

GridPP
UK Computing for Particle Physics

GridPP Project Management Board

Project Status

Document identifier:	GridPP-PMB-174-ProjectStatus.doc
Date:	28/04/2016
Version:	1.0
Document status:	Final
Author	PMB

Introduction

The Oversight Committee last reviewed the GridPP project in March 2015 at the end of GridPP4 and the start of GridPP4+. That one-year extension is now complete and this report summarises the project status as the four-year GridPP5 project is starting in April 2016.

The LHC Run-2 started in June 2015 and reached record proton-proton collision energies of 13 TeV with an intensity increase provided by a beam structure containing 2240 proton bunches. By the year-end-technical-stop (YETS) in December 2015, 4 fb^{-1} of data had been accumulated, providing tantalising hints of new physics. The machine is now being re-commissioned to continue Run 2 with a goal of increasing the intensity further using 2748 bunches and accumulating a total Run-2 integrated luminosity of 25 fb^{-1} by the end of the year.

The UK Grid infrastructure has continued to provide an excellent service as part of the Worldwide LHC computing Grid (WLCG) and the resources have been heavily used, though with slightly more headroom when the LHC was not running.

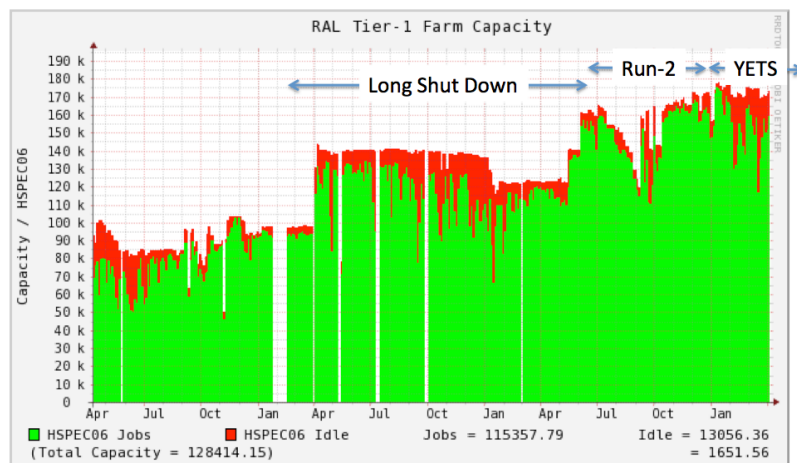


Figure-1: CPU usage at RAL from April 2013-2016. Green (Red) areas show used (idle) capacity.

Over the past year, the 34th and 35th GridPP collaboration meetings were held with themes on “Run-2”, at QMUL, and “Embracing the Future”, in Liverpool. The review of the GridPP5 proposal was completed and funding at a level of approximately 90% of flat cash has been awarded for the period April 2016 to March 2020.

Wider Context

WLCG resources have continued to be delivered according to the annual cycle wherein resources pledged at the end of August are deployed by the following April. At the beginning of March 2016, WLCG reported to the LHCC Referees that no particular problems were anticipated this year. The UK procurement/deployment have allowed the UK pledges to be delivered on time.

GridPP currently provides about 11% of the total global resources. Figure-2 below shows that the Tier-1 at RAL delivered 10% of the global LHC Tier-1 resources. In fact, this is rounded up in the figure from a numerical value of 9.6%, which is entirely consistent with the 9% reported in previous periods. Figure-3 shows 12% of the LHC Tier-2 resources (down from 14% last time) were delivered by the UK over the same period. This change is predominantly due to substantial increases in the German Tier-2 contribution, from 11% to 17%. The difference between the UK Tier-1 and Tier-2 global fractions reflects the significant leverage of Tier-2 resources from the institutes that GridPP funding has enabled.

Developed by CESSGA 'EGI View': / sum_normcpu / 2015-4-2016:3 / TIER1-VO / lhc (x) / GRBAR-LIN / 1

2016-04-06 02:06




Figure-2: CPU delivered to LHC VOs (April 2015 to March 2016) by Tier-1.

Developed by CESSGA 'EGI View': / sum_normcpu / 2015-4-2016:3 / COUNTRY_T2-VO / lhc (x) / GRBAR-LIN / 1

2016-04-06 02:06




Figure-3: CPU delivered to LHC VOs (April 2015 to March 2016) by Tier-2s.

During the last 12 months GridPP has been active on several scales in the efforts to consolidate e-Infrastructures and promote more shared use across communities.

At the EU scale members of GridPP have continued their support of the activities of EU-T0 (an association of European funding agencies) principally in promoting policies for a shared infrastructure (across particle physics, astronomy, cosmology and astro-particle physics), and in positioning for participation in the immediate H2020 call for the European Open Science Cloud (EOSC). The EOSC is set to be a major theme over the next decade, and it is therefore important to be involved to be best placed to reap any eventual benefits for all of STFC science.

Within the UK there is (independently) a strong movement to harmonise the national e-Infrastructure across research councils. This is in part driven by the move to an R-UK following the Nurse review. We have participated strongly (including leading) in many of the National e-Infrastructure (Nel) meetings centred upon the RCUK e-Infrastructure group and the “Project Directors Group”. As part of this work we have been central in helping to prepare the estimates of the capital and staff resource investments required to underpin STFC computing in the next 5 years that were submitted to BIS in January 2016.

Within the STFC domain members of GridPP and STFC SCD have been centrally involved in promoting the coordination of computing across PPAN science areas and the national facilities. A face-to-face meeting was held in November 2015 and three liaison phone meetings have been held since. Through this work STFC communities have begun to discuss how to work together in the future. In order to establish the spirit of this initiative GridPP has made efforts to provide some small, but important, resources to several astronomy communities. We have enabled LSST¹ to use 150,000 core hours for galaxy shear analysis and have a joint GridPP-LSST post (Edinburgh). We are currently enabling Euclid to use resources. Lux-Zeplin has now been enabled to use GridPP for routine work and we have a joint GridPP-LZ post (Sheffield). We have contributed to the SKA project through a joint GridPP-SKA post (Cambridge) and in supporting preparation of the (now submitted) AENEAS H2020 proposal to design a European Science Data Centre. STFC-SCD will participate in this directly, and GridPP in general will remain engaged to help SKA developments where this is helpful.

The UK continues to receive support from EGI.eu and EGI Engage to provide services that are critical to WLCG (APEL accounting; Grid Operations Centre Database (GOCDB); and leadership of international security operations and policy development). The UK currently chairs the EGI Council (Dovey, JISC) and Executive Board.

GridPP is also engaged with and contributing to other H2020 initiatives. These include EUDAT, the Indigo Datacloud project and the AARC Authentication and Authorisation for Research project in collaboration with Geant and others. GridPP is also working with the HEP Software Foundation, (a collaboration of HEP software projects designed to promote standards, reduce overlaps and provide common utilities).

¹ The Large Synoptic Telescope Project

GridPP4+ Status

Since the last OC meeting the GridPP4 project was granted a one-year extension known as GridPP4+. Over this period, the production systems at both the Tier-1 and 2 sites have continued to function well and have contributed to the successful processing of Run-2 data. GridPP4+ hardware funding was spent at the Tier-1 and Tier-2 sites and the equipment procured is currently being commissioned.

The project progress has continued to be monitored by Quarterly reports from the various Work Packages and tracked against the GridPP ProjectMap, which contains a series of on-going metrics and a set of annual milestones. The ProjectMap was revised at the start of the GridPP4+ phase in order to remove the milestones delivered during the earlier GridPP4 phase. The evolution of milestones and metrics over the GridPP4 and GridPP4+ period is shown in Figure-4 and the current Project Map shown in Figure-5.

	Q112	Q212	Q312	Q412	Q113	Q213	Q313	Q413	Q114	Q214	Q314	Q414	Q115	Q215	Q315	Q415
Metric OK	148	157	161	158	169	171	169	159	168	171	166	164	164	164	164	169
Metric close to target	17	14	17	9	7	6	5	17	9	7	10	8	11	11	12	7
Metric not OK	18	12	5	15	5	5	8	6	5	4	2	6	3	2	1	1
Not able to be measured	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Suspended	21	22	18	19	19	18	18	19	19	19	23	22	22	19	19	19
Awaiting input	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Metrics Total	204	205	201	201	200	200	200	201	201	201	201	200	200	196	196	196
Milestone achieved	23	28	33	34	40	46	54	60	62	62	65	73	76	15	15	19
Milestone close to completion	1	1	0	0	5	3	0	1	2	4	3	2	0	0	2	0
Milestone overdue	4	2	2	1	1	0	0	0	0	1	1	0	0	0	1	1
Milestone not due/ metric n/a	50	46	46	46	36	33	28	20	17	14	12	7	6	18	15	13
Milestones total	78	77	81	81	82	82	82	81	81	81	81	82	82	33	33	33
Total	282	282	282	282	282	282	282	282	282	282	282	282	282	229	229	229




Figure-4: Evolution of Milestones and Metrics for GridPP4 & GridPP4+ to Q4 2015. The lower left plot shows the number of metrics that are currently met (green) and not met (red). The lower right plot shows the evolution of the number of milestones that have been met (red flags late ones; the drop in 2015 is where the completed GridPP4 milestones were removed from the count). The data is presented numerically in the table at the top.

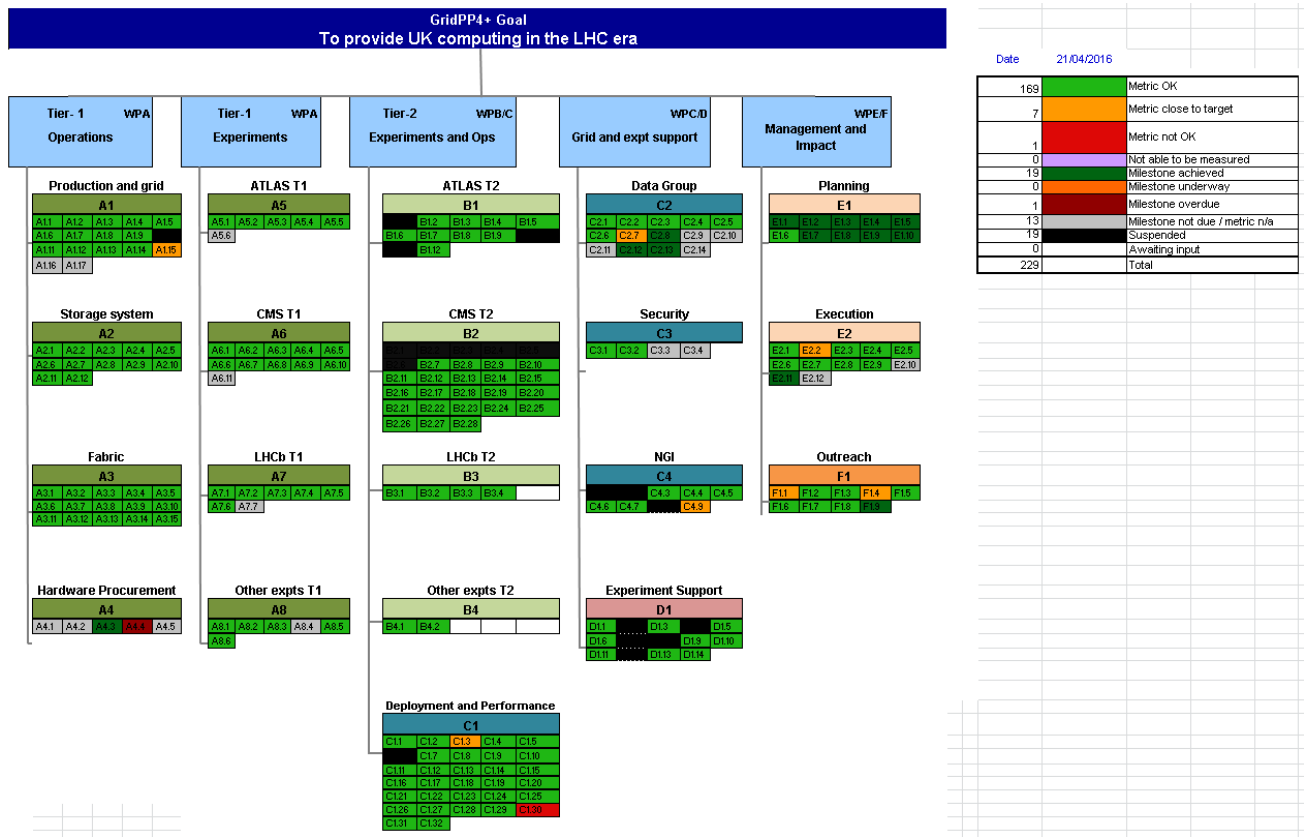


Figure-5: The GridPP Project Map at the end of Q4 2015

Overall the majority of metrics are being met or are close to target, with just seven metrics unsatisfied this quarter.

They are:

- A1.15 ‘Percentage of GridPP4 Staff in post’: The achievable maximum staff count at the Tier-1 was reduced to 18.5 during GridPP4+ (owing to an increase in overheads). However, as the Tier-1 staff-count needed to fall to 17.5 by April 2016 it was not feasible to hit that target.
- C1.3 ‘Fraction of HEPSP06 available utilised in the quarter’: This was down to 63% this quarter because of downtime at various sites due to power outages from flooding or air conditioning problems.
- C1.30 ‘NorthGrid average SAM availability’: This was low this quarter, again affected by the prolonged power outages due to the Lancaster floods.
- C4.9 ‘The number of sites marked red by EGI’: This was affected by UCL consistently failing due to the absence of GridPP funded effort at that site. GridPP has addressed this issue by reconfiguring the site to be a CPU only “Vac” site and the situation has improved.
- E2.2 ‘All quarterly reports received two months after the quarter end’: Some reports have been more difficult to produce due to key staff taking on new duties. Replacements are now in place so should not be a problem in Q116.
- F1.1 ‘Number of events attended by GridPP stand/posters’: This was low but has become a reduced priority for GridPP as “dissemination” has refocused on “impact”.
- F1.4 The number of news items on the GridPP web site is currently below target at 8 in the last year. However as before the focus of the Dissemination Officer has been to engage new VOs and a great deal of progress has been made in that area.

Five milestones have been delayed. These are:

- A1.16 & A1.17 'External review of Tier-1 operations' and 'Strategic operations plan for Tier-1': This will be rescheduled to 2016 now that the planning for GridPP5 is finalised. There seemed little point in performing these tasks until the uncertainty associated with GridPP5 funding had been resolved.
- A4.4 'Tier-1 Capacity Orders placed': Increased bureaucracy associated with procurement using SBS caused delays in tendering. Responses were received at the end of December, orders placed early in January, and delivery made before the end of the financial year.
- C3.2 The banning security challenge was delayed but was run in Q116.
- C3.3 "Security recommendations for the future". With the extension to GridPP4 by one year this metric will be delayed until GridPP5.

The quarterly reporting has continued to be useful to track issues and focus the different groups on their priorities. The overall performance of the different work packages in GridPP has been very successful and this has been reflected in the project map and reports. The quarterly report format and structure will be reviewed and overhauled for the start of GridPP5.

All detailed quarterly reports are available on the web at:

<http://www.gridpp.ac.uk/gridpp4-quarterly-reports/>

Risk register

The GridPP risk register was recently reviewed for the start of GridPP5 and the current version summary is presented below in Figure-6 and the full version is available at:

<http://www.gridpp.ac.uk/collaboration/docs/pmbdocs/>

Ref.	Risk Description	Owner	Inherent Risk			Residual Risk		
			Likelihood	Impact	Total	Likelihood	Impact	Total
1	Castor Storage System Problems	AS	0.4	60	24	0.3	60	18
2	Expt. s/w runs poorly on the grid	PG	0.5	60	30	0.5	50	25
3	Tier-2s are not fit for purpose	SL	0.4	60	24	0.1	60	6
4	Outage of UK T1	AS	0.4	60	24	0.1	50	5
5	Loss of experienced personnel at T2s	SL	0.7	50	35	0.6	45	27
6	Recruitment retention problems at RAL	DK	0.9	60	54	0.6	50	30
7	Failure of T1 to meet SLA or MoU	AS	0.6	50	30	0.2	50	10
8	Security problem affecting reputation	DK	0.7	80	56	0.6	50	30
9	Loss of GridPP service due to security	DK	0.3	50	15	0.1	50	5
10	Mismatch between budget and hardware costs	DB	0.25	50	12.5	0.25	20	5
11	Core service funding insufficient	DB	0.5	40	20	0.2	10	2
12	Breakdown of NGVEGI infrastructure	JC	0.3	30	9	0.3	30	9
13	Insufficient travel funds	DK	0.4	25	10	0.15	25	3.75
14	Expt. Computing models change	PG	0.3	40	12	0.1	40	4
15	Critical middleware no longer supported	DC	0.4	40	16	0.1	40	4
16	Significant loss of custodial data at the T1	AS	0.1	75	7.5	0.1	50	5
17	Unplanned infrastructure costs	PG	0.25	30	7.5	0.2	10	2
18	Loss or damage to hardware at T1	AS	0.05	100	5	0.05	100	5
19	Disaster at T1 leads to prolonged outage	AS	0.05	80	4	0.05	70	3.5
20	Insufficient funding at T2s for h/w	SL	0.2	40	8	0.2	30	6
21	Loss of EGI.eu	DB	0.2	20	4	0.2	15	3
22	Insufficient Network Bandwidth	PC	0.4	50	20	0.3	50	15
23	Technology Mismatch	DC	0.2	50	10	0.05	50	2.5
24	Insufficient VO/user support effort	PG	0.7	50	35	0.55	50	27.5
25	Financial Uncertainty	DB	0.5	60	30	0.4	60	24
26	Failure to deploy or operate hardware	JC	0.5	65	32.5	0.3	65	19.5
27	Conflicting opinions amongst GridPP stakeholders	DB	0.3	50	15	0.3	50	15
28	Over contention for resources	DB	0.3	40	12	0.1	40	4
29	Capital vs Resource at the Tier-1	DB	0.8	50	40	0.5	60	30
30	Tier-1 CEPH Project fails	AS	0.7	60	42	0.5	60	30
31	Failure of achieving further intergration within PPAN community	JC	0.5	60	30	0.45	60	27

Figure-6: The GridPP Risk Register Summary

The current risks that are somewhat elevated (orange in Figure-6 above) are discussed below:

Risk-5: Loss of experienced staff at Tier-2s.

The prolonged uncertainty of funding for GridPP5 caused some staff to seek alternative employment. Further losses could adversely affect the operation of the Tier-2s. New working models are being developed to help mitigate this by providing a lower maintenance mode of working at sites with little manpower. However this places more demands on the remaining larger sites.

Risk-6: Failure to retain or recruit key technical staff at RAL

Losing key staff could impact the projects ability to meet deliverables and key milestones. This is a potential problem that has to be dealt with rapidly by STFC management by speeding up recruitment in the event of a critical loss. It has been noted that recruiting new staff is increasingly difficult due to non-competitive salaries and conditions. Retention has been affected by the uncertainties in GridPP5 funding.

Risk-8: Security problem affecting reputation

There is a continuous need to be vigilant on all security issues. Security problems may bring down the whole of GridPP or even WLCG, given that all Sites run very similar software. A major outage is likely to attract significant outside interest and GridPP is likely to suffer damage to its reputation. Attacks of various types are becoming more common and the time required to patch and update can take up significant operations time. We will continue to act on many fronts to reduce the threat by ensuring software is up to date and to contain and handle incidents quickly when they occur. The UK (GridPP) leads the EGI (WLCG) Security Policy Group and the Software Vulnerability Group. The UK security team operates across GridPP sites to promote best practice and many staff regularly benefit from security training as part of HEPSYSMAN events. The security officer will coordinate all security issues and lead the distributed team.

Risk-24: Insufficient VO/user support effort

We have dedicated support posts for the major LHC VOs, but the increased number of new users and groups is stretching the support teams. A dedicated GridPP-support email list has been setup to help with this. The new user groups have each been assigned a 'Champion' who takes responsibility to ensure they are progressing well. The status is reported weekly at the GridPP Operations meeting.

Risk-29: Difficulty with budgets due to Capital vs Resource limitations

Classification of different types of equipment as capital or resource can sometimes change causing problems with budgeting. It is important to have good frequent communications with STFC to track such changes and limitations. FY16 levels have not yet been confirmed.

Risk-30: Tier-1 CEPH project fails

If the CEPH project (to replace CASTOR for disk storage) fails during scale testing, we would have to fall back to CASTOR, which we expect would not cope well with the increased load. Once CEPH is put into production, new problems may arise leading to an unacceptable degradation of performance or data loss. Progress is going well, and initial testing is under way but close project management by the Tier-1 team is in place to mitigate the risks escalating.

Risk-31: Failure of achieving further integration within PPAN community.

If GridPP fails to work with and engage others within the PPAN community GridPP would not benefit from potential shared infrastructure investment and cost savings. There could be reputational damage, and lack of endorsement that the GridPP infrastructure service is relevant to a wider community. Effort is being made to bring on new users but staff resources are being stretched.

The remaining risks in the table above are green, which indicates they are not regarded as significant dangers at this time. Nevertheless, we discuss the four highest of these below:

Risk-2: Experiment software runs poorly on the Grid.

In general we have very good communications links with the major LHC experiments. However ALICE require specialized storage access protocols not well supported by the future ECHO (CEPH-

based) storage project. Interim solutions will be provided but there is a risk that in the longer term the Tier-1 cannot support them without additional expenditure and effort. However, ALICE is a small part of GridPP and we are in active discussion with them.

Risk-7: Failure of Tier-1 to meet SLA or MoU commitments

Failure to meet our WLCG obligations would damage our reputation. Many systems have been put in place to cover call out, change management, disaster recovery and hardware procurement. These continually focus operations staff on the MoU commitments to ensure excellent performance. However recent staff losses, difficulties with financial approval for procurements, and the capital-resources issue, have made this difficult.

Risk-25: Financial Uncertainty

GridPP5 has now been funded so should provide some stability in the short term. There are still longer term uncertainties with funding following the recent CSR settlement, and due to plans for the Research Councils post Nurse-review.

Risk-26: Failure to deploy or operate hardware

This risk has increased due to the concern that the changes in BIS requirements and SBS procurement systems delay or complicate the next procurement round, which could result in late delivery of equipment. This would make meeting the WLCG pledge levels difficult.

Tier-1 Status

Overview

This section of the report covers the GridPP4+ period from February 2015 to March 2016.

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

- Ongoing routine stable operation of the production service supporting the first year of LHC Run 2.
- Procurements and deployment of hardware to meet the 2016 MoU commitments
- Continue to seek to broaden the funding base by engaging with new user communities and funding opportunities.
- Continuing development of new services (such as cloud and object store) in order to meet potential needs of new user communities

Fabric and Infrastructure

Tenders were carried out for disk and CPU capacity to meet the Tier-1's 2016 MoU commitments. There have been increasing constraints on the Tier-1's procurement process imposed by government policy and shared service centre (SBS) operational procedures and support capability. Several months were lost while a viable procurement process was established. Ordering and delivery was therefore later than planned (approximately 13PB of disk and 106KHS06 of CPU) but was finally completed in March 2016. Newly procured capacity is projected to be online by June 2016 – too late to meet the April MoU commitments. Nevertheless, both disk and CPU MoU commitments can still

be met by delaying the phase out of equipment after scheduled end of life, and delaying deployment of the full volume of the new CEPH storage system which is not yet required to deliver production capacity. The R89 machine room has provided a stable operating environment during this period.

The procurement process (through a new Framework agreement and tender management system) yielded a bid from only one disk supplier and it was fortunate that this bid was both acceptable and competitively priced. Usually the Tier-1 procures two independent disk and CPU solutions as a risk mitigation strategy against downstream batch related hardware problems – this year's single tranche of disk will increase downstream hardware operational risk, particularly (although not exclusively) in the last year of GridPP5 (FY19/20) when it comes off maintenance.

STFC has recognised the increasing challenge it faces in procurement through SBS this year and has put in place a new procurement specialist in house and established new relationship management processes between STFC and SBS in order to minimise issues caused by procurement. Nevertheless it is clear that the Tier-1 will immediately need to establish clear plans for its procurements. This is likely to be hampered by STFC's uncertainty regarding its full year budgets (although the much reduced scale of planned FY16 procurements may simplify matters).

BIS has recently clarified its IT procurement policy guidelines. Above a threshold of £100K all IT procurement must be approved by BIS and it is now clear that STFC will need to obtain BIS/cabinet office approval for the Tier-1 procurements in GridPP5. This can be a lengthy and uncertain process and will need to commence shortly.

Despite a general perception that overall hardware reliability is improving, large-scale batch related hardware problems still crop up from time to time. In January 2016 a review of recent operational experience of one of the two FY11 generations of disk servers indicated a hardware fault correlated with disk drive manufacturer. The servers began to experience correlated multi-drive failures under high load, problems beginning to occur after 3.5 years of operation. The same disk drives are also present in one of the 2012 disk generations (now just over 3 years old), which also appear to show early indications of similar problems. Several mitigation strategies are being pursued (firmware updates, drive replacement with other (also old) compatible drives that have shown better reliability, load reduction) in order to continue operation for some time until the equipment can be phased out. While one generation appears particularly unreliable we have also concluded that the second 2011 generation is sufficiently reliable that its operational life may be extended by a further year from 4 to 5 years. This is a useful reminder that we should not be complacent with regard to hardware (particularly disk). An unreliable batch of servers places a huge operational load on the team managing server failures and data loss incidents. The current planned four-year lifespan for disk remains appropriate and risk mitigation in terms of multiple procurements remains a well-justified process.

The tape service is currently operating on a mix of T10KC and T10KD generations of tape drive, the C drive writing media at 5TB capacity and the D using the same media at the higher 8.5TB. LHCb and ATLAS were placed on the C media and CMS and other experiments on the D media. However, in order to minimise projected resource expenditure in FY16 (and accommodate the expected unplanned demand of 5PB from the DiRAC² collaboration) a new "cost neutral" solution was found where the T10KC service is closed in 2016 and all experiments are placed on T10KD drives. Financial resource requirements are reduced but offset to a large part by further (capital) tape drive purchases. Migration of ATLAS has now commenced and LHCb will follow in due course. Further T10KD tape drives will be purchased in FY16 to handle the increased load.

² Distributed Computing utilising Advanced Computing <http://www.dirac.ac.uk/>

Up until now it has been possible to accommodate the unplanned (in GridPP5) DiRAC tape media request out of temporary headroom in tape capacity allowed by tape media migrations, however a crunch point is likely to come in late FY16/early FY17 when a migration to the next generation of tape media (T10KE) will be necessary. As yet we have insufficient detail about T10KE costs and media capacity to allow us to judge if we can continue to meet all tape media requests from planned budgets. This will also be affected by LHC data taking rates and a successful conclusion of DiRAC's efforts to use the Tier-1 tape store for disaster recovery.

A further issue is that the T10KE tape media migration (budgeted for out of the hardware budget) will require >£400K resource in FY17 (and a further £150K in each of FY18 and FY19) to fund tape media purchase – unless it proves feasible to reclassify tape media as capital (it was last ruled by finance to be resource) an increased hardware resource budget will be necessary to meet LHC tape media costs. This matter will need to be addressed during 2016 in order to be able to plan budgets for FY17.

The Tier-1 network is becoming increasingly complex and challenging to operate. Development is not as rapid as required and faults are proving hard to trace and rectify. The rollout of IPV6 has been delayed owing to limited effort available within the site network team and the Tier-1's own network upgrades have been delayed by operational issues with the Tier-1 network. It is likely that more effort will need to be committed to networking at the expense of other Tier-1 development projects. In the autumn GridPP plan to review Tier-1 network operation and plans and assess if changes are required.

Production, Operations and Service

Tier-1 operations remained generally reliable over 2015 and early 2016 although some load related issues on the CASTOR tape service impacted CMS and a number of network issues affected all experiments. Average RAL (ops test) service availability since January 2015 remained high at 99.6% compared to a WLCG target of 97%. Average RAL availability for 2015 for the LHC VOs was 98.8%, comparable to the previous 12 months covered in the previous reports. Reliability broken down by VO was follows:

	1Q15	2Q15	3Q15	4Q15	1Q16
ALICE	100%	100%	100%	100%	100%
ATLAS	99%	99%	99%	99%	98%
CMS	99%	99%	98%	96%	97%
LHCb	97%	99%	99%	99%	98%

Figure-7: Quarterly VO availability at Tier-1

Middleware, Grid and Cloud Computing

CPU utilisation was high through most of 2015. Farm occupancy averaged 89% over the period (compared to 81% in 2014) and job CPU utilisation efficiency remained unchanged at 84% (identical to 2014). Since January 2015, RAL delivered 10% of global LHC Tier-1 CPU computing (see Figure-2 earlier.)

The Tier-1 OpenNebula IaaS cloud platform is now in production operation. This work has been jointly funded by STFC's Scientific Computing Department and by GridPP. Many of the Tier-1's routine test and development systems are now deployed on the cloud. The service is also being used by LOFAR who are working to deploy some of their workflows. Other communities such as the H2020 Indigo-Datacloud and Westlife projects are also funding effort to support their use of this private cloud and ISIS too are deploying workflows on it. While OpenNebula has allowed a stable service to be successfully deployed it is likely that in the longer term OpenStack will be required in order to provide missing functionality and better compatibility with other projects.

The CEPH service (called ECHO) is now nearly ready for large scale testing by the experiments. ECHO is intended to replace the Tier-1's CASTOR disk service. Before tests can begin the recently delivered disk hardware must be deployed into the service and CEPH itself must be upgraded to the most recent stable release. This is expected to be complete by August after which a phased ramp up of load tests, beginning with the least critical data will commence.

Management, Business Processes and Communications

Owing to a change in STFC overheads rates for GridPP4+, the maximum affordable (given GridPP4+ flat cash settlement) staff count at the Tier-1 fell from 19.5 FTE to 18.5 FTE. Four (partially funded) staff left during the period, impacting expertise in both the CASTOR team and particularly the Database team where numbers are now below the critical level necessary to assure service. Recruitment and retention is identified as a high corporate STFC risk owing to deteriorating pay and conditions for staff. Actually staffing level at the Tier-1 over the GridPP4+ period (FY15) averaged 17.8 FTE, somewhat lower than planned for GridPP4+. As things stand current levels are close to the 17.5 FTE planned in the first two years of GridPP5.

The Tier-1 team were minor partners in the recent AENEAS proposal submitted to the Horizon 2020 program. AENEAS is a project to develop a design for European Science Data Centre for SKA. If funded, AENEAS will provide 6PM of effort to work on Computing Requirements for SKA. Given the close alignment between SKA and LHC Run-3 requirements the intention is to part fund the Tier-1 manager role in order to work on technology foresight and system architectures.

The Tier-1 manager is currently working reduced hours (30%) owing to a recent change of duty within Scientific Computing Department. This has impacted several activities such as continuous service improvement and several other activities have had to be picked up temporarily by technical staff.

Deployment Status

The last year of operations has seen steady changes and updates within the GridPP infrastructure. This section describes the main areas of progress and issues encountered. At the outset it can be reported that GridPP sites were ready for LHC Run-2 in June with new hardware in place and major infrastructure changes completed.

Resource-wise, since Q115 GridPP sites have seen an order of 8% increase in the available HS06 (the CPU benchmark use in WLCG) and a 20% increase in disk storage available. Both figures represent temporary peaks because equipment purchases with GridPP4+ funding have recently been put into production, but old out of warranty equipment (which across many sites is seeing increases in HDD failures, dead motherboards, PSUs and memory particularly on WNs) has in most cases not yet been removed. The sites achieved a very fast turnaround on purchases once grants were released.

The majority of sites have seen gradually increasing utilization levels across reporting quarters despite the rising resource levels. This has ensured that GridPP delivery to WLCG remained well aligned with the UK pledged resource shares. However, some sites experienced a drop in CPU time utilization in Q215, which was a consequence of the WLCG VOs increasing their use of multicore jobs (wall clock time utilization remained good). The accuracy of some figures taken from the EGI accounting portal around this time was questionable as the number of CPU cores was not correctly captured by EGI – the pre-production accounting portal containing new code did a better job. Even with accurate capturing of the job accounting data it must be noted that HS06 as a benchmark is diverging from actual experiment code performance measures. Recognising this the WLCG Management Board in October 2015 created a new task force to look again at benchmarking options.

Across the reporting period availability and reliability for sites has generally remained well above the WLCG target of 90%. All occasions when the target was not met have been followed-up and reported to the WLCG Project Management Board. Within these there are some incidents of note. In April 2015 a minor network change at the RAL Tier-1 triggered wider problems across the network and led to significant disruption of services – a post-mortem revealed the incident resulted from the incorrect connection of a cable that led to a network loop. Glasgow suffered from power cuts in Q215 and a severed fibre disconnecting its server room from their site access router in Q3. In September Oxford suffered an air-conditioning failure that led to the site being unavailable for a week – during this time the GridPP Nagios service was run from a backup instance at Lancaster. Lancaster itself suffered in December due to a site wide power cut related to severe flooding in the area. Cambridge was also affected by flooding when university networking equipment had to be powered down to prevent damage. Unrelated to the flooding, but leading to a notable downtime, Cambridge moved its cluster to a new machine room in Q3. It should be noted that the distributed nature of the GridPP Tier-2 infrastructure naturally mitigates the impact of these individual problems.

As in 2014, GridPP sites continued during 2015 to move towards better supported batch systems and CEs. The RAL Tier-1 removed support for CREAM CEs in May 2015. ARC and Condor have been popular replacements for CREAM and torque/maui respectively, and UK leadership in this area has been shared in talks presented at numerous WLCG events. As might be expected, the updates have not been entirely trouble free – the ARC sites (such as Glasgow and Bristol) had accounting issues with their new setup for several months in early 2015.

The last year has seen good progress in the area of operations related to Tier-2 evolution. Since Q3 2015, this activity has formed a standing item for discussion at the weekly GridPP operations meeting and continues to report at a biweekly 'GridPP technical meeting'. A GridPP operations project was setup in the CERN JIRA service to track progress particularly in relation to Vac (a simplified approach to VM management developed within GridPP at Manchester). The UCL Tier-2 site was one of the first to be tracked with a ticket in September as it was updated with a new Vac-in-a-Box installation. LHCb promptly ran tests of two payloads per dual processor VM. In October a new GOCDB 'uk.ac.gridpp.vcycle' service was introduced and support for APEL-sync records in Vac added. Liverpool enabled 126 VM slots with Vac in early December and almost at the same time the GridPP DIRAC³ VMs began to make use of a new dirac.gridpp.ac.uk service, which simplified the VM contextualization. At this time support for metadata (EC2/OpenStack) and WLCG machine features (via HTTP rather than ISO image and NFS) were also introduced. By March 2016 GridPP had 5 production Vac sites and jobs for the pheno VO had been successfully demonstrated at Manchester. Progress with standard cloud installs has gone on in parallel at several sites during the last year, for example Glasgow got an OpenStack installation working in Q4 2015. It is worth pointing out that on the back of this experience GridPP sites have had strong engagement with the RCUK Cloud Working Group.

Across the UK we have made some progress with IPv6, but not as much as we would have liked. Whilst some sites like Brunel have moved all Computing Elements, Storage Elements and monitoring to IPv6 capable links, many have only enabled dual stack on their network monitoring perfSONAR nodes. The remainder have been struggling to get their campus networking to allocate IPv6 addresses to them. Of the countries contributing to WLCG the UK is ahead of the curve on IPv6 but increased focus in this area is needed in the year ahead. There is little else to mention about networking during the period as it has functioned well, reflecting our excellent relationship with JANET/JISC and the benefits of the "DRI" investment in site/campus networking a few years ago.

With an eye to the future, Imperial worked with JANET in Q215 to setup a logical routing layer (VRF⁴) to perform tests of a connection to LHCONE. A similar link for the RAL Tier-1 is in progress. This work is being undertaken to ensure that GridPP is in a position to react to any future global requirement of the experiments.

There have not been any significant upgrade campaigns in the last year of operations. There has been a push to have all sites correctly publish the number of cores they have available – this being important for accurate multi-core accounting. In recent months consideration has been given to removing SL5 from all machines. During 2015 the majority of UK SL5 machines were either upgraded or removed from service. There remains a handful to be dealt with during the next month. Alongside this transition from SL5, the UK has had involvement in middleware readiness testing with the experiments for CentOS7. A push for CentOS7 is likely at the end of Run-2.

The GridPP core operations team has continued contributing to various WLCG working groups such as multi-core deployment, network monitoring and http proxy discovery. Many of these led to excellent contributions at the WLCG workshop and associated CHEP conference in Okinawa in April 2015. Our team continued to look at glxexec implementation (for traceability of workload submission) for Worker Nodes but could not find a good solution. Significantly in February 2016 a WLCG decision was taken to freeze further deployment of glxexec. This now poses a risk to ongoing support of the glxexec mechanism and is likely to have an impact on the GridPP DIRAC service used to support non-LHC VOs. GridPP DIRAC has been making use of multi-user pilot jobs and glxexec is essential for the

³ Distributed Infrastructure with Remote Agent Control <http://diracgrid.org/>

⁴ Virtual Routing and Forwarding

credential handling – indeed from June 2015 onwards our ops team worked closely with sites to progress the adoption of the GridPP pilot role. The WLCG decision was made because there are simpler options emerging for isolating jobs – such as using cgroups, VMs and containers – and take up of glegleexec has been slow because of the investment of effort needed by the experiments to make the approach compatible with their job submission frameworks. To keep the DIRAC approach supportable in the medium term now requires the investment of several months of effort to explore alternative security frameworks. In a similar thread, WLCG is currently revisiting the use-cases for the information system components of the Grid with a view to simplification. Like the decreasing support for the LHC File Catalogue (LFC) and Workload Management System (WMS) and freezing of glexec deployment, this is likely to have most impact on non-LHC VOs supported by GridPP.

The GridPP security team has seen a boost with the appointment of a new security officer. The UK, through GridPP, has always been a strong contributor to security efforts in WLCG and EGI – it has been seen positively that of the five security incidents reported within EGI in the last year, 2 came from the UK. It is a credit to the team that they were spotted and dealt with very rapidly. There has been a steady stream of OS and middleware vulnerabilities announced. This is a serious and growing concern for GridPP because frequent patching requires manpower that is needed elsewhere. To give a sense of the problem here is a list of some Q2 2015 vulnerabilities: Zen vulnerability hypervisor memory corruption issue; DIRAC SQL injection vulnerability; VENOM: QEMU vulnerability (May 2015); Persistent XSS in OpenStack - Horizon admin dashboard (June 2015). In August 2015 there were another three issues related to dCache, VOMS and libuser respectively. We then saw critical vulnerabilities announced in nss, nss-util, and nspr libraries in November, and various Java CVE's (with a maximum score) and OpenSSL in December. Moving on to February sites had to respond to a serious glibc remote code execution and also a Linux Kernel vulnerability. A distributed denial of service attack on JANET services in Manchester in December led to access problems with the UK regional monitoring dashboard making it almost unusable for several days.

Keeping sites secure is a core driver for our involvement in WLCG work that is looking at options for improved pooling of data and use of analytics methods to give early alerts through Security Operations Centres (SOCs). We are also making essential contributions to Cloud job traceability work and have engagement with EGI on a “Security Threat Risk Assessment – with cloud focus” project; our staff are taking a keen interest in attempts to provide a usable federated identify management approach. Alongside these technical approaches to security, we are aware that there is a changing risk assessment picture (more targeted phishing; incidents on commercial clouds that our user communities are starting to use more often; moves to more standard software products) that needs to be dealt with through awareness campaigns and training, and for this reason our biannual HEPSYSMAN meeting was again accompanied by a security training session.

The storage group has continued supporting DPM as part of the DPM Collaboration – in February a DPM critical bug in v1.18.10 introduced new drain commands that caused data loss at some sites but fortunately this did not impact any UK sites. The storage group have supported efforts within VOs for data cleanups and catalogue synchronization activities. In February they assisted members of the SNO+ collaboration with their data model.

Users' Reports

ATLAS

Since the last OSC, ATLAS has collected and processed record volumes of data reflecting the new Run-2 trigger rate. It has successfully implemented its new analysis model, with a xAOD format for analysis and regular processing of new or improved xAOD by means of analysis trains. To do this it has made extensive use of Tier-2 resources for simulation, group analysis and individual user analysis, and the Tier-1 resources primarily for reprocessing, simulation and data curation; however, as planned, the boundaries between roles at the Tiers are dissolving to an extent, with increased use of remote access to data from Tier-2s. ATLAS is rolling out its event service model, where sites can access data across the network on an event rather than file level. This aids the reduction in endpoints for data transfer, and is accelerating the shift to large Tier-2s with significant data storage and smaller Tier-2s providing CPU and disk cache. In the immediate future, Tier-2 sites need >400TB of disk storage, but in effect sites with less than 1PB will be effectively be providing a rolling cache.

The UK continues to perform well, delivering roughly 14% of the global joint Tier-1 and Tier-2 capacity by CPU consumption (reflecting the UK authorship share). ATLAS has also continued to make good use of beyond-pledge resources made available by the institutes from their own funds. It has allowed UK groups to exert direct influence on the development of analyses by the provision of over-pledge capacity for work using the institutes' leveraged internal resources.



CPU consumption Good Jobs in seconds (Sum: 3,787,081,325,920)




Figure-8 Sharing of ATLAS Tier-1 and Tier-2 CPU consumption 1 April 2015- 1 April 2016.

The resource at major Tier-2 sites continues to be invaluable to the ATLAS operations. With the declining human resource at smaller sites, ATLAS is working to develop modes and practices that make most efficient use of those sites with the reduced effort.

ATLAS continues on-going work to explore opportunistic resources and new styles of computing, with the UK in the vanguard. The Grid continues to dominate (80% of CPU); but HPC resources make a useful contribution to some workflows, particularly Monte Carlo generation, and there is an almost equal contribution from various cloud provision (see Figure-9). The High Level Trigger farm is used as a processing resource outside of running periods, and we have various demonstrator cloud resources. Indeed, the UK has been using a commercial OpenStack cloud provider (DataCentred) at no cost on a demonstrator basis, allowing them to tune their offering for commercial clients. The GridPP cloud also provides a significant fraction of the scientific cloud use by ATLAS.

Naturally, the UK provides the majority of the ATLAS Vac resource thus far, which is provided on a production basis.

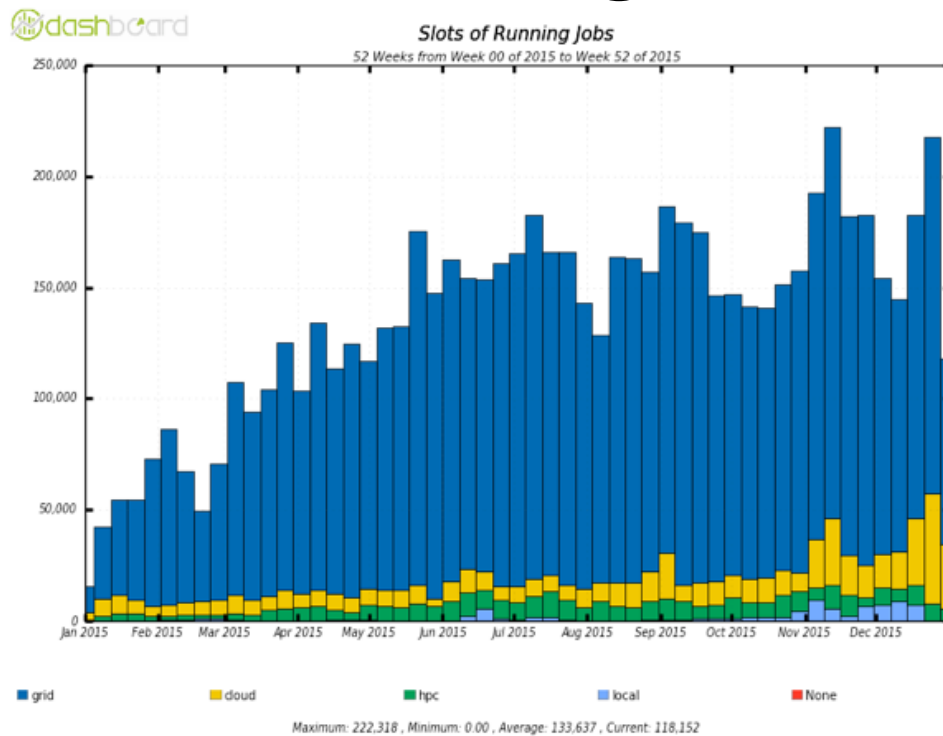


Figure-9 Sharing of ATLAS production between Grid and other forms of provision.

ATLAS Vac sites

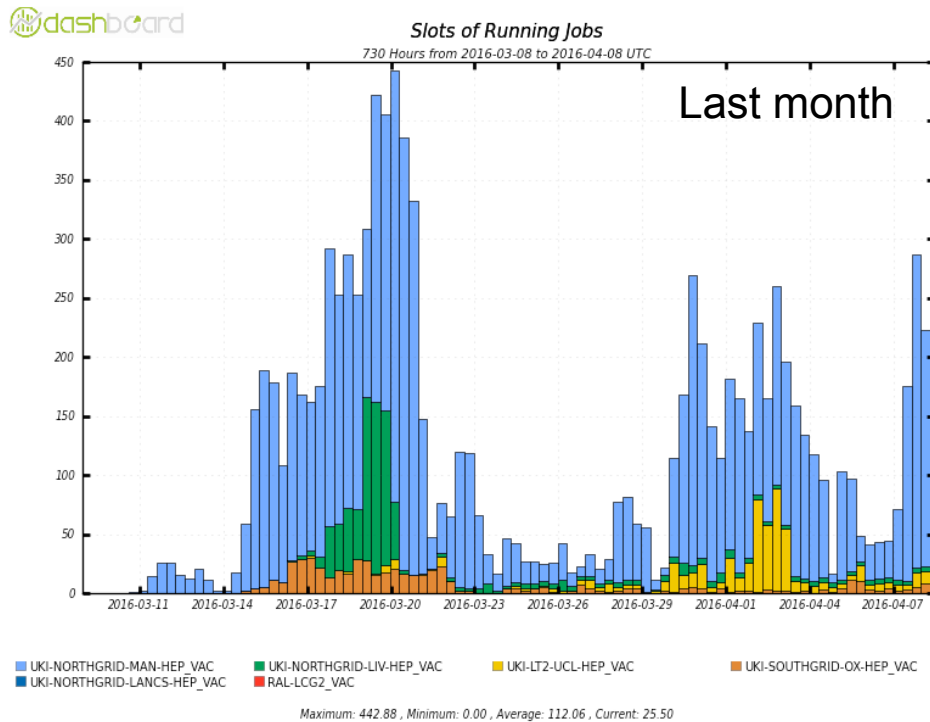


Figure-10 The UK has been providing Vac resources for ATLAS on a production basis.

CMS

CMS has also successfully taken data at the higher trigger rates of Run 2. It has also been a period of implementing the planned evolution of the computing model. CMS computing is expecting to be resource limited during Run-2 and these changes are allowing CMS to be more flexible and to utilise the resources available to it in a more efficient manner. Many of these changes have a direct effect on how UK sites are utilised. We describe three such changes:

1. Dynamic Data Placement
2. Any data, Anytime, Anywhere
3. Global Pool

1. Dynamic Data Placement

CMS data is not accessed uniformly and there are approaching 9 orders of magnitude between the number of times that the most and least read data are accessed. This was measured for 2014 data accesses. This means that the previous CMS model of having all data distributed evenly with sites supporting particular physics groups made little sense. Instead a system is in place, where sites just devote the majority of their storage to be managed centrally. This is shown in Figure-11. This enables popular data sets to be replicated widely and less popular data sets to be deleted from some sites. Which data are popular varies as a function of time and so this process is dynamic.

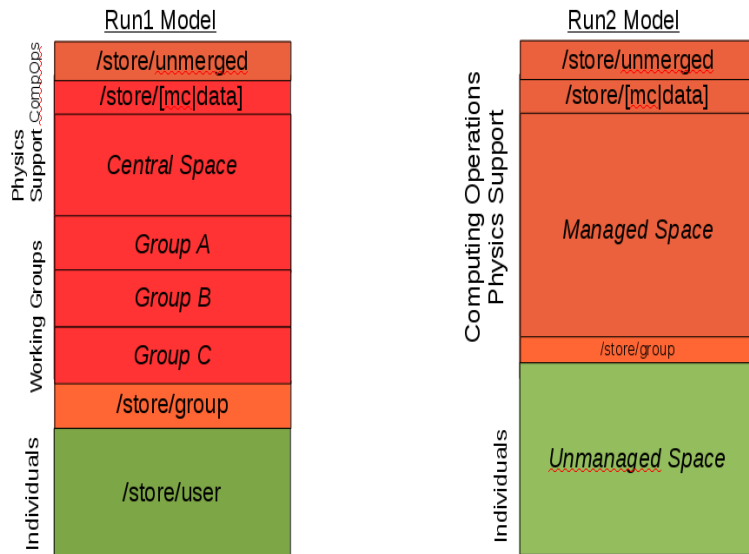


Figure-11: Comparison of how storage at sites was configured in Run-1 and is now configured in Run-2

2. Any Data, Any time, Anywhere

The CMS Any Data, Any time, Anywhere (AAA) project uses the xrootd protocol to allow remote access to data. Through intelligent access of the objects within the data and tuning xrootd this only incurs an inefficiency of a few per cent and enables sites that are busy to be supported by less busy sites that are close by in networking terms. Within CMS this is referred to as “overflow”. A job that has been queuing at a site for more than 6 hours is moved to a nearby site provided one can be found that is less busy. UK sites recently had overflow enabled between them.

Global Pool

In Run 1 CMS had different submission infrastructures for different activities. These have now been merged into a single Global Pool. This allows increased flexibility to run different types of workflow at different sites. This and the separation of disk and tape at the Tier-1 centres, have also blurred the distinction between different computing tiers. It also allows CMS to prioritise different workflows in an agile manner, as can be seen in Figure-12.

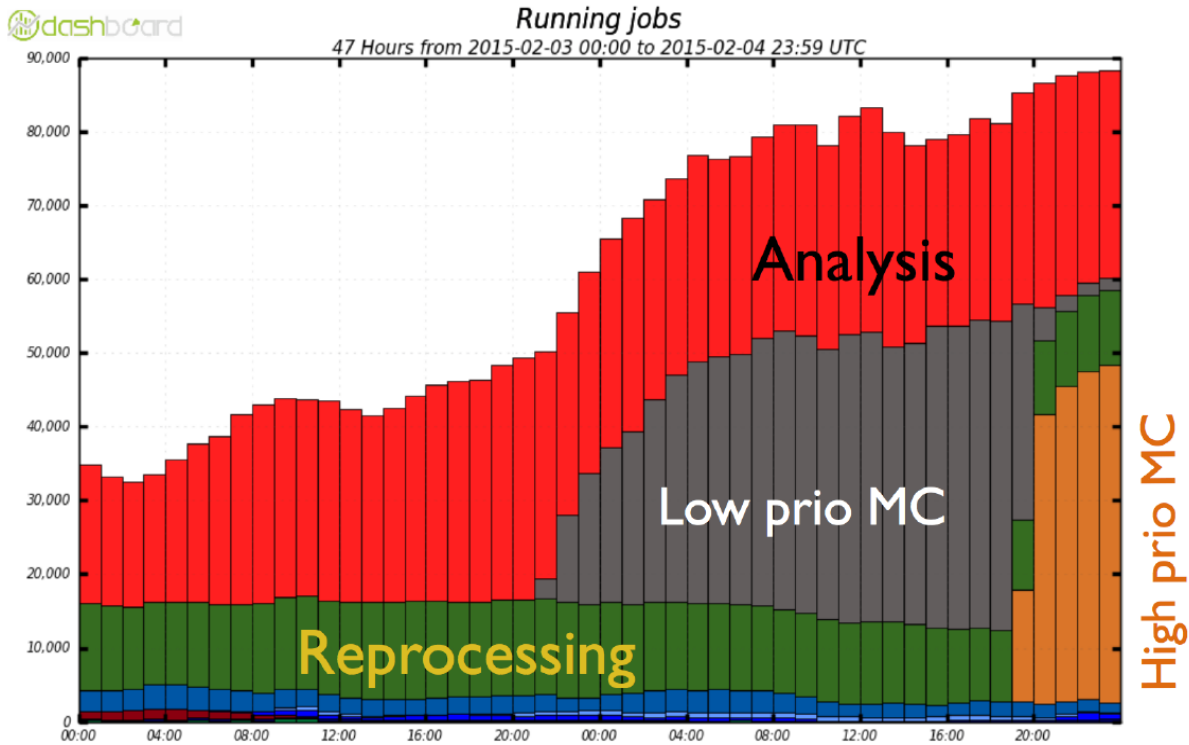


Figure-12: Usage of CMS resources through the Global Pool showing how CMS can prioritise workflows (in this case high priority MC) in an agile manner.

CMS Usage of UK Tier 1 and Tier 2 resources

GridPP is pledged to provide 8% of CMS' Tier-1 resources and 5% of CMS' Tier-2 resources. Figures 13 & 14 show the usage of the Tier-1s and Tier-2s (grouped by country) over the last year. It can be seen that the UK provided just under 8% of CMS' Tier-1 resource and around 6.5% of CMS' Tier-2 resource. The difference between the pledged 5% and the 6.5% delivered is largely through equipment provided outside of GridPP and the efficient running of resources in the UK.

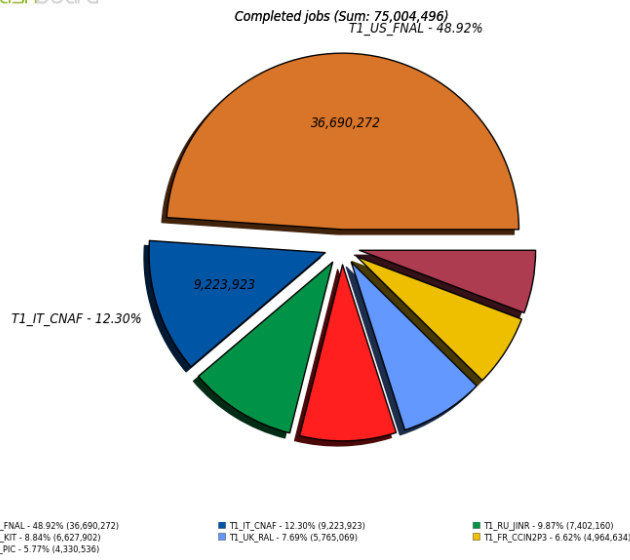


Figure-13: CMS' usage of Tier-1 resources (measured as number of completed jobs) from April 2015 to April 2016.

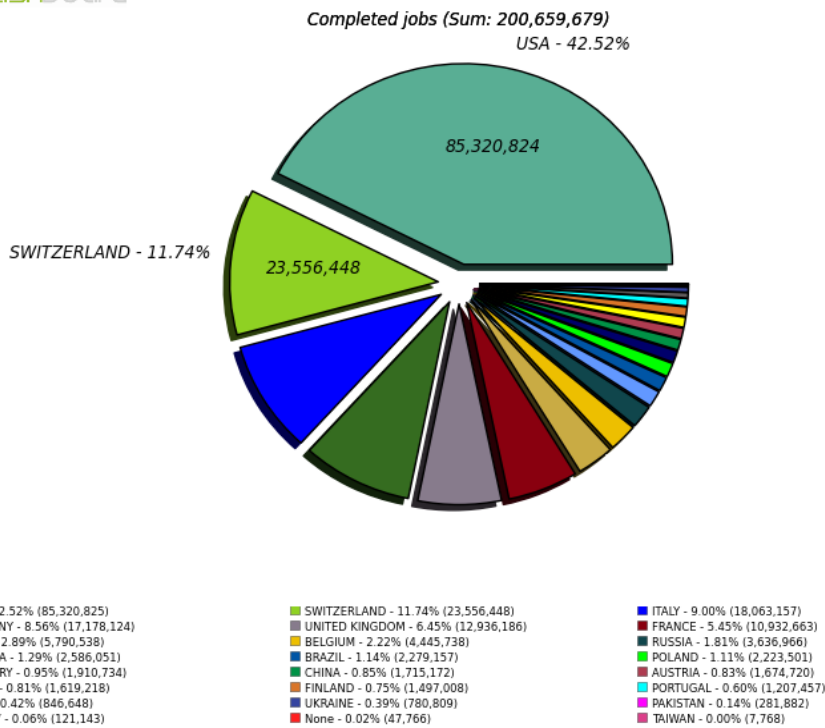


Figure-14: CMS' usage of Tier-2 resources (measured as number of completed jobs) from April 2015 to April 2016, grouped by country.

LHCb

The UK continues to be jointly the largest national collaborator in LHCb along with Italy. We are ~20% of the collaboration for Tier-2 purposes and ~30% of the collaboration within Tier-1 countries. During the last period the UK has continued to be a very reliable provider, both at Tier-1 and Tier-2 sites. This reflects both capacity but also as importantly the reliability and availability of resources. Figure-15a below shows CPU usage at Tier-1 sites (left) and Tier-2 sites (right) during 2015. Similarly Figure-15b below shows the disk usage at Tier-1 sites (left) and Tier-2 sites (right) during 2015. From the LHCb point of view the UK pulls its weight and responds quickly and positively to any issues that arise. There have been no significant problems due to UK GridPP resources.

LHCb has now officially changed its computing model to parallel that of the GPDs, with the expectation that more Tier-2 sites will participate in data processing as well as simulation, and that nominated Tier-2D sites will additionally hold real and simulated DST data for user analysis jobs. The introduction of Tier2-Data centres (Tier-2 centres with data) has proceeded smoothly and successfully. We now have Tier-2 sites in several non-Tier-1 countries contributing and with ~2 PB of disk capacity provided by the end of 2015. As usual, the UK T2D sites at Manchester, RAL, and Imperial have proven to be extremely effective and led the way. More UK T2D sites are expected to join (Glasgow and Liverpool).

LHCb has evolved its computing model for the approaching Run-2 to include three categories of data: (1) The normal prompt data stream to be processed immediately, (2) a parked stream to be kept for later processing and (3) a turbo stream, which is processed within the trigger farm. The prompt stream will be reconstructed after a delay of a few hours, to allow calibration and alignment to be completed. This means that no further reconstruction will be needed, as has been the case in the past, when a full reconstruction took place immediately after the end of data taking in any year. The parked stream may be needed from 2017 onwards. The turbo stream is mainly for charm physics, and assumes that all necessary information can be created in the trigger farm itself, and therefore no further offline processing is needed. These changes will help to reduce the overall CPU requirement, and LHCb continues to use substantial non-pledged resources in Russia, and to make full use of its trigger farm.

LHCb is also developing the means to parallelise its framework, Gaudi, and is very much assuming many core-devices will be the norm in future. LHCb is also adapting its workload management system to support virtual machine architectures associated with the desire to use heterogeneous resources. This activity is being led by the UK in conjunction with the virtualized infrastructure provided by GridPP sites and using the virtual machine lifecycle managers, Vac and Vcycle, developed by GridPP. These developments are already giving LHCb access to resources equivalent to another larger Tier-2.

Following C-RSG recommendations to all experiments, LHCb has put considerable effort into developing better data replication strategy based upon data popularity. In fact LHCb has been very efficient anyway, but still expects to save some data volume.

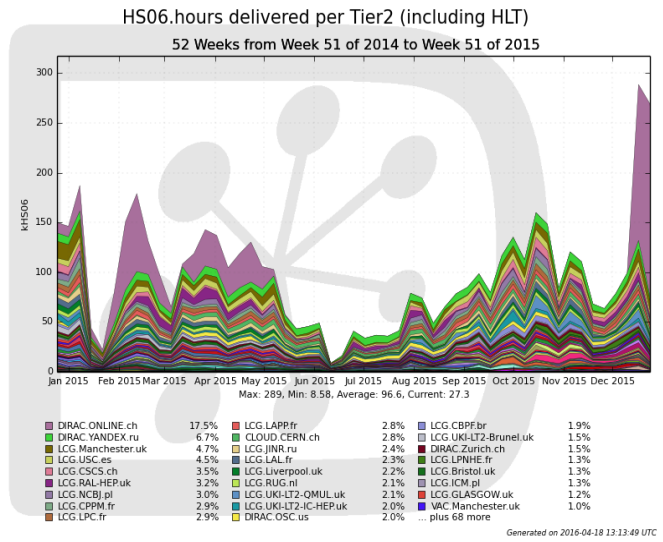
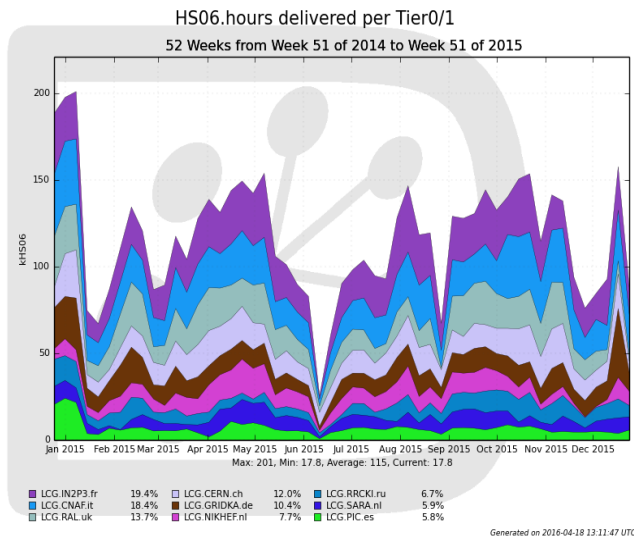


Figure-15a: LHCb CPU usage at Tier-1 (left) and Tier-2 (right) in 2015.

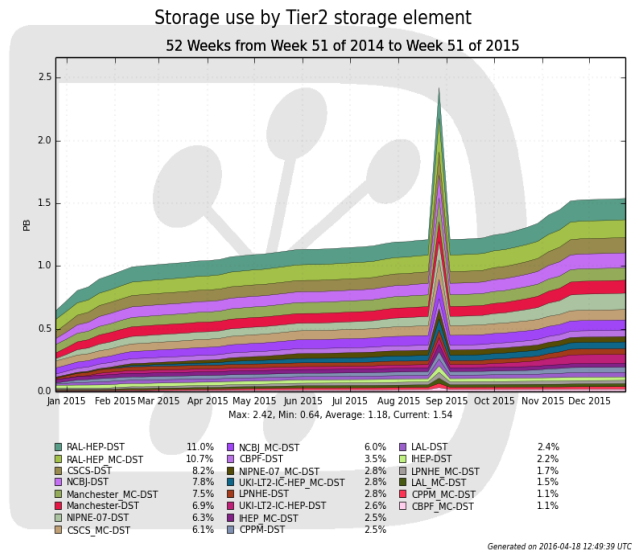
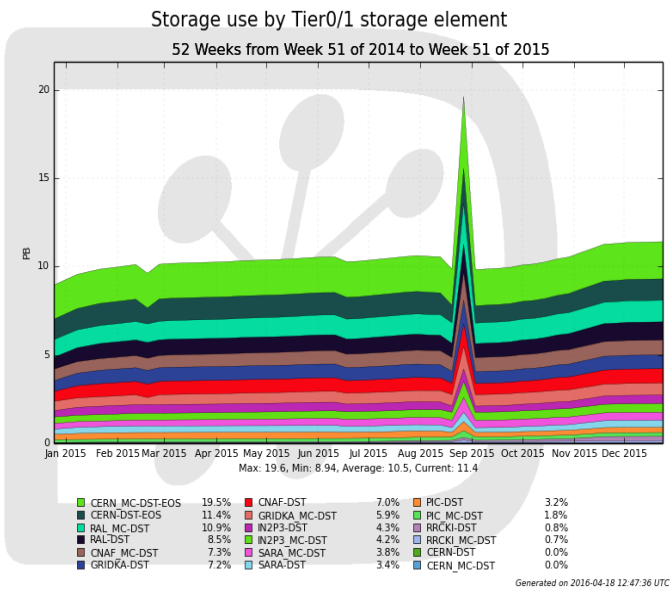


Figure-15b: LHCb disk usage at Tier-1 (left) and Tier-2 (right) in 2015.

Other VOs

We are making steady progress with our other VO communities (i.e. other than ATLAS, CMS or LHCb). Excluding ALICE, the highest other user, the fraction of GridPP Tier-1 resources used by “others” increased from 3% in 2013 to 5% in 2014. At Tier-2s the figure rose from 12% to 14% with Biomed, followed by Phenogrid as the largest users in this category. Recent topics of interest include:

- *Migration to Scientific Linux 6:* Our SL6 migration in 2014 went without major incident for the other VOs, and there was no reoccurrence of the VOMS renewal issues that affected many VOs in early 2013 - when user VO memberships expired with little warning due to a middleware update causing the resetting of all membership accounts to a default date. Problems that did occur during the reporting period include SNO+ writing 1TB files to several sites (e.g. Sheffield) causing disk pools to fill and Nagios tests to fail. Biomed made strong demands on resources such that their job storms at times have saturated a couple of sites (the sites have since implemented maximum job quotas on the VO or in some cases banned individual users – particularly where the users also submitted multi-threaded jobs to single core resources).
- *Infrastructure for smaller VOs:* Workload and data management strategies for our smaller VOs are still a concern, especially as the LHC VOs move away from tools such as the Local File Catalogue (LFC) and WMS. To mitigate major risks in this area we have been developing a new job management and data management service based on DIRAC. DIRAC as a service was in test during 2014 and has received positive feedback from several of our early adopter VOs but particularly Landslides, T2K and CERN@School. In December 2014 a proposal to ramp up use of DIRAC such that it becomes a formal production service was agreed; see the “Impact and Dissemination” section below for further details.
- *Data management:* More work is needed to address increasing issues associated with use of shared storage resources whereby problems for other users of the same space (or other members of the VO) occur. It is hoped the use of the GridPP DIRAC service will also help in this area. DIRAC uses its own file catalogue (the DIRAC File Catalog), and also has built-in metadata functionality, removing the need for users to implement a separate metadata system.
- *Software deployment:* GridPP has continued to encourage the smaller VOs to move their software deployment mechanisms to use the CERN Virtual Machine File System (CVMFS or CernVM-FS). ILC has been one of the most recent VOs to move to using CVMFS, and the Tier-1 made a final decision to switch off its NFS software server in September 2014.
- *Middleware changes:* In mid-2014 GridPP operations sent out warnings to our VO communities about a major change that may affect them. In November 2014 there was a move away from lcg-utils and gfal towards the newer gfal2 and gfal2-utils tools; the older versions then became unsupported. This change required the VOs to actively adapt their frameworks which some, but not all, have done (the old rpms remain installed in parallel to the newer ones but are not available when we move to SL7).

For more on recent engagement with new user communities, see “User community engagement” in the “Impact and Dissemination” section below.

Impact and Dissemination

Overview and status

The focus of impact and dissemination activities over the course of GridPP4+ has been on continuing the New User Engagement Programme, described below. The GridPP Dissemination Officer works partly for GridPP and partly as an STFC Public Engagement Fellow working on "CERN@school and GridPP: harnessing the power of the Worldwide LHC Computing Grid for research in schools and beyond". The CERN@School project that has now become part of the newly-formed Institute for Research in Schools, a charitable trust supporting school students and teachers to develop authentic research in schools. This is allowing us to leverage this work and fully integrate it with GridPP and will give us access to other school projects such as data from classroom based Timepix silicon pixel detectors, the space-based Langton Ultimate Cosmic ray Intensity Detector (LUCID), the Monopole and Exotics Detector at the LHC (MoEDAL), and the TimPix project that uses Timepix data taken during Tim Peake's Principia mission on the International Space Station. During GridPP4+ the emphasis has changed more from dissemination towards impact and two-way engagement with industry etc.

The New User Engagement Programme

Until recently, smaller VOs were presented with a wide array of tools and interfaces for accessing GridPP resources. While this approach offered a great deal of flexibility, sometimes this choice could be overwhelming to new users. Experience had also shown that new groups would often develop (and redevelop) their own ad-hoc scripts and custom tools with the help of local GridPP Tier-2 Sys Admins, placing additional demands on manpower to essentially reinvent the wheel. It was therefore decided that a more uniform approach, using a standard user interface and toolkit, would offer a more scalable solution as GridPP sought to widen its impact outside of particle physics. Industry partners - SMEs in particular - simply do not have the time or the staff to develop custom Grid solutions, so this was seen as a priority over the more traditional "news" aspects of the Dissemination Officer role. With the help of the GridPP Deployment team, and drawing on his experience of working on the CMS experiment (i.e. a larger VO), GridPP has assembled and properly documented a suite of tools to make it much easier for new users, such as SMEs and others, to quickly get started using our Grid resources. This body of work has become known as the "New User Engagement Programme", and this one-time investment in a streamlined package of tools and user guides will make it much easier to attract SMEs and others to take advantage of GridPP resources as we move towards GridPP5.

- *Infrastructure*: Larger VOs (i.e. the LHC experiments) typically have whole teams dedicated to developing proprietary solutions for their grid-based needs. In order to support smaller VOs, the GridPP Collaboration has committed resource to develop the GridPP DIRAC service, particularly at Imperial. DIRAC supports the complete Grid workflow including job submission, data management and metadata handling. Adding metadata functionality has proven particularly useful for smaller VOs, as this has not traditionally been covered by LCG toolkits. By offering a GridPP-developed solution for smaller groups, the entry barrier for new users is significantly reduced and the scope for achieving wider impact has been increased.
- *Tools*: In addition to DIRAC, three tools have aided these efforts. CVMFS allows VOs to easily deploy software to worker nodes via https. RAL hosts a GridPP-supported Stratum-0

CVMFS repository where new user communities may deploy their software. The GridPP CernVM is a Virtual Machine (VM) image based on CERN's CernVM technology that anyone can download and use to create a grid User Interface (UI), giving them access to the Grid (and, handily, any software deployed via CVMFS). In addition to removing a huge barrier to interacting with the Grid, the GridPP CernVM also serves as a "virtual Grid node" where users can compile and test their software, data and workflows before using it on the Grid. This greatly reduces the time needed for development and testing. Finally, the Ganga software suite provides a powerful yet user-friendly Python-based interface to the Grid that can interface with the GridPP DIRAC service. GridPP is also developing Cloud interfaces to its resources and these coupled with increased virtualisation of services should also make it even easier to attract new communities to our resources.

- *The new GridPP website and documentation for new users:* A significant amount of work during GridPP4+ revolved around the development of a new public-facing website for the GridPP Collaboration. This launched in December 2015 at <https://www.gridpp.ac.uk>. Based around the FoundationPress framework, the new website is mobile-friendly and aimed at informing and engaging potential new users with the GridPP Collaboration. The new website also hosts the GridPP UserGuide at <https://www.gridpp.ac.uk/UserGuide>. Powered by GitBook, this provides new users with a guide to getting on the grid using a GridPP CernVM and the GridPP DIRAC service, as well as offering advice on subjects such as getting a Grid certificate and data management. While it will never replace the support new or existing users can receive from the GridPP community (via support mechanisms such as the JiscMail GRIDPP-SUPPORT mailing list), it nevertheless offers a single-path to Grid access for complete beginners and will continue to evolve over the course of GridPP5.

User community engagement

The last year has witnessed significant new interest from wider academic communities in using GridPP resources. An initial slow start in moving data for HPC DiRAC from Durham to the RAL Tier-1 (due to the large number of small files) has been replaced with a sustained present rate of about 12 TB a day from Durham and Leicester. The limiting issue now becomes available bandwidth at the source sites. Excellent progress was achieved for Galaxy Dynamics work done in conjunction with researchers at the University of Central Lancashire. Interest was also received and followed up for several other astrophysics communities including: the Low Frequency ARray (LOFAR) who are working with the RAL Tier-1 for access to their cloud and new Echo service via S3⁵; The Large Synoptic Telescope Project (LSST) are working very closely via GridPP support at Edinburgh and Manchester (the users are currently running jobs at 8 GridPP sites); LIGO – who plan to use Condor glidein jobs to create a Condor pool on top of the RAL Tier-1 pool; recently in 2016, Euclid⁶ who are moving initially to a 'Central UI' approach to get jobs running quickly. Support of Lux Zeplin (LZ), a dark matter experiment based in the US for whom GridPP's contributions included running Monte-Carlo for their Technical Design Report, has been very well received.

The increase in VO support activity has led to two changes. Firstly, the creation of a new GridPP-Support email list, and secondly the establishment in May 2015 of a standing item at GridPP Operations meetings where the status of new VOs can be regularly reviewed with reports being generated by a "GridPP VO champion".

Some members of the UKQCD community made contact in 2015 to resurrect their VO (it was previously hosted by GridPP). They ran simulations, which led to a publication on improvements to

⁵ S3 refers to the interface used by the Amazon Simple Storage Service

⁶ Euclid is an ESA astronomy and astrophysics space mission

glueball calculations. Following this they submitted a consolidated grant application that made reference to GridPP support. GridPP also featured in an additional funding request for PRaVDA who previously ran simulations in 2015 on the GridPP infrastructure to model interactive proton computerized tomography (similar work was done via the Cambridge based GHOST project). SNO+ data transfers to a GridPP SE started in Q1 2016 and the SuperNEMO collaboration have also re-engaged. Initial discussions have been had with interested members of DEAP-3600 and discussions with the Dune collaboration are about to be arranged. IceCube are exploring GPU usage at Manchester.

Whilst the number of supported active VOs has been growing, the Operations team have also during this period undertaken a VO cleanup campaign to remove old (e.g. Hone (the H1 experiment), obsolete and duplicate VOs from the configurations at sites. Finally, work continued for established VOs such as pheno and T2K. The most significant use of GridPP resources over the last year for non-LHC VOs has come from ILC (3%); pheno (2%) and biomed (2%).

In anticipation of further engagements the core Operations team have recently reviewed key documents maintained by the collaboration including the GridPP User Guide – this received a complete usability test via a summer student in 2015. A VO-GridPP MoU proposal was drafted in January 2016 and some example templates to guide initial interactions between GridPP members and new communities has been created within the last month.