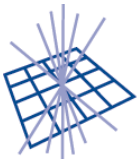

GridPP Project Management Board

Lessons Learnt

Document identifier :	**GridPP-PMB-157-LessonsLearnt.doc**
Date:	**17/05/2011**
Version:	**1.0**
Document status:	**Final**
Author	**PMB**

Introduction

Given the great success of the project as evidenced by the ability of physicists to analyze the data from the LHC from the day it started, there is probably not very much that with hindsight we would have done very differently within the project. In fact many lessons were learnt and adjustments made as we went along which resulted in the PEGASUS project describing our working practices as '*Collective Agility*' as discussed below.

In the next two sections we give, in no particular order, some of the lessons we have learnt in terms of general observations along with some of the more technical issues. This is followed by our interactions with STFC, including the Oversight Committee and the Office. Finally we record some of the positive interactions, including external reviews.

General Lessons Learnt

Tier-1

The RAL Tier-1 has become a world-class facility that performs well when compared with peers world-wide. We believe the close working relationship with GridPP has proved to be a successful model and that a more arms-length relationship would not have yielded such a high degree of success.

There were many initial problems with the new R89 building at RAL. The formal feedback from the GridPP Tier-1 review stated: "*The panel notes the many of the 2009 problems at RAL were related to the new building R89. The panel urges STFC to note that this design-and-build approach failed to produce a building on time or with the expected specifications. To make this type of approach work, a much more effective oversight of the project is required by, on behalf of, STFC."*

Tier-2s

We believe that our decision to try and engage all institutes in our Tier-2 framework has been exonerated despite pressure from funding committees to reduce the number of separate sites. At issue is the trade off of engaging more people and leveraging more external resources from them against the danger of spreading diminishing manpower too thinly. We estimated that between 2005 and 2010, the institutes contributed approximately £18M to GridPP in terms of capital costs (machine rooms etc), hardware not provided by GridPP (SRIF, CIF etc), electricity and additional manpower. We hope to maintain this inclusiveness although in GridPP4 we have compromised somewhat by introducing the concept of 'Analysis Sites' and 'Production Sites' where the former are better provided with GridPP funded manpower. The distributed UK Tier-2 infrastructure has proved to be a success from both the technical and financial points of view, and the direct engagement of physicists and institutes has been invaluable in ensuring a user-led project.

User Engagement

Engagement between the providers and the users is essential, both in the project as a whole and at the site level. In GridPP there was the famous 'credibility gap' whereby Grid funding could not be used to fund Grid people in experiments and Grants funding could not be used to fund experiment people working on the Grid. In fact having people working at the interface was essential (e.g. the funded experiment posts at RAL and user support posts elsewhere) and this is now recognised in GridPP4.

ATLAS

The lessons learnt for ATLAS are the following. First, it is important that ATLAS takes a significant direct role in the job scheduling and data placement. Second, data integrity on Tier -1 sites remains an issue, and the solution is to design operational systems to take this into account; fortunately, the computing model had anticipated this from a resource and replication point of view. Third, the UK will continue to need 'in house' effort to develop dashboards, liaise with ATLAS centrally for operations and especially for user support.

Middleware Development

Although not a GridPP specific lesson, it is clear that the centrally provided middleware is too limited, especially in the area of data management. This has led the experiments to build their own bespoke solutions with little commonality. This has meant that it is very difficult for new communities or experiments that do not have the manpower to develop their own middleware. More effort should have been funded to work on the common underlying data management middleware.

Metrics

Although it is desirable to set milestones and metrics at the start of the project, these, especially metrics, should not be set in stone. The user requirements evolve with time. The metrics need to be provided by the users (experiments) rather than the sites. In general the users have the best tuned sources for accounting and performance; but even these evolve rapidly. The providers/sites need to know what is being measured but these need to be responsive to user needs. GridPP has always allowed the metrics in the Project Map to change and some have been dropped if they have become unsuitable or unmeasureable. These changes have of course always been agreed by the Oversight Committee.

Security

It is well known that managing the risks related to Information Security is a continuous process. Security policies and procedures need to be regularly revised to address the current requirements and threats. Security vulnerabilities that come to light in the Grid middleware also have to be handled appropriately and fixed before these turn into real security incidents. GridPP has international leadership roles in Security Operations, Software Vulnerability Handling and Security Policy preparation for EGEE, EGI and WLCG. We learned that building trust between many diverse computer centres and countries requires perseverance and diplomacy. The achievement of all collaborating Grids and sites in WLCG operating securely together in a common security domain and all following common policies and procedures has been a satisfying outcome. This will need constant attention and is indeed being taken forward into GridPP4, EGI and WLCG under GridPP leadership.

UK e-Science

We got the impression from attending many 'All Hands' meetings that the rest of UK e-Science was not particularly interested in GridPP despite our successes in our own field and others such as bio-informatics. This is particularly ironic given that many of these projects have now fallen by the wayside and others are looking towards 'Clouds' that are functionally very similar to parts of our infrastructure and not necessarily cheaper. The dichotomy arose because in order to ensure success LCG needed to commit to the EDG-EGEE middleware at an early stage, imposing a pragmatic but possibly less than ideal solution to the problem. With hindsight and the delay of the LHC more could probably have been done to engage these communities.

Doom Bar

As noted by PEGASUS, Doom Bar always helps with any problem.

Technical Lessons Learnt

Disk is more difficult than CPU

At the start of GridPP, the emphasis was on Monte Carlo simulation production that dominantly used CPU resources. These were relatively easy to provision and operate. Now that the LHC is running the emphasis has changed to analysis that needs very large amounts of disk (and tape) storage. Disk is much more manpower intensive to operate and there have now been several procurement issues. As a result the Tier-1 and some Tier-2s have started buying storage from two independent vendors with different disk/controller combinations in order to reduce the effect of failures of whole batches.

Networks are better than envisaged

The architecture of the wLCG is based on MONARC simulations from the late 1990's that vastly underestimated the bandwidth and reliability of the academic networks and the costs of dedicated links. We have leant that networks are a powerful resource that should be optimised technically and financially in

balance with CPU, disk and tape. Presently, the LHC experiments are re-examining their computing models to make better use of networks. One likely consequence will be a reduced reliance on custodial quality disk storage as data becomes more dynamic.

Testing

It is clearly necessary to test all aspects of the infrastructure as completly as possible but we have learnt that even if all of the separate components seem to be working, the complete system may not or may not be performant. We have learnt that end-to-end throughput tests, which first started with the so-called 'Steve Lloyd' tests, through the ATLAS Hammercloud tests to the recent ATLAS sonar tests, have proved essential in detecting either mis-configurations or bottlenecks at sites. The STEP09 full-scale end-to-end tests before the LHC switch on were an essential step that led to the Grid's success for real data.

Air Conditioning and Power Fails

Air condition and power supplies to computer rooms can and usually do fail, especially in the early days of a new computer room. It is essential to have automatic shutdown of systems when very high temperatures are reached.

Interactions with STFC

Interactions with the Office

GridPP has always had a very good relationship with the office staff at STFC who have always been very helpful and encouraging. We also believe that the staff have also found GridPP relatively easy to deal with. This is partly due to GridPP having been reasonably funded in terms of management effort. We believe that for large projects such as GridPP it is vital that there be properly financed project leaders and project managers and that key academics are bought out as required to oversee the project.

The GridPP Oversight Committee

GridPP has always had a very good relationship with its Oversight Committee and is grateful for many helpful ideas and supportive comments over the years. This relationship has evolved as the project has matured. At the start while the project was still being formed and there were many strategic decisions to be made, the Oversight Committee had close involvement and helped shape the project and its reporting, through for instance, the design of the Project Map. As the project has matured to the point where we are in continuous production mode, the Oversight Committee has needed less influence and STFC has sensibly decided to go for a lighter, less frequent review process.

Finances

The decreasing financial flexibility and increasing financial uncertainty of the lifetime of GridPP3 have been increasingly problematic. We recognise that these are largely as a consequence of factors beyond the control of STFC. However, being asked on several occasions to change spending plans by up to a million pounds with short notice, inevitably leads to problems and inefficiencies. The current rigidity in the system between 'capital' and 'resource' and between one financial year and the next, are particularly costly.

JeS Forms

The requirement to provide individual JeS forms for all posts from all institutes prior to approval was a hindrance and partly a waste of time as it was clear that the project would have to be re-scoped following approval and posts would probably be cut or redeployed. The requirement to specify exactly who would do what, where, many months before the project started meant that some decisions had to be made much earlier than they needed to be and hence were not necessarily optimal. JeS forms involve a huge amount of internal bureaucracy and were only necessary for the final posts that were approved. However we do appreciate that this process was improved for GridPP4.

eVAL

eVAL is a huge bureaucratic nightmare completely wrongly configured for capturing input from large collaborations (which is the encouraged way to research these days). The whole system should be set up by project with input by individual institutes only where absolutely necessary. It is crazy to be organised by individual project grants, many of which are very similar.

Final Reports

We have always appreciated the very sensible decision that we have not had to provide final reports for all the individual project grants since the project has been continuously reviewed by the Oversight Committee and everything that could be known is already known by the office. We very much regret that this decision has been reversed and we are now faced with extra pointless bureaucracy.

SSC

Luckily the disaster that is SSC has been somewhat mitigated in GridPP as a large fraction of the funding goes through the universities. However the RAL expenditure and the GridPP travel budget involve SSC and despite GridPP having a thorough understanding of its own finances the inability to get information from SSC in a timely or useful manner meant that we were flying blind at the end of the project and unable to optimise the use of budgets, e.g. travel.

Positive Interactions

The Grid Works

Clearly the most positive outcome is that the Grid works. Since the LHC turned on in late 2009 the data has flowed almost seamlessly from CERN to the Tier-1s and Tier-2s and been analysed by thousands of physicists around the world. From initial scepticism physicists have come to accept the Grid as an everyday tool for their work and it is rarely actually noticed or commented on nowadays having been incorporated in everyone's workflows.

Knowledge Exchange

From the outset, GridPP decided to let other projects and disciplines use the resources if they were not saturated by LHC work. In fact 10% is currently set aside for non-LHC activities. As well as being a benefit to GridPP itself in terms of testing, especially in the early days, this policy has allowed GridPP to contribute to projects such as Malaria and Avian Flu studies as well as enabling collaborations with imense and iLexIR at Cambridge (Camtology) and Econophysica (mathematical models for commodity trading).

As well as the LHC experiments ALICE, ATLAS, CMS and LHCb, GridPP also supports many other UK particle physics experiments and activities such as T2K and UK phenomenology. In addition GridPP provides, or has provided, resources for a broad range of activities from other disciplines for example Earth Sciences and NanoCMOS, a total of 95 Virtual Organisations.

We learnt early on that in order to do effective dissemination you need dedicated people with the right background and training – not best efforts from other members of the team. In addition you need to pay proper professional graphic artists etc for publicity material, logos etc.

PEGASUS

From 2006 to 2010, an EPSCRC funded project called PEGASUS at the London School of Economics studied GridPP. See *Particle Physicists Engagement with Grids: A Socio-technical Usability Study* (http://www.pegasus.lse.ac.uk/index.htm). PEGASUS describe GridPP's management practices as '*Collective Agility*'. They conclude:

"*We suggest that the particle physicists, while not following any pre-defined agile methods, are aware of the challenges they face and have made deliberate and substantial effort to achieve a suitable development process. … The agency and knowledgeability of members of the Grid project are central in this. While no one serves as the mastermind of the project, the interaction and coordination among them give rise to a 'collective mindfulness' with 'a rich awareness of discriminatory detail and a capacity for action'. It takes effort to maintain this collective mindfulness, without which distributed agility would not be possible or sustainable. Therefore, while agility can be described as an emergent property of the distributed collaboration, such emergence is very much enacted, not without deliberation and reflection, as it is instantiated in day-to-day practices*."

The key organisational practices in GridPP identified by PEGASUS are:

- *Draw upon past experience to handle new tasks;*
- *Continuous reflection and learning;*

- *Extensive communications within and between different groups, with an emphasis on face-to-face informal communication;*
- *Work with power users; cultivate user communities;*
- *Project leader articulates clear vision and shared goals;*
- *Use high level milestones and deliverables to create momentum, but be ready to change them;*
- *Share knowledge by mailing lists, wiki, blogs, etc;*
- *Cultivate community bonding and shared identity;*
- *Develop trust, loyalty and mutual support;*
- *Motivate and rely on good people;*
- *Maintain high level of transparency within the project;*
- *Allow mistakes and unsuccessful explorations;*
- *Allow parallel solutions to compete*

UK e-Science Review

In 2009 there was an international review of UK e-Science. Overall the report concluded that the programme is by all accounts and measures "*World leading*" and "*Empowering*". As such the successes of GridPP were mentioned explicitly within the report along with the NGS as having "*been highly successful, providing many users with access to more computing power than they could otherwise easily obtain*". This is not all that GridPP does however and the technology and expertise developed by the project within the area of large-scale, distributed data management and analysis was also highlighted. As with many of the other projects GridPP has been active in working with the the business community and the collaborations with imense and iLexIR at Cambridge (Camtology) as well as Econophysica (mathematical models for commodity trading) at QMUL were lauded as examples of the effect academic projects could have on the wider community.