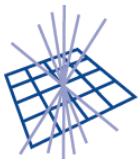

GridPP
UK Computing for Particle Physics

GridPP Project Management Board

Project Status

Document identifier :	**GridPP-PMB-150-ProjectStatus.doc**
Date:	**02/12/2010**
Version:	**1.0**
Document status:	**Final**
Author	**PMB**

Introduction

Since the last Oversight Committee, the LHC has gone from strength to strength with the √s = 7 TeV proton-proton run ending on 4 November before the transition to heavy ion collisions. The plot below shows the integrated luminosity delivered to the experiments and the rapid rise in the last few weeks of the run due to the increasing number of bunches and reduced bunch spacing.

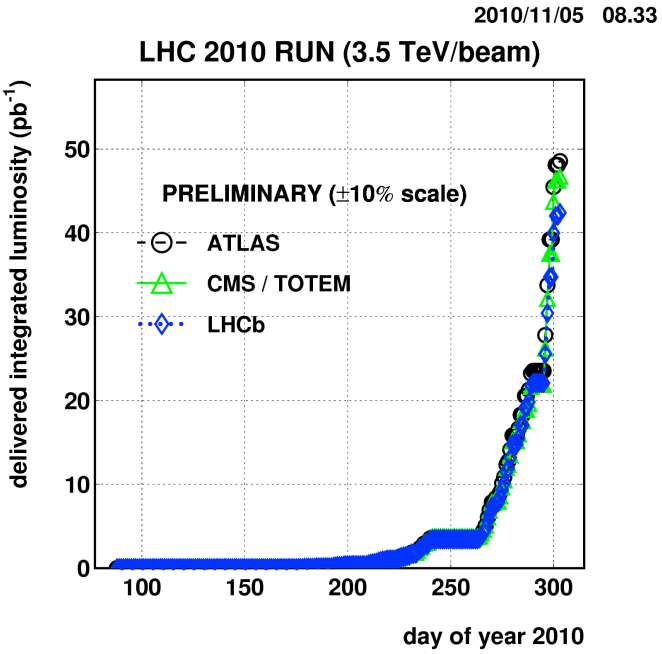

Figure-1: Integrated luminosity delivered to the experiments in proton-proton collisions during 2010 run. It should be noted that the similar figure presented to the last Oversight Committee showed the first 145 days and an integrated luminosity of 18 $\underline{nb^{-1}}$ compared to 50 pb^{-1} here.

This exceptional start-up of the LHC and the experiments has demanded, and has received, outstanding performance from the world-wide computing Grid, which has handled data rates at and above the first year design values. GridPP has played a significant role in this success. Figure-2 shows the RAL Tier-1 usage over the last year, where red is unused capacity and green is used. The two steps represent the installation of the two tranches of the annual procurement, just in time for the rising demand in July.

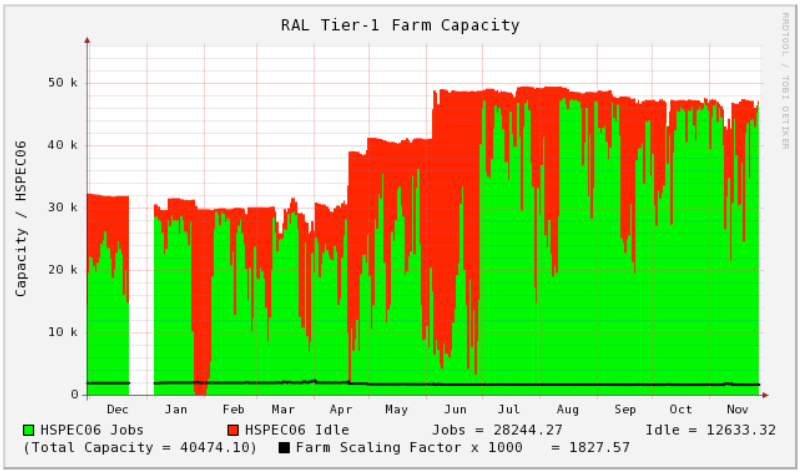

Figure-2: The RAL Tier-1 usage over the last year. The green areas show HS06 in use and the red indicates idle capacity.

Over this period, the GridPP4 proposal was taken through the remaining steps of the peer-review process, including the site visit. The project was accepted with a total budget of £28m over four years and funding initially approved for the period April 2011 to March 2013. The uncertainty introduced by the Comprehensive Spending Review has meant that only some elements of this award could be formally committed and the outstanding issues in this area continue to be a concern.

Project Status

Since the last OC meeting, quarterly reports have been produced covering second and third quarters of 2010, with the second quarter ProjectMap presented at the GridPP25 meeting. The latest ProjectMap covers the period to the end of the third quarter of 2010 (September 2010) and is online at: www.gridpp.ac.uk/pmb/ProjectManagement/GridPP3_ProjectMap_10.xls, while the reports for each quarter are at: www.gridpp.ac.uk/pmb/ProjectManagement/QuarterlyReports/reports.html. The overview of the latest ProjectMap is shown in Figure-3 below.

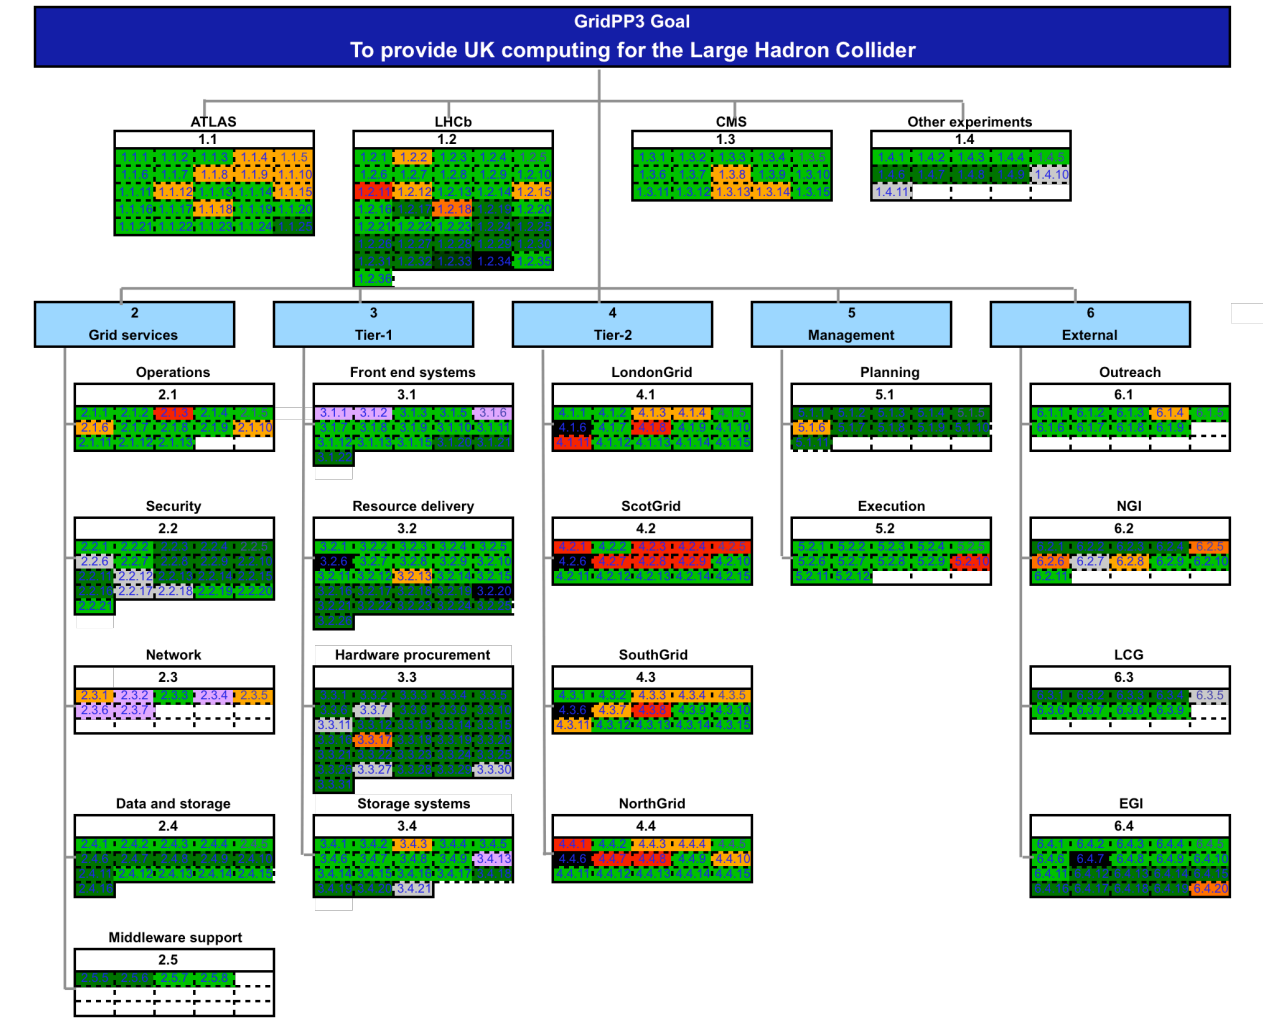

Figure-3: *The GridPP ProjectMap as of the third quarter of 2010.*

The table below shows the status of the ProjectMap for each quarter since the start of GridPP3.

	Q208	Q308	Q408	Q109	Q209	Q309	Q409	Q110	Q210	Q310
Metric OK	99	142	155	172	184	185	182	182	187	183
Metric close to target	24	47	39	32	22	21	25	26	35	34
Metric not OK	41	32	32	21	27	22	24	22	13	16
Not able to be measured	27	22	11	10	3	17	7	6	5	7
Suspended	0	6	6	9	12	9	5	6	6	8
Awaiting input	34	5	12	10	3	0	2	0	0	0
Metrics total	225	254	255	254	251	254	245	242	246	248
Milestone achieved	11	22	32	42	57	64	71	84	98	105
Milestone close to completion	0	0	0	0	0	0	2	8	5	5
Milestone overdue	2	7	13	17	4	7	7	3	5	0
Milestone not due / metric n/a	101	80	69	60	58	50	42	35	16	12
Milestones total	114	109	114	119	119	121	122	130	124	122
Total	**339**	**363**	**369**	**373**	**370**	**375**	**367**	**372**	**370**	**370**

Table-1: Evolution of the ProjectMap Milestones and Metrics to Q3 2010.

Substantial data flows from the experiments have now begun to test the Grid. Despite this, the number of 'not OK' metrics is lower than at the time of the last OC, with the 'not OK' and 'close to target' metrics reflecting a range of issues including:

- Some of the Tier-2 sites are still procuring and installing their latest round of hardware, so do not yet meet their pledged amounts of disk. This is a consequence of the delay of the 2nd tranche of Tier-2 funding from FY09 to FY10. This is discussed in more detail in the section on Deployment Status.
- CPU utilisation across the Tier-2s was lower than the target 50% overall, however the average was calculated using the capacity at the end of the quarter when in fact much of the capacity was only installed towards the end.
- Issues with sites in ScotGrid: there were problems at Durham due to staff changes and familiarity with Grid systems; ECDF had GPFS file system problems and downtime for upgrades; and Glasgow suffered from Nagios probe test failures and site BDII stability problems.
- ATLAS file losses due to server failures. This reflects issues with a particular generation of disk server, which is now being marked as read only and will be temporarily removed from service.
- LHCb also saw problems with disk servers and SRMs at the Tier-1, leading to low efficiencies and downtime. The LHCb Castor instance at RAL was also down for 2.5 days during the upgrade to 2.1.9 for LHCb. However, after a dip at the start of data taking when much LHCb production was run at CERN, the UK met 30% of LHCb worldwide production computing needs in the third quarter of 2010.

Overall, there are now no 'red' milestones, and only 5 milestones that are passed their due date but close to completion. Key points with these include:

- Issues with the second tranche 2009 disk procurement have been addressed with the vendors, and the hardware was accepted and the bill paid in November 2010, but this was after the end of the third quarter so that milestone remained orange.
- Progress is being made with detailing how GridPP and the NGS will work together as the UK NGI, including looking at Operational Level Agreements (OLAs) between the NGI and EGI. However, milestones in this area are not yet complete.

- LHCb have a milestone for successful reconstruction and stripping of real data at the UK Tier-1 to be sustained for > 1 month. Although, strictly, this has been met, the recent load problems seem to have been caused by workflows consisting of simultaneous reconstruction, merging and user analysis, so this milestone is set at 'orange' until this year's data is reprocessed.

Risk register

The GridPP3 risk register has been revised and updated by the PMB, and the latest version can be found at: http://www.gridpp.ac.uk/pmb/ProjectManagement/GridPP3_RiskRegister_4.xls

Key issues and changes for the risk register are:

- **R1: Recruitment and retention difficulties**. This remains high risk. Although the Tier-2 grants for GridPP4 have now been issued, reducing the risk of Tier-2 staff leaving, there is still a major concern over staff on fixed-term contracts at STFC. Nine SLA-funded GridPP posts are due to finish at the end of March; GridPP have requested that these be extended but there hasn't yet been confirmation of this.
- **R2: Sudden loss of key staff**. The issue with SLA staff also feeds through to risk 2, which has been increased to a likelihood of 4. If uncertainty about the extension of fixed-term contracts continues, it seems very likely that we will lose key staff at STFC, particularly in the Tier-1. In addition, the Project Manager will leave in January – plans for her replacement are currently underway.
- **R5: Service insufficiently resilient with respect to storage**. This has been raised to a likelihood of 3, given the issues we have seen with storage as the load has increased. However, the impact has been reduced to 3, as the recently introduced 6-month disk buffer gives more flexibility in this area if there are problems with a particular generation of storage.
- **R11: Changes to the LHC schedule**. This risk replaces the previous R11 (Further delays to the LHC). It has been set at likelihood and impact of 3 each, putting it in the high-risk category. Discussions are currently underway as to the possibility of LHC running in 2012, which will have a major affect on experiment requirements for this period.
- **R15: Lack of funding for network provision costs.** The HEFCE review of JISC is underway with the consultation period closing early November. It is not yet clear what impact this will have on the cost of JISC's services to the research councils, but it is possible that it will lead to charging of higher users. This is not yet imminent, so likelihood remains at 2 and impact at 4, but is flagged as possible future high risk in the longer-term future.
- **R27: Uncertain long term funding**. This risk has decreased to 3x3 now that some interim funding for GridPP4 has been agreed. In addition, the CSR outcome was relatively positive but some uncertainty remains around STFC posts and capital expenditure.

LCG Status

LHC commissioning proceeded this year at an unprecedented pace. Experiments showed their readiness in the exploitation of the 7 TeV data and were ready to follow the increase of luminosity with more complex triggers. Analysis proceeded very rapidly in all experiments with results being submitted for publication within days. The 'brilliant performance' of the WLCG was explicitly quoted by Sergio Bertolucci, CERN's Director of Research as a key factor in 'the spectacular start-up'. The 2010 target integrated luminosity for proton-proton collisions was exceeded on 13th October and the proton run ended on 4th November. Heavy Ion running started on the 7th November and is continuing with ATLAS, ALICE and CMS recording large amounts of data. The performance and resolution of the detectors has been impressive. The experiments have already exercised much of the Standard Model and are starting to explore beyond it.

During the first 6 months of LHC data at a collision energy of 7 TeV the WLCG service has shown itself to be able to manage, process, and analyse the data as planned, thanks in large part to the significant testing activities and data challenges of the past several years. The net result of this success is the ability of physicists to access physics data within hours of if being recorded, and the early publication of physics results. The key points to note from this early experience are:

- The resources are being used in accordance with the planned computing models, with slight variations in the light of experience. Globally, well in excess of one million jobs per day are being supported, and the equivalent of 100,000 CPU-days/day being delivered on average; around 13% of this is in the UK. These numbers are significant as they are at the scale anticipated for the first years of data taking in the experiment technical design reports.
- The networks are also being used at the scales anticipated; data have been available at Tier-2s for analysis within hours (far quicker than anticipated). The reliability of the OPN has been excellent, partly due to the reliability of the underlying networks, but also due to the redundancy now built in: With the deployment of the backup UK link, all Tier-1 sites now have at least 2 routes to CERN with a high bandwidth.
- Tier-2 sites are used heavily, and provide the majority of the analysis power. The frequently expressed concern that too much analysis would be done at CERN has not materialised; the number of people doing analysis is significant with ~200 for each of LHCb and ALICE, and 500 or more for CMS and ATLAS.
- Less CPU power was required at Tier-0 and Tier-1 sites than anticipated for a nominal year of data taking. This will change as the luminosity continues to increase but, in the interim, this has allowed for multiple reprocessing of the data as the calibrations have improved, with direct benefit for the physics output.

General Issues that have arisen include a lack of resources for ALICE (both pledged and actual), though opportunistic use has allowed ALICE to use significantly more than the pledge in the UK, and higher than expected heavy ion data rates in the Tier-0 that has resulted in heavy ion data being shipped directly to the Tier-1s. Data models that expect complete datasets of thousands of files to be at a site have difficulty coping with inevitable file loss and corruption. Investigations are under way to evaluate alternative caching models where missing files are brought in dynamically to complete the dataset. Existing and former GridPP staff are involved in several of these efforts. Globally, there has also been larger than expected numbers of breaks of service at the Tier-1s. These are for a variety of reasons and so there is no obvious way to reduce them. However, there are opportunities to reduce their duration and their impact on the experiments and a programme of sharing best practice among Tier-1s is under way.

In summary, the computing Grid has stood up well to the challenges of real data, which is just as well given the excellent performance of the LHC and the detectors.

EGEE/EGI Status

Shortly before the last OC meeting the EGEE project came to and end and the transition to the EGI European Grid Infrastructure started. EGEE ended on April 30th, with the operation of the EGEE infrastructure formally being taken over by the European "National Grid Initiatives" (NGI) that together make up EGI. For the next 4 years, the operation of EGI is co-funded by the European Commission via the EGI.InSPIRE project. In practice this means that the EC subsidies the central coordinating activities in Amsterdam and co-funds a number of activities in the NGIs themselves. The NGIs are typically dominated by the relevant National contributions to wLCG, especially in the Tier-1 countries. This alignment should ensure that the EGI infrastructure continues to deliver the services required by wLCG. In the UK, the NGI is an organizational umbrella covering GridPP and NGS operations. The UK commitments to the EGI.InSPIRE project are coordinated by John Gordon, who reports to an NGI operations management

board made up of representatives from GridPP and NGS.

The EGI coordinating body in Amsterdam, called EGI.eu, is supported by subscriptions from the NGIs, with JISC committed to meeting the UK subscription. The creation of the foundation was achieved later than originally hoped for and this created some difficulties during the transitions at the end of EGEE. However, the operation of the infrastructure has continued largely without incident and has contributed directly to the success of the LHC physics run. This smooth transition was in large part due to the dominant role of the wLCG sites in the operations. EGI.eu is now almost fully up and running with a staff of almost 20 people in Amsterdam. The foundation is gradually taking over the coordination of all EGI activities including the future evolution of the Grid services and installed software. EGI.eu formally reports to a Council made up of representatives from each of the funding NGIs, however, an Executive Board appointed by the Council oversees day-to-day operations. Neil Geddes is the UK Council member and has been appointed to the EGI Executive Board until February 2012. David Britton was recently appointed as the UK member of the EGI.InSPIRE Project Management Board representing the UK, Ireland and the Netherlands.

Current issues with respect to EGI and EGI.eu are:

- Formalizing the engagement with Grid user communities, including wLCG, following the hiatus in moving form EGEE to EGI. This is gradually being addressed as various committees and formal processes are put in place but needs to be monitored carefully over the coming year. GridPP and wLCG are represented directly and indirectly at many levels in the EGI structures and, hence, have good visibility of developments.
- Formalizing the mechanisms for defining and supporting the developments required to evolve the software infrastructure for EGI. EGI itself will not support extensive software developments and must collaborate with other groups such as the European Middleware Initiative (EMI) who are funded to do this. The EC has expressed its commitment to supporting the required software developments, but the mechanisms for doing this are not yet tested.
- Ensuring that all NGIs continue to support EGI.eu at the previously agreed level, in light of changing economic circumstances. Essentially all NGIs have paid the full fees for 2010. However, in September the EGI Council did not approve the increase in the fees required under the EGI.InSPIRE agreements (a majority supported the proposed increases, but the required 2/3 majority was not achieved). A plan to address the potential shortfall was drafted by the Council Executive Board but failed to pass a vote in the September council meeting. Discussions are now underway with EGI.InSPIRE as to how to operate within the available budget.

Tier-1 Status

This report covers the period from June 2010 to November 2010. The Tier-1 operations continued to be very smooth and reliable up until late August. However as data rates continued to increase a number of load related problems have required more urgent attention. Most issues have now been resolved but storage hotspots and disk server exceptions remain an operational concern.

During this period, the priority areas of work have been:

- Service operation during LHC data taking
- Final commissioning of 2009 disk hardware, and tenders for 2010 equipment.
- Preparation for, and implementation of, upgrade to CASTOR from 2.1.7 to 2.1.9 (only ATLAS now remains to be done and is scheduled for early December).

Fabric and Infrastructure

The (delayed) acceptance of the second 2009 disk delivery was successfully completed. This delay did not affect the ability to meet MoU commitments due to our policy of retaining a buffer of disk against such problems. Delivery of the 2010 disk and CPU procurements has commenced. Target date for commissioning of new equipment is March 2011.

The 10Gb OPN backup link to CERN was commissioned and the RAL site moved to a 20Gb link to SuperJanet. Work continued to roll out the Quattor Fabric management system with an increasing number of systems now entirely managed by Quattor, freeing up some effort. Solutions to provide a virtualisation environment for our service systems are being evaluated.

Operation of one of the 2008 generations of disk servers (about 1PB) has recently become very problematic (unknown cause at the moment). This generation will be withdrawn from service for assessment/re-certification. The shortfall in capacity will be temporarily made up from the buffer held against procurement problems, yet again demonstrating that this is essential.

The R89 machine room has provided a stable operating environment during this period.

Production, Operations and Service

During the months of June, July and (most of) August little pressure was placed on Tier-1 infrastructure despite LHC data taking. The Tier-1 continued to improve its overall performance metrics and operations continued to be smooth and generally trouble free. However since August there has been a considerable increase in workload by LHCb and more recently by ATLAS as demonstrated in Figure-4.

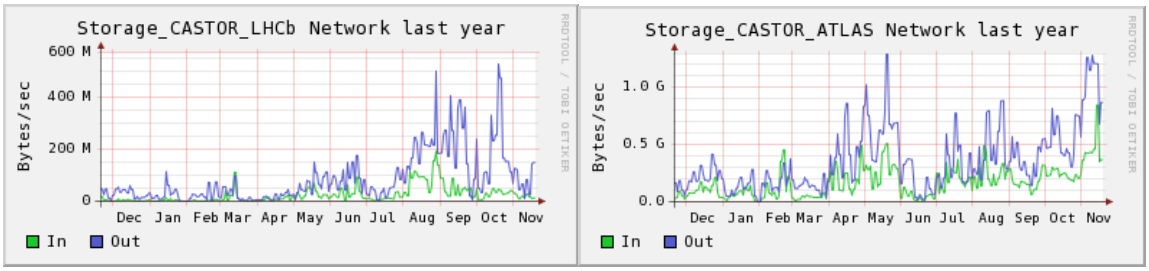

Figure-4: *LHCb (left) and ATLAS (right) workload over the past year.*

The Tier-1 availability as seen by the ops VO was 97% in the second quarter of 2010 and 99% in the third quarter, compared to the WLCG target of 97%. Availability to the individual VOs, compared to the average of all Tier-1 sites, was:

	June	July	Aug	Sep	Oct
ALICE	100% (98%)	99% (96%)	100% (98%)	93% (91%)	97% (97%)
ATLAS	94% (95%)	98% (92%)	98% (91%)	99% (95%)	99% (95%)
CMS	100% (98%)	99% (97%)	100% (96%)	99% (93%)	98% (97%)
LHCb	97% (90%)	94% (85%)	79% (84%)	77% (71%)	74% (91%)

Table-2: RAL Tier-1 availability by experiment compared to global Tier-1 average.

Generally the Tier-1 delivered above average availability to the VOs, usually exceeding its target of 97% however:

- LHCb operations were difficult, partly owing to issues related to the unexpectedly high loads and partly due to a few issues following the upgrade of CASTOR to 2.1.9 (LHCb were the first VO to upgrade). Work on changes to accommodate the higher than expected load is still ongoing.
- Routine operational disk server exceptions (where a server has to be taken out of production for a short time for hardware or O/S intervention) have had a high impact on VO operations (particularly ATLAS). One generation of disk will be temporarily removed from service until the fault is resolved.
- There were three data loss incidents in the period caused by this generation of disk. All followed a single drive failure from which the RAID hardware should have allowed recovery without data loss. Investigations are underway to understand the underlying cause.
- Callout rates have increased, but remain within the team's capacity to handle.

The long planned upgrade of CASTOR to release 2.1.9 (from 2.1.7) commenced after extensive testing. So far, three of the four CASTOR instances (LHCb, ALICE and CMS) have been upgraded. Although there were a few teething problems with the first (LHCb), generally these have gone well, completing ahead of schedule with very few problems.

Management, Business Processes and Communications

- The Tier-1 procurements are on track but disk and CPU prices were higher than expected. A number of factors contributed to this, primarily: exchange rates, global component shortages, low purchase volume and an increase in the required specification.
- Recent guidance requires BIS approval for ICT projects above £100K, and Cabinet Office approval for ICT projects over £1M. It is now believed that this does not apply to scientific computing but formal clarification has not yet been received. Such constraints may have delayed (or prevented) future Tier-1 hardware purchases and thus have had a negative impact on our ability to meet commitments.
- The Tier-1 is on track to achieve its FY10 planned expenditure provided new T10KC tape drives become available to the schedule we expect and the ICT spend is allowed.

  Tier-1 staffing remains at full compliment although a number of fixed term appointments are due to terminate in March 2011. A request has been made to STFC to extend all these posts to March 2012, however this has not yet been approved.
- A new resource management process has been put in place to better manage our tracking and allocation of disk, CPU and tape capacity. This works closely with the User Board to ensure resources are properly managed and allocated to meet experiment requirements and MoU commitments.

Deployment Status

Overview

Over the last 6 months the Grid has run well in the UK and globally. There have been several OS kernel vulnerabilities that have required a coordinated response. Rollout and testing of CREAM CEs has continued but the LCG-CE remains the recommended production version. Little impact has been observed from the continued transition of services and coordination to EGI. Within the UK NGI work has continued on developing joint and integrated support and coordination processes.

Deployment & Operations

Over this reporting period the retirement plan for gLite 3.1 has become much better defined; some services have been upgraded already and the rest will be go during 2011. GridPP sites have continued to rollout CREAM CEs in parallel with LCG-CEs to allow the experiments to test and smoothly transition. ATLAS is now using CREAM as its primary CE where it is available. The 35% of GridPP sites not yet supporting CREAM have plans to deploy in the coming months – these sites are typically the smaller ones with less manpower or where integration is more complicated (for example due to shared clusters). In parallel with this work most Tier-2 sites have been commissioning new hardware purchased with the 2nd tranche of GridPP3 funding. There have been a couple of new ARGUS/glexec instances deployed but little further testing by the LHC experiments and it is speculated that the WLCG MB will extend the use of Multi-User pilot jobs beyond the end of the year.

Security updates have been a hot topic recently with a succession of RHEL5 OS level vulnerabilities that potentially allow ordinary users to escalate their privileges. The incidents have demonstrated a security coordination process that works well and has shown that sites react promptly. In response to a request from the sites to have a statement about balancing availability of resources with risk to the site, the GridPP PMB issued a statement (item-5 in http://www.gridpp.ac.uk/pmb/minutes/101108.txt) on expectations in this area. In addition to real incidents, a new round of security challenges has begun with results for the WLCG Tier-1s showing a marked improvement compared to previous exercises – sites are assessed on communication, containment and forensics. The Tier-2 sites will be tested in the coming months.

Overall feedback from users and infrastructure monitoring has been positive with sites meeting the WLCG availability/reliability targets with a few exceptions: In August all UK sites were initially recorded with results below 95% but this was later understood to be due to problems with the monitoring itself and was subsequently revised upwards. There have also been occasional problems with the WMS services upon which the WLCG monitoring relies prompting a review of the Nagios submission framework and a closer look at the UK Nagios testing setup. The Liverpool Tier-2 has been badly affected by an omission in the Nagios monitoring, which failed to recognise the OS installed – though addressing this in the software is a relatively simple change, having that fix available in a production release has taken a long time. From October the WLCG Management Board started reviewing Tier-2 installed capacity versus pledged resources using the gstat tool; consequently the GridPP deployment team have been proactive in checking and fixing site published data – the wLCG MB did not raise any pledge concerns against the UK.

The LHCb upload issue affecting transfers from worker nodes (reported in the last update) has been resolved for all but one site. Apart from occasional issues with the experiment software areas LHCb and CMS report good satisfaction with the infrastructure. ATLAS has also had a good experience though there have been more issues to resolve including installation of Squids at Tier-2 sites, the setting of a minimum available disk (50TB) in spacetokens for sites to be considered for analysis work and the continued clearing up of dark data. Whilst there have been few serious issues, it has recently been observed that some sites are unable to cope with network traffic with their current WAN connections and traffic shaping is being implemented as a work around. Manual adjustments of the FTS channels to Tier-2 sites have also been required to alleviate occasional backlogs at some of the larger sites.

As sites have deployed more disk servers and as experiment usage of Tier-2 disk has increased, disk outages and data loss at Tier-2s has become more noticeable. In response the GridPP deployment team are working with the experiment representatives to ensure the recovery/deletion process for each VO case is well documented.

WLCG wide there have been advances in the approach taken to tracking and resolving network problems. An Information Officer is now in position, who acts as a single point of contact in resolving problems, improving and seeking input on LHC experiment and site needs. This individual will also advise on strategy for top-level BDII deployment. EGI held a technical forum in Amsterdam in September where plans for the infrastructure and EMI were outlined. The main change in this period has been the

introduction of a new interface to GOCDB4. There were some bugs to be overcome and sites continue to feedback on the changes.

Tier-2 MoU Commitments

Delivery of Tier-2 resources against the pledges due in July 2010 is shown in Table-3. CPU pledges are generally met. Edinburgh actually far exceeds its commitment but the GridPP formal site allocation reflects the minimum guaranteed fairshare (of the shared cluster) rather than the actual delivery. All theTier-2 as a whole meet the WLCG disk commitments, though not all the individual site pledges to GridPP are yet delivered. In some cases this reflects delays in funding from GridPP or intentional delays as the existing disk was not being fully used. To understand the ongoing situation GridPP has looked at the projected installed capacity for 1st April 2011 (based on purchasing and commissioning that is currently taking place) and this suggests no cause for concern. QMUL for example will be at 118% of its pledge, Lancaster 200% and Glasgow is already at 125%. Durham will not be deploying more disk as their current deployment meets the more specialised VO usage needs for that site and the GridPP MOU will be adjusted to reflect this.

	CPU		Disk	
	HepSpec06	2010 MoU	TB	2010 MoU
Brunel	4417	151%	450	306%
Imperial	8676	93%	605	139%
QMUL	9791	180%	319	42%
RHUL	3160	79%	286	169%
UCL	2604	585%	50	455%
Lancaster	6342	217%	381	93%
Liverpool	8318	326%	286	65%
Manchester	18482	295%	207	116%
Sheffield	3908	379%	213	107%
Durham	5700	184%	30	19%
Edinburgh	1118	42%	100	164%
Glasgow	20264	204%	474	44%
Birmingham	3344	231%	195	109%
Bristol	1836	278%	110	500%
Cambridge	2260	197%	101	94%
Oxford	3564	175%	199	98%
RALPPD	12624	194%	698	192%
LondonGrid	28648	129%	1710	112%
NorthGrid	37050	290%	1087	89%
ScotGrid	27081	172%	604	47%
SouthGrid	23628	200%	1303	149%
Total Tier-2	116407	187%	4704	96%

Table-3: Disk and CPU resources at GridPP Tier-2 sites compared to the pledges due in July 2010.

Users' Reports

ATLAS

ATLAS has continued to make heavy and effective use of the UK Grid resources. While the job success rate at the RAL Tier-1 is slightly below the desired target, the causes are now more evenly split between many small causes, and the previous problems associated with accessing the storage are much reduced. ATLAS has by now developed a degree of flexibility to deal with such failures, which now merely decrease the job efficiency a little. Perhaps more troubling is the rate of disk server failures at the Tier-1. Although ATLAS requested some investigation of this, the Tier-1 had independently began a study. After initially concluding that there was no cause for concern beyond statistical fluctuation, it is now becoming clearer that there may be some hardware failure issues to do with a specific generation of disk purchase. ATLAS is investigating the best use of the 'suspect' storage, allowing it to remain useful without threatening the operational efficiency. ATLAS commends the response of the Tier-1 staff when file losses have occurred; their efforts to recover such files have been diligent, and ATLAS has actually requested less of an effort in most cases, as it is easier to remake many files by other means.

The Tier-2s have generally performed well, both for production and for analysis. There has been a sustained effort within ATLAS to develop metrics that can track the site performance in these areas using minimally intrusive tests and passive measures. Inevitably, much needs to be learned about these metrics, as they are mixed with external factors; we no longer have the luxury of dedicated tests run on relatively idle sites, as was the case a year ago. As such, the subset of important and reliable metrics still remains to be established, and will continue to evolve, but the information is available for sites through the Ganga Robot webpages. We note that some sites continue to have problems with their storage systems, and continued effort in this area will be needed as datasets and file systems grow.

ATLAS is trialling a new means of data distribution based on the popularity of datasets for the Tier-2s. It is hoped that this will increase the utilisation of the Tier-2s. Initial studies have been promising, but have revealed more work is needed on the brokering.

CMS

The period since the last Oversight Committee has been extremely exciting for CMS. During this period the UK has played a full role in CMS computing. Currently all physics data (with the exception of Heavy Ion running data) is copied to all Tier-1 sites, however this is unsustainable in the medium to long term and custodial copies of primary datasets are allocated to specific Tier-1 centres. The UK Tier-1 was the primary custodial site for the "e-gamma" dataset in run 2010A and the "jets" dataset in run 2010B (the "e-gamma" dataset ceasing to exist as an individual dataset once the trigger thresholds were increased for run 2010B).

The Tier-2 centres have continued to support 5 analysis groups: Electroweak (Brunel), Exotica (RALPP), E-gamma (Imperial), SUSY (Imperial) and Trigger Studies (Imperial & Brunel) and these have provided heavy usage of the Tier-2 sites throughout this period. The other role of the Tier-2 sites in CMS is Monte Carlo event generation and here the UK has augmented the capacity of the CMS centric Tier-2 sites by using spare capacity in other sites and treating them as "opportunistic Tier-2 sites". These are labelled as (one of the several different sorts of) Tier-3 sites in CMS terms so that they are not subjected to the same continuous testing as regular Tier-2 sites. The great benefit that these sites bring to CMS can be seen in the plots below. Figure-5 is the cumulative number of events analysed at UK sites since the last Oversight Committee meeting and is dominated by Imperial, RALPP and Brunel. Figure-6 is the total "good" CPU used at UK sites over the same period and includes all activities (including the Tier-1 activities and the Monte Carlo production at the opportunistic Tier-2 sites) and clearly shows the importance of the contribution made by other sites. It is important to note that this is only possible because Monte Carlo production is much less arduous than data analysis.

In general CMS is very happy with the UK Computing performance. There have been some minor issues including: ALICE jobs blocking the queues meaning that reprocessing of custodial data had to be moved from RAL to FNAL on the eve of an international conference; two minor data losses at the Tier-1; the efficiency of analysis jobs at Brunel; and there is still room to improve efficiency at the other UK sites. However the Tier-1 and Tier-2 sites have been able to deal with issues as they arise and overall operations have been very smooth for CMS during this period.

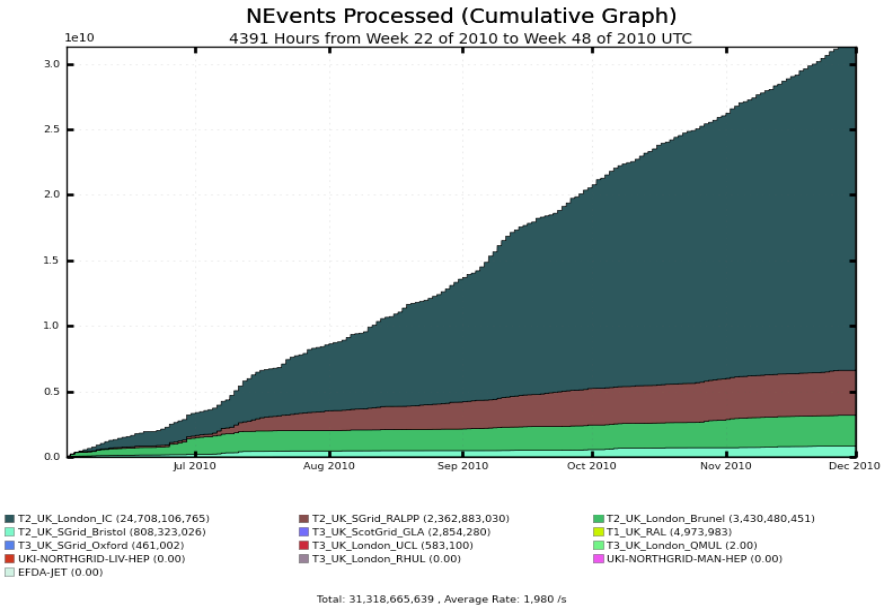

Figure-5: CMS Physics events analysed in the UK since the last oversight committee

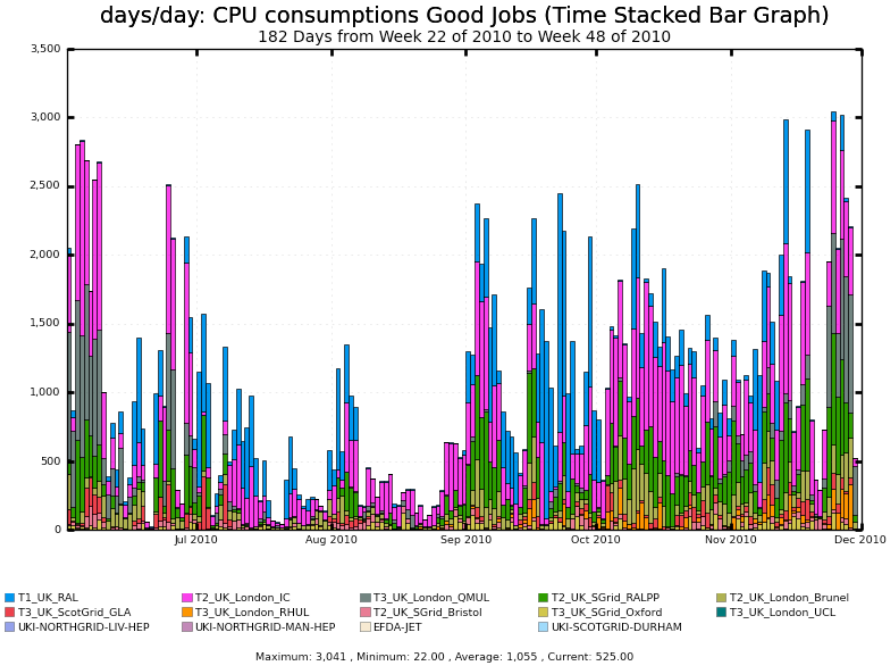

Figure-6: UK CPU hours used since the last Oversight Committee, showing the Tier-1 (blue), Tier-2s (pink, green and olive) and valuable contributions from Tier-3 (opportunistic Tier-2) sites.

LHCb

LHCb recorded an integrated luminosity of ~38 pb^{-1} during the proton-proton running at a centre-of-mass energy of 7 TeV. The LHCb detector performed well with all sub-detectors achieving high (>98%) efficiencies and with the major UK components of the VELO and RICH playing a major role in vertex location and particle identification.

The data taken so far has been successfully processed at CERN and the six external Tier-1 sites of LHCb. Another reprocessing of the data started during the last week in November in preparation for the winter conferences and using the latest calibrations and improvements to algorithms. The running conditions of the LHC posed some challenges to the data processing due to the fact that the number of interactions/crossing was up to six times larger than the design value for the experiment. In preparation for recording 1-2 fb^{-1} in 2011, some changes to the LHCb computing model are anticipated with the possibility that reconstruction will also be performed at Tier-2 sites in addition to the Tier-1.

The period since the last Oversight Committee meeting has revealed a few issues with LHCb processing in the UK. There has been a particular problem with disk-server reliability at the RAL Tier-1 and LHCb continues to have concerns. Unexpectedly high network needs in the LHCb production workflows and various internal RAL configuration issues, also caused significant job failures during August - October. The UK Tier-2 sites have spent a significant amount of time understanding and solving a recurring problem with uploading LHCb data to Grid storage.

The upgrade of the LHCb Castor instance to version 2.1.9 at the Tier-1 in September went smoothly once some re-configurations were made. After the disk servers were upgraded to 64 bit SL5 and the SRM front end machines were upgraded in mid November, the RAL Tier-1 storage has seen increased stability for LHCb. Since July, the UK has processed 18.3% (~614k jobs) of the total number of batch jobs submitted, continuing to be well in line with the UK authorship fraction of LHCb. Ganga is used to submit jobs to the Grid through DIRAC and the job rates have an interesting “spiked” nature as can be seen in the plot below.

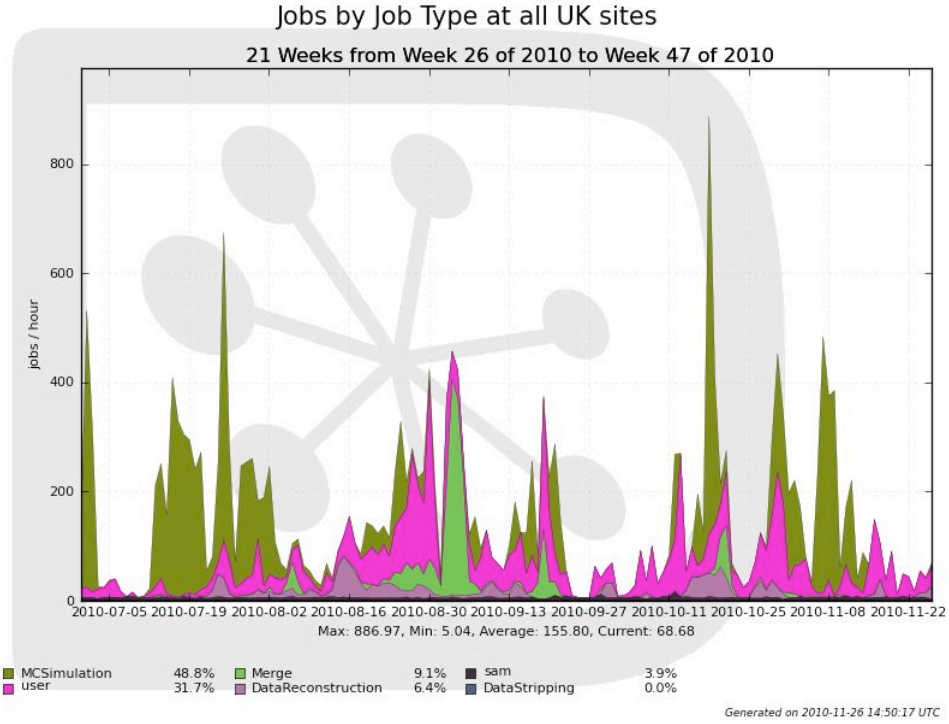

Figure-7: LHCb batch jobs at all UK sites.

Other Experiments

As with the other LHC experiments, ALICE has been very active in exploiting the RAL Tier-1 centre with high CPU utilisation through July and August. A ramp-up occurred again in October in the lead up to the heavy ion run, which started on 7 November and is currently underway.

SuperB has been active in preparation for the Technical Design Report in 2011 and has made good opportunistic use of the Tier-1 centre. Similarly, both H1 and MINOS were active over the summer months for Monte-Carlo production.

The MICE experiment concluded Step I of their programme at the ISIS facility in August. In the latest computing model of the experiment, the RAL Tier-1 centre is used to store custodial copies of the data with the RAL PPD Tier-2 site used for express calibration jobs. The University sites are used for Monte-Carlo production, skimming and user analysis. GridPP effort has provided support for the data mover system, which automates data upload to the Grid.

ILC has consumed only modest amounts of CPU at the Tier-1 in the last few months, but Monte Carlo production for the Compact Linear Collider (CLIC) Study will be starting soon and there will then be an analysis phase of these data from April 2011.

Support has been provided for the maintenance and upgrading of the LFC for SuperNemo and a new host machine has recently been deployed.

Impact

Outreach

The autumn is always very busy for GridPP dissemination, and this year was no different. The event season began with the annual UK e-Science All Hands Meeting in Cardiff and the EGI Technical Forum in Amsterdam, in the week 13-17 September. At both these conferences GridPP shared a stand with the NGS to promote the newly established UK NGI, with a screen displaying the RTM, a pop-up GridPP poster and handouts. The events were considered an excellent opportunity to showcase the work being done in the UK both technically and politically through the formation of the UK NGI.

In late October GridPP contributed a stand, posters, plenary and session talks to the CHEP conference in Taipei. This is one of the most important scientific meetings for the GridPP community and the UK has always had a highly visible presence. The GridPP dissemination team also helped to organize and blog from the Citizen Cyberscience Summit at Kings College in September. This brought together a high profile list of speakers, scientists and volunteers using the various distributed volunteer computing networks such as SETI@home and LHC@home, and drafted a “citizen cyberscience manifesto” on how to take the field forward.

Individual members of the collaboration regularly give outreach talks. Examples within this period included a seminar to the KCL undergraduate physics society by Neasan O'Neill, the GridPP Dissemination Officer entitled "From Web to Grid: The Future of Scientific Computing?" and a talk by Steve Lloyd on the LHC and Grid computing as part of the QMUL Alumni Reunion Weekend.

The GridPP website continues to be actively updated, with 12 news items since June and a regular stream of blog postings aggregated by planet GridPP, including posts from, among others, the Tier-1, ScotGrid, NorthGrid, SouthGrid the storage group and NGS since the last OC. The GridPP news section has been revised and the back-end administration updated, allowing users to view news by year as well

as the latest news, and is an interesting archive of the project's history containing 280 news stories since GridPP began.

GridPP also began to work closely with NGS, NeSC and the Software Sustainability Institute to help form a coherent and sensible dissemination strategy for the UK NGI. This has resulted in a draft logo and a website is currently being designed. The EC-funded GridTalk project, which was led by the QMUL outreach team, was successfully concluded with an excellent final review in Brussels. Its successor, e-ScienceTalk, began on 1 September and funds Grid outreach staff in the UK at QMUL and IC.

Knowledge Exchange and Collaboration

GridPP held a well-received and interesting session on knowledge exchange at GridPP25. Speakers were: Becky Parker, a teacher at the Langton Star Centre, who spoke about her work with the CERN@school project; Stuart Purdie from ScotGrid, on making your Tier-2 work for non-HEP users; and Alex Efimov, a QMUL technology transfer officer who has previously consulted for STFC and who spoke about opportunities for knowledge exchange with industry.

GridPP is now working further with Becky Parker, who wishes to use Grid resources as part of a multi-school project to gather and analyse cosmic ray data. CERN@school has set up a VO on the GridPP infrastructure and is currently bidding for funding from various sources with input from GridPP. Also as a result of the discussions at GridPP25, talks have begun with the technology transfer team at QMUL to look into possible industrial collaborations.

A new collaboration was recently started with the Digital Curation Centre in Edinburgh. The DCC is interested in how GridPP members manage data and information to support research, what improvements are called for, and what lessons our experience might offer other projects facing similar issues. Their project will include a survey of GridPP staff, interviews with selected participants and diaries of data use.