24.4%	39.4%	41.2%	56.6%
SouthGrid	1.7	5.4	5.9	13.7	15.2	11.7%	3.5%	20.0%	4.8%	88.6%
NorthGrid	2.7	4.3	4.8	67.1	67.9	4.4%	6.7%	18.8%	2.1%	50.4%
Total Tier-2	8.7	24.3	66.2	144.7	142.5	15.6%	9.2%	26.9%	13.2%	60.7%
RAL Tier-1			136.2	88.4	121.1			45.9%	50.0%	46.6%

Table 4: Disk resources made available and accounted by LCG from 05Q1 to 06Q1.

Site	Inbound Outbound [MBytes/sec]		SRM	Site	Inbound Outbound [MBytes/sec]		SRM
<i>London Tier2:</i>				<i>NorthGrid:</i>			
Brunel	57	59	DPM	Lancaster	800	500	dCache
IC-HEP	80	190	dCache	Liverpool	88	22	dCache
IC-LeSC	156	95	DPM	Manchester	320	320	dCache
QMUL	118	172	DPM	Sheffield	144	414	dCache
RHUL	59	58	DPM	<i>SouthGrid:</i>			
UCL-HEP	71	63	DPM	Birmingham	317	461	DPM
UCL-Cent	90	309	DPM	Bristol	117	291	DPM
<i>ScotGrid:</i>				Cambridge	293	153	DPM
Durham	193	176	DPM	Oxford	252	456	DPM
Edinburgh	276	440	dCache	RAL PPD	397	388	dCache
Glasgow	414	331	DPM				

Table 5: SC4 Tier-1 to Tier-2 disk to disk transfer performance status (June 2006).