GridPP Project Management Board

GridPP5 – Deployment, Operations & Support

Document identifier:	**GridPP-PMB-165-GridPP5-DOS**
Date:	**19/02/2013**
Version:	**1.1**
Document status:	**Final**
Author	**Jeremy Coles/GridPP**

Contents

1. Introduction

This document has been prepared to capture the roles and activities that are evident at this stage in GridPP4 with a view to informing the GridPP5 proposal in the areas of Deployment, Operations and Support.

At the core of GridPP there are three primary areas needed to sustain a fully functional computing infrastructure. These encompass procurement, deployment and testing of hardware and software at sites, operational monitoring and following up and resolving problems, and finally supporting the site staff and user communities. These are augmented by experts working in specific key areas that are essential to the secure and efficient running of a Grid; including security tasks, accounting, and storage and data management.

Within GridPP we have developed a very efficient approach that divides essential core tasks across staff already needed by the project to undertake deployment and support work at sites – this allows them to contribute in several areas according to their skills and maximise their benefit to GridPP, WLCG and to the experiments.

The next section provides more insight into the areas. It is followed by sections that describe how the GridPP project has evolved to cover the tasks within these areas; how current structures interact with international partners (for example WLCG and EGI) and the experiments; and how wider non-LHC and non-physics are supported. The final section uses a table to summarise the effort allocated across GridPP institutes according area.

Appendix A provides short summaries of the relevant posts in the areas of Deployment, Operations and Support as funded by GridPP in January 2014. Appendix B extends the thinking in several sections with a focus on the future of Data Storage and Management.

2. The areas – Deployment, Operations & Support

Whilst there are overlaps between the areas of deployment, operations and support, it will help to distinguish the main type of activity that each one covers.

2.1 Deployment

Deployment covers the activities that together take us from a situation of having hardware in physically distributed machine rooms to a usable distributed infrastructure. This starts at the **fabric level** where we agree on hardware needs and then investigate, select and implement solutions for computing cluster management (such as puppet – which is coming to replace the widely used cfengine) on top of the purchased physical machines. GridPP staff contribute to evaluating and documenting approaches to cluster management and then ensure integration with the wider software environments that are being used. This area is never static because hardware changes, software products mature or lose support and new more efficient approaches are adopted. In addition (and in a similar way) the base operating systems put onto the hardware develops over time and these must be quickly patchable when vulnerabilities are identified.

With the bare machines setup the next level of deployment concerns the services that are to be run. GridPP sites deploy traditional batch scheduling systems, fabric monitoring and additionally Grid and experiment/VO middleware and software. Together this allows for the translation of specific computing tasks so that they can be effectively shared, run and monitored across a distributed infrastructure. The middleware comes in many forms but our sites generally support a common set that has until recently been packaged for installation by consortia such as the European Middleware Initiative (EMI) and further tested and integrated by the European Grid Infrastructure (EGI) project. GridPP has by necessity been involved in various aspects of the packaging and testing work to ensure that the solutions are fit-for-purpose in the HEP computing environment – checking for bugs and functionality improvements, ensuring that solutions scale sufficiently and where necessary providing for infrastructure redundancy. The experiment software – which is continually evolving to cope with increasing demands and improved methods - also needs to be tested and validated outside of the developer's test clusters, and this is again an area where deployment effort is required. We perform integration testing, validate functionality and also take part in setting strategic directions and final rollout. Throughout contributions are made to documenting the processes and results.

In summary deployment is a multifaceted activity that is essential for GridPP. Where possible we work with partner projects, such as WLCG, HEPiX, EGI and others to share the work that needs to be done. The actual work spans assisting with packaging products, integration and code testing on resources that are provisioned for the task, right through to ensuring the widespread rollout of an accepted upgrade or component. GridPP has often been at the leading edge of such work and demonstrated its own efficient approach by being the first to fully implement technologies such as CVMFS and have its sites transitioned to SL6.

Table 1 below gives an overview of the **production** core infrastructure services run across GridPP sites – it expands on a similar table in the document "Tier-1 in GridPP5". It should be noted that all sites are expected to run instances of: gLite-APEL/APEL; site-BDII; SRM; CREAM-CE/ARC-CE and as part of WLCG these are complimented by perfSONAR; glexec; squid; UI and ARGUS (see the Technical Overview document for explanations of these). Some also run extra functionality services such as MPI (Glasgow and QMUL) and cloud interfaces, or support experiment run services such as job submission factories. To aid testing, many sites also run non-production instances of services. For completeness it should be mentioned that there are other Grid services run outside of the UK upon which our sites rely, this includes the Global Grid User Support (GGUS) helpdesk, an Operations Portal (that captures information that for example allows sites to configure their services for specific VOs), LHC VO Site Availability Monitoring (LHC SAM) as well as management tools like the WLCG Resource Balance & Usage (REBUS) system.

Service	**What it Does**	**Hosting site(s)**
APEL	Collects site accounting data from across all WLCG/EGI and	RAL Scientific Computing Division

	OSG sites.	
NGI Argus Authorization Service	Provides the authorisation service for distributed services (e.g., user interfaces, portals, computing elements, storage elements)	RAL Tier-1 (NGI)
CernVM File System (CVMFS)	A network file system based on HTTP and optimized to deliver experiment software in a fast, scalable, and reliable way.	RAL Tier-1 – Stratum-0 (non-LHC VOs) RAL Tier-1 - Stratum-1 (LHC VOs).
File Transfer Service (FTS)	Provides a reliable data movement service for transferring files between Storage Elements. In addition, it provides features for administration and monitoring of these transfers.	RAL Tier-1 (All)
Frontier distributed database caching system	Provides ATLAS conditions data to jobs running across the Grid. Provides squid caches to local ATLAS and CMS batch jobs	RAL Tier-1 (ATLAS)
Grid Operations Centre Database (GOCDB)	It provides a repository for site data such as name, contact information, services available, downtime information and project membership	RAL Scientific Computing Division
LCG File catalogue (LFC)	Offers a hierarchical view of files to users, with a UNIX-like client interface. Depends on a back-end Oracle database.	RAL –Tier-1 (non-LHC VOs)
MyProxy	The MyProxy service provides a certificate repository that allows proxy certificate renewal for jobs.	RAL Tier-1 (All) Oxford (SAM)
NGI-SAM (Site Availability Monitoring)	Sends out regular Nagios test jobs to sites and uses the results to provide problem alerts.	Oxford Lancaster (backup)
perfSONAR dashboard	Used to collect and present results from UK site networking tests performed by a mesh of site perfSONAR installations.	Imperial (NGI)
Top-BDII information service	The BDII makes available information about the presence and configuration of resources on the GRID.	RAL Tier-1 Manchester Imperial
Virtual	A system for managing	Manchester (Admin&Core)

Organisation Management Service (VOMS)	authorisation data. An Admin interface allows VO admins and members to manage their membership or VO. A Core database component used to generate Grid credentials for users when needed.	Oxford (Core) Imperial (Core)
Website	The GridPP website runs various services for the project coordination including a twiki documentation system.	Manchester
Workload Management System (WMS)	Accepts user jobs, assigns them to the most appropriate Computing Element, records their status and retrieves their output. The Tier-1 provides a WMS service for the UK and beyond.	RAL Tier-1 Glasgow Imperial College

Table 1: Core infrastructure *production* services run across GridPP sites. A ‘core’ service is one run for the benefit of the whole infrastructure.

2.2 Deployment considerations for GridPP5

During the years of GridPP5 there will be many changes in hardware technology (for example further evolution of many core chips), hardware provisioning and usage (for example via Virtual Machine and cloud environments), software and connectivity (including a move to IPv6) and in the experiment computing model frameworks (increasing demand for multi-core queues and whole node scheduling). To keep up with such changes it is absolutely essential that GridPP maintain an effective deployment capability that is well coordinated with partner projects.

In 2011, WLCG established a set of Technical Evolution Groups (TEGs) that sought to simplify the set of software that must be deployed, operated and supported with the aim of reducing the overall operation and support costs of the project. It was understood that by moving away from HEP specific solutions (through use of standard software, protocols and tools) the WLCG infrastructure would become easier to maintain, operate and evolve. But, it was also noted that where there is a unique need, such as with globally federated data distribution and access, the target would be to produce tools and services that might benefit other communities who may soon face similar problems. The discussions have already led to improvements in Run-1, drawing heavily on operational experience gained in early data taking. The TEGs together with wider work within WLCG lead to projections that impact deployment, operations and support during GridPP5. Deployment needs are constantly evolving and the following observations are relevant:

- Overall simplification of the software and middleware to be deployed both in terms of the functionality and deployment. There is an aim for near stateless sites needing minimal operational effort;

- There will be code development that involves parallelism at all levels from multi-core and multi-threading down to instruction-level parallelism, vectorisation, and the use of coprocessors. Re-engineering of frameworks and algorithms will demand changes in our hardware and site configurations;
- Sites are becoming ‘functionality bound’ rather that strictly 'tiered’. The functions put different demands on sites according to these categories: data archiving; large-scale data serving; batch capability (high i/o or not) and running of infrastructure/experiment central services;
- Reconstruction and reprocessing activities need to remain close to data archives in order to avoid large network data transfers. It is expected that data serving sites will continue to deploy with large reliable disk storage and significant network capabilities;
- Improved network capabilities have already allowed a significant evolution of parts of computing models and associated simplification. Continued deployment of upgrades incorporating new techniques to manage resources are expected. Already the File transfer service (FTS) no longer needs SRM and works with a mesh model;
- The activities of event generation and the full (or fast) simulation steps are compute-intensive and have relatively little i/o and can be successfully deployed on opportunistic resources such as cloud or HPC systems with little or no local storage;
- Opportunistic usage will increase as site configuration needs reduce via standardization and improved working (for example through web caches and CVMFS). This is much easier once no long-term site state is required;
- The cost of commercially procuring computing may become competitive and lead to hybrid cloud models. Virtual data centers may also become possible;
- Provisioning of opportunistic resource (e.g. HPC and Cloud/IAAS) with transparent interfacing to the experiment distributed computing environments is becoming possible;
- There are suggestions that campuses should deploy secure science data enclaves to meet needs of HEP and other data-intensive disciplines.

Against these broad shifts in the infrastructure, there are going to be site level advances and considerations including:

- That site build systems are moving to open source tools – Grid middleware is being embedded in an open-source environment;
- There are marked increases in computing and storage needs in coming years. To keep up computing center architectures (layouts and fit for new server formats) may need to evolve;
- Sites will have to deploy hardware according to performance trends – presently these are driven by accelerator add-on boards; system-on-a-chip (more integrated memory, graphics and CPU) and micro-servers (cost efficiently). Product heterogeneity is increasing;
- CPU performance gains have been coming from increasing cores per CPU and so the ratio of i/o streams per disk spindle is increasing and may soon lead to congestion. This may require more disks to be deployed to sustain i/o rates;
- Projections and experiment computing model evolution expects evolution in networking capability (e.g. T1s and large T2s at 100Gb/s);
- Non-negligible CPU consumption increases have been seen with Monte-Carlo

event generation of hard processes. More efficient access to and use of multiple cores to process a single event is needed;
- The drive for simplification of deployment mechanisms (which reduce costs) will have to balance against machine limitations in areas such as the total memory required in each machine, the number of open files and/or database connections needed, the number of independent (and incoherent) i/o streams, the number of jobs handled by batch schedulers and so on;
- Both the capabilities and support outlook for batch systems is becoming a concern.

The deployment area will need to move with project technical developments. Already we can observe changes that are creating:

- A move away from resource brokering to pilot job frameworks;
- A desire to redevelop job traceability implementations;
- A drive for direct submission to VMs via cloud management software (i.e. remove the CE layer and interface to cloud API directly) ushering in private clouds (and leveraging wider support communities) and cloud bursting opportunities;
- An expectation that some sites will run hybrid cloud/batch systems;
- Use of stateless data services for smaller sites (via squids) – with the aim of removing complex storage management software that needs installing and managing;
- A further push for data federation and remote data access to optimize storage deployment (jobs can query data over the WAN). Data being placed according to usage and demand;
- Moves to off-the-shelf solutions utilizing common protocols for large-scale data movement (moving away from gridftp towards http) thus reducing the need for dedicated effort and simplifying deployment;
- Improved (common) software distribution mechanisms (e.g. CVMFS);
- A simplified information system.

As these lists demonstrate there will be substantial evolution of site hardware, underlying fabric and infrastructure, the deployed (experiment and middleware) software and the way sites are used (with a breaking down of the tiered structure) during GridPP5. The continued need for a skilled team who can deploy, operate and support changes to the WLCG/GridPP infrastructure is clear.

2.3 Operations

'Operations' describes a wide variety of work that ensures the deployed distributed infrastructure is usable and functioning optimally. This is achieved through various levels of monitoring, tracking of problems and pre-empting issues. Operations deals not just with picking up on and addressing problems, but also with making sure that mechanisms for quick detection are in place and that infrastructure wide concerns, either actual or potential, are dealt with effectively. The operations work within GridPP coordinates across all areas to ensure that smooth running and a positive user experience is maintained as much as is possible.

The day-to-day monitoring work is based around open-source products that have steadily evolved from earlier phases of the Grid project such as comprehensive Nagios probe-based testing. These probes feed data into multiple dashboards – to give a view on site availability and reliability from an operational perspective and also that of the experiment/users. At the time of writing a consolidation task is underway within WLCG to simplify the monitoring environment and automate as much as possible. Inevitably in a project as large and international as WLCG (and EGI) there are constant changes required at sites and it is the job of the operations team within GridPP to not only distil and disseminate the information efficiently, but also to form a UK view on deployment and operational matters and thus influence the direction and strategy across WLCG (and EGI). This applies from the small (e.g. what protocols to use for storage access) to the large (e.g. when to switch to a new Operating System or adopt a new cryptographic hash function) scale. It is also the job of the operations area to remain aware of matters affecting competition for resources (i.e. massive increases in demand) and to guide non-LHC communities in approaches they might use to mitigate impacts on their own deadlines. When major interventions are required, the operations team liaise with sites to plan and schedule work to minimise disruption to users.

A critical component of operations is being alert to, and developing responses to, security vulnerabilities and incidents that put our resources at risk. These responses must be coordinated with partner projects and this makes security policy a vital and key component of operations work. At a time when major changes are occurring to the infrastructure with the integration of new computing approaches (both academic and commercial such as cloud) it is essential that GridPP continues to get this area right.

2.4 Operations considerations for GridPP5

Most of the points mentioned under 'Deployment considerations for GridPP5' have direct repercussions for the way in which the GridPP infrastructure will be operated. The need for continued effort in all the core task areas as developed in GridPP4 is evident. WLCG recognised this with the establishment in 2012 of a team drawn from across the WLCG countries to provide necessary operational direction. The GridPP team have been significant contributors to the 'WLCG Operations Coordination Team', and have actively led several of the task forces it has established (e.g. for CVFMS migration and multi-core workflow testing). In GridPP5 we must continue to actively contribute to the WLCG wide activity where possible.

Operational tasks are likely to increase during Run-2. Here is a list of some areas where we can already see new demands:

- The coordination of integration activities; overseeing the deployment and testing of new and updated services; coordination across multiple infrastructures;
- Absorption of some EGI tasks (this depending on EGI project evolution);
- Managing fallout from the gradual move away from strict tier roles of the Monarch model – as there is divergence in the capabilities needed across sites;
- Increased running of data federations (services), and ensuring better data placement and dynamic caching achieves its goals (e.g. monitoring 'popularity' indexes and subsequent loads on resources);

- Optimisation of services with the more closely integrated experiment software stacks;
- Managing and responding to overspill (e.g. prompt data reconstruction migrating to T1s in event of congestion T0);
- Expanding the role for Tier-2s as experiment workflows evolve – better optimization of load and throughput (via job mix) will be needed;
- Handling the refined granularity of production and analysis payloads;
- Increased processing times and memory footprints (higher energy and pile up) impacting site efficiency;
- Extension of the CPU accounting portal to storage accounting – allowing better utilization of disk (for example through 'active data' lookups);
- New approaches to data caching. Many WLCG major service incidents have surrounded Oracle databases. Reducing the complexity and demands in this area will make operations easier. For example sharing conditions data via caching mechanisms (CVMFS and Frontier);
- Consolidation of file catalogues;
- Increased use of software-defined networking;
- New security risks and the impacts of particular security measures and developments.

During GridPP5 we will see start to see the benefits of improved distributed computing methods and consolidated monitoring. However, reductions in operational need here are likely to be offset by new needs introduced with new hardware and technology approaches, higher site capacity targets and increased throughput demands.

2.5 Support

Support is the last area to be described but is a key function of any successful production service. It can be sub-divided into operational support, deployment support and user/VO support. All of these are dealt with by way of helpdesk tickets that are triaged and assigned to people within the project for follow-up. The tickets can come from anyone who either needs assistance or spots a problem – for example from confused users, experiment operations teams, the GridPP operations on-duty team or from system administrators. GridPP is part of a production infrastructure and it is expected that tickets be acknowledged within a few hours (during the working day) and usually that progress/resolution is demonstrated quickly. Tickets raised need to be tracked and for this we rely on the Global Grid User Support (GGUS) facility in Germany – GridPP members have played a significant role in developing the functionality, processes and procedures used by GGUS. When ticket progress is not apparent or rather slow then the support staff escalate the matter as appropriate or offer suggestions.

Operational and deployment support is usually handled by the GridPP Tier-1 and 2 site staff or by the GridPP operations team or area experts. Problems are typically due to configuration issues, fabric failures (including machine room failures) or represent requests for the site to implement a change. Most tickets are resolved quickly but some are manifestations of cross-site issues and take longer to resolve. Where progress cannot be made the team either bring in international support or reassign the ticket to people who can investigate more deeply – for example the developer of a specific product. Given the critical role of data storage and management for WLCG, GridPP established a storage group and members of this resolve a good deal of the tickets raised in the UK.

User support requests may come directly from individual users or represent a community/experiment issue that needs further investigation. Often the problem relates to actual grid job failures (something may have timed out or a certificate may have expired), but for a new Grid community the request may be very generic and seek advice on the best strategy for their users to submit jobs and manage data or represent a request for their enablement within VOMS, on a central service or at sites. As can be seen in Figure 1, the weekly number of new support tickets being created over time (December 2011 to February 2014) for GridPP sites and experts to deal with fluctuates over the whole period between about 10 and 40.

With decreasing levels of support from EMI product teams it has been necessary to pick up additional responsibilities for the support and maintenance of certain middleware components upon which our sites rely. Tarball packaging for User Interfaces and Worker Nodes are two recent examples, along with increased support of the Disk Pool Manager Storage Element as part of the DPM collaboration.

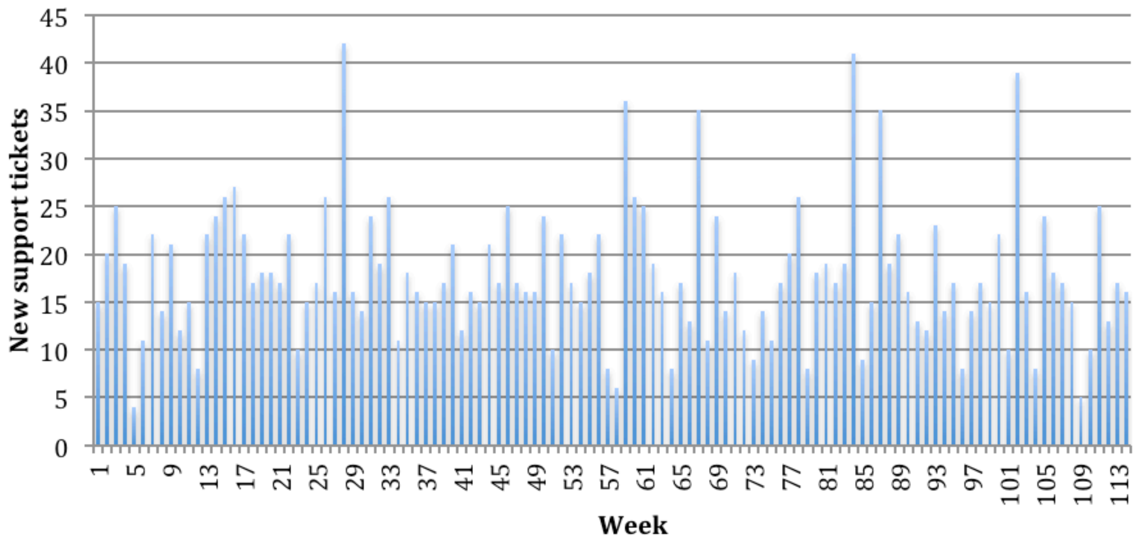

Figure 1: The number of new UK directed support tickets opened each week in GGUS between 1st December 2011 and 1st February 2014. Weeks with the highest numbers show deployment and operations campaigns. The lowest numbers are seen around the Christmas and New Year holidays. The whole service is constantly evolving which is why there is no clear downward trend over time.

2.6 Support considerations for GridPP5

Resources in the future will be increasingly heterogeneous and accessible with new interfaces – from direct submission via cloud APIs to data transfers using newly implemented protocols such as http. Whilst the evolution will lead to an overall simplification of the infrastructure and the ability to draw upon much wider support communities, the era of stateless sites and self-healing monitoring based solutions is still some way off. Support in all its forms remains essential. This support will be needed to help sites deploy and investigate issues particularly related to data management and storage. Although GridPP's main 'users' are the experiments, there will remain a close connection between GridPP site staff and local experiment staff and between experiments and sites hosting specific physics groups, group and individual Grid analysis work and outputs. Security will remain high on the list of areas where support must not be compromised.

As Figure 1 shows, even as the infrastructure has improved the support load has not dropped substantially. This largely reflects the fact that we do not reach a steady state with the infrastructure, as there is continuous change as things develop to become more capable. Thus, although complexity (and thus maintenance costs) will likely decrease as open-source cloud software replaces parts of the bespoke Grid software, and the experiment software stacks and submission frameworks converge (reducing the diversity of support needed), this is unlikely to translate into a direct reduction in the support activity.

Support requests could actually increase; for example if the wealth of experience developed in producing this environment (that allows distributed access to resources from large communities of users) is called upon by other disciplines; especially if there is also a shift towards the consolidation of resource provision across sciences or between different types of computing need. We have seen significant advantages in having compute clusters and associated support located close to users and this has been an important factor in building trust within the LHC user community. Maintaining current support effort is essential for continued stable operation of the infrastructure and may also become key to allowing any required evolution to shared resource provisioning.

3. GridPP task areas

The majority of deployment, operations and support efforts within GridPP come under a core team within which tasks are allocated.

3.1 The core tasks

The core Operations Team is led by the GridPP **Production manager** (who reports to the Project Management Board) supported by 8 **Tier-2 site experts** (one from each of the sites supporting all experiment activities). In addition there are several posts part funded by EGI that contribute to the core team.

The **core team** tasks have been distilled as:

Staged rollout – this activity, to which we contribute, has been coordinated within EGI but with the end of EMI there has been an increasing move towards using standard tools and relying on Extra Packages for Enterprise Linux (EPEL) repositories. The current WLCG approach is being developed as part of a 'middleware readiness' task force. GridPP is involved in this task force and will be expected to contribute effort (which includes allocating hardware for testing, setting up dedicated queues and documenting findings) for testing middleware in due course. This area is increasingly used for integration testing of components. Software is continually evolving due to functionality changes and patching of bugs. The effort required for this task will increase in GridPP5 as products become less centrally coordinated and new approaches (such as cloud provisioning) make it to production.

On-duty – this is an essential task and consists of one person (rotated on a weekly basis) from the operations team (and more recently beyond) undertaking a series of tasks each day – including reviewing dashboards, raising tickets against observed problems, and following up on tickets.

Ticket follow-up – there are typically 25-50 live tickets at any one time across GridPP sites. These vary from user support requests to requested infrastructure updates and changes. These tickets are reviewed regularly by a member of the operations team, and at a national level on a weekly basis – during which issues arising more widely are discussed. One member of the operations team takes responsibility for reviewing open tickets but the team as a whole contributes ideas to resolving more difficult problems.

Regional tools – Site staff rely upon a number of tools and services run to help diagnose problems. One member of the team takes responsibility to ensure that the monitoring framework is running and kept up-to-date (for example the Regional Nagios service run at Oxford with a backup at Lancaster). They act as our representative within international discussions aimed at evolving the tools. The ops-team also work to provision tools to assist site staff and users in their work – by for example providing better configuration management and checking tools.

Security – Maintaining a secure infrastructure is critical and a team of several core-ops members support the GridPP security officer in their role. The team meet on a bi-weekly basis to agree strategies for the UK as a whole to follow (for example in coordinating rollout of authentication services), to review and follow-up on incidents, to form a view on vulnerabilities and discuss and recommend updates to security policies. This security activity needs to be carried out in an international context. The UK security officer is required to perform regular weekly duties as part of an operational rota at the EGI Computer Security Incident Response Team (CSIRT) level, this CSIRT being the operational security team used by WLCG in Europe. In the event of an incident the team will ensure a coordination structure is in place and run effectively. Other tasks include checking site patching and reviewing automated security test results.

Documentation – Large quickly evolving projects produce varying qualities of documentation that can quickly become obsolete and this impacts the overall efficiency of staff and the level of service that can be provided. The ops-team assign responsibilities for documents across its members and also provide tools to encourage regular review of critical documentation. One member of the team regularly reviews documents and follows up on missing or out-of-date areas. The GridPP website is maintained under this area.

Monitoring – During the last 6 years there has been an explosion of monitoring solutions and the ops-team take part in feeding back to the developers with regards to changes and new features, ensuring common understanding of results and so on. In recent months the emphasis has moved to improved consumption of the data and therefore distributable scripts have been developed to pull data into dashboards. GridPP ops members have played a role in this work and also now take a part in a WLCG monitoring consolidation task force aimed at simplifying the monitoring infrastructure.

Accounting – For the projects (GridPP, WLCG and EGI), infrastructure management and users it is important to be able to track resource usage. CPU usage (and increasingly storage) data is published centrally. It is not trivial to keep the published data accurate as much depends on benchmarking the performance of hardware and ensuring the data flows are working. Therefore one of the core tasks concerns the benchmarking of resources and following and contributing to developments in this area. As there is a close association with the metrics used by GridPP for funding allocations, this task also aims to keep that metric information up-to-date.

Core services – There are a number of national services that the project must run for the infrastructure to function. These services are distributed across the sites to which the ops team members belong and cover such things as network monitoring (perfSONAR), VO Management Servers (VOMS) and Workload Management Servers (WMSs). The effort to deploy and run these services, to understand the results and also contribute requirements and expertise to the developers comes from several core operations team and Tier-1 staff members.

Interoperation – GridPP sites rely upon middleware developed within various international projects and within the experiments; they also contribute to larger infrastructures such as the European Grid Infrastructure (EGI). As such it is important that GridPP remains fully exposed to and contributes input and direction to those projects. Therefore a core task role with responsibility to represent our sites in these contexts exists. Typically this requires attending several meetings a week and coordinating UK participation. Within the UK we also have interoperation needs (as happened with the NGS and will future closer working with other infrastructures such as DiRAC).

Wider VOs – GridPP endeavours to support other communities outside of its core LHC VO user communities. These range from HEP organisations to other academic disciplines through to industrial partnerships. Introducing these communities and addressing their needs, not to mention actually enabling access at sites and core services is not a negligible undertaking. Therefore, a core task exists to explicitly drive the liaison with these communities and follow-up on issues they are encountering. We also investigate and help VOs test tools and develop strategies that may reduce their support loads (for example by running a job submission framework - DIRAC).

Networking – The distributed data distribution used for WLCG computing coupled with the high i/o demands of many tasks has led to a constant need to review site LAN capability and configurations, institute connectivity to JANET and wider considerations about network peering. Work done within this area extends beyond the core-ops team and covers areas such as network forward looks, network monitoring and the evaluation of technologies such as IPv6. Better than expected network connectivity is one of the driving factors behind improvements in the experiment computing models and networking as a core operational area is of increasing importance. The area task has for much of GridPP4 been covered within the ops-team under Core Services.

In GridPP4 a lead for each task was identified from amongst the Tier-2 site experts funded at the sites the experiments deemed more critical for their operations. These sites received an additional staff member who in addition to contributing to core operational tasks had a responsibility to interact more closely with experiment computing work. There was also effort part funded by EGI - posts funded in this way have mostly contributed to the tasks that we run in conjunction with EGI (e.g. on-duty, staged rollout and middleware testing/tracking).

In addition to the core team tasks, there are a number of dedicated tasks undertaken in other well-defined areas but overlapping with the main core areas of operations and deployment. These are explained next.

3.2 Data storage and management

The role of the storage and data management team is to support the infrastructure for storing and moving data in the UK. Coordinated via a weekly meeting, the team liaise with system administrators at Tier-2 sites, to provide support for the sites' storage services. This support is not just troubleshooting, but also testing new releases to ensure upgrades do not affect the services, and to get the most cost-effective performance out of the transfers.

Staff at the sites generally have a lot of expertise in running data services: they can install and upgrade head nodes and pool nodes, handle procurement and monitoring. Where the expertise of the storage team is needed, site staff will email the *storage list* and request assistance with a problem – conversely, the support team will advise people with respect to things that they cannot do themselves. This support model also means that it is possible to run (small) GridPP sites with as little as one person, as individual site staff do not need to know everything.

GridPP's storage and data management group has shown some remarkable successes during GridPP4. Not only has the team been successful in supporting a "Big Data" infrastructure which among other things enabled the discovery of the Higgs boson, they have also led the WLCG storage "evolution" task force, they have presented work at HEPiX and CHEP (as well as to other conferences including organising several cross discipline and industry 'Big Data' related workshops) and established an enviable reputation within WLCG for expertise. GridPP's storage and data management support list now contains approximately 80 Tier-2 site members in 10 different countries

The core team is split across three posts, whose duties, in addition to the core duties of supporting storage at sites and responding to incidents, are:

1. Data storage support:
- Investigating and supporting filesystems;
- User community support for non-LHC communities;
- Next generation data storage – looking at what is the most cost effective way to deliver storage of the required performance and data "durability".

2. DPM support:
- General Disk Pool Manager (DPM) support (DPM is the most popular Grid storage middleware in the UK. As a part of an international agreement to support DPM the UK must commit a certain level of effort;
- Toolkit support: in addition to general DPM support, the UK has contributed the "toolkit," an administrative tool for performing maintenance and troubleshooting of a DPM installation;
- Liaising with the international DPM support team;
- Liaising with WLCG's storage teams, including reporting to the GDB (Grid deployment board, WLCG's steering group).

3. Data transfer support, Tier-2 coordination, and VO liaison:
- Optimising and tuning of GridPP transfers and network stack;
- Monitoring of the data traffic;
- Supporting of data transfers from Tier-2s to the RAL Tier-1, or to Tier-1s or other Tier-2s outside the UK.

In addition, all three team members are expected to disseminate their work, both in peer reviewed conferences as well as by less formal means such as presentations and blog posts.

The core team focuses on DPM support (19 instances across 13 sites in GridPP.) Currently there is no dedicated effort to support StoRM (7 instances across 5 sites) and dCache (2 instances across 2 sites). Appendix B provides an indication of areas where Data Management is expected to develop during the period of GridPP5.

3.3 Security areas

A **Security Officer** post exists within GridPP to ensure sufficient expertise is available to advise upon/manage the various security risks involved and to maintain the availability of Grid services. The GridPP security officer is required to participate actively in the EGI CSIRT, this being the only operational security team used by WLCG in Europe. This work includes regular weekly operational duties on a rota within the EGI CSIRT. The tasks include the investigation of security related issues and taking a lead in formulating the GridPP/WLCG/EGI strategy in response to them. The role also includes providing advice on training and dissemination of security best practice. Some of the day-to-day responsibilities, while recruitment of a new security officer takes place, have been distributed amongst the GridPP security team but much of the required work has been put on hold, such as assessments of cloud security models, provision of training events and most importantly the UK has not been participating in the EGI CSIRT duty rota and our international colleagues are complaining. Trying to cope with no Security Officer in post (during the last few months) has taught us that if possible it is better to spread the single role over more than one staff member and we realize that the GridPP security team is very beneficial for handling major incidents, with or without a security officer.

Security vulnerability handling. Software vulnerabilities represent a major weakness to our services and have the potential to result in security incidents and service failure if they are not handled properly and fixed. It is therefore important for any distributed computing/data infrastructure to have a team and a process whose role is to coordinate the reviews of identified vulnerabilities and ensure tracking of problems and developer responses to them. Clear distributable alerts following risk assessments are also key to effective blocking of exposed/known vulnerabilities. This role was first created by GridPP as a major contribution to WLCG/EGEE and EGI. During GridPP4 this role has been funded by EGI.eu/EGI-InSPIRE. We perform the role primarily as a UK responsibility for the benefit of WLCG. It is not yet clear how best to fund this important role during GridPP5, but the tasks will continue to be required for the secure operation of WLCG and GridPP.

International Security Coordination – Distributed computing infrastructures rely on trust between institutes, countries, providers, users and developers being maintained. This trust is established and maintained by well thought out security policy and by appropriate security operations and procedures. As the infrastructure evolves and new technologies are deployed it is therefore essential to further develop, evolve and disseminate security policy in conjunction with national and international partners. GridPP has been a leader in this role for many years. All of the current security teams used by WLCG (security operations, vulnerability handling, policy coordination, IGTF/EUGridPMA for Authentication) were proposed and started by the GridPP security team and this has allowed us at many levels to give a firm direction to it.

Security Task evolution in GridPP5

All of these security tasks will need to continue in GridPP5. The GridPP team wishes to continue providing these services to WLCG. Particular challenges will arise from the deployment of new technologies, such as federated cloud services, the move towards more use of federated identity management, varying levels of assurance in authentication, aggregation of national identity attributes with community-based authorisation attributes and roles. All of these will need changes to security best practice, policies and procedures.

4. Interfaces to the LHC experiments

4.1 Meeting structures

As explained earlier, GridPP4 funded 5 ATLAS and 3 CMS sites to have site staff who would work closely with (and have involvement in) experiment operations. LHCb had less of a requirement on Tier-2s, nevertheless there is an overlap with core-ops members.

There is a weekly GridPP operations meeting that involves representatives of the LHC experiments where problems and issues are discussed in a more general environment. This is consolidated by more in-depth experiment specific UK operations meetings run by the experiment operations representatives. WLCG also runs a short daily operations review meeting, which has both core experiment and GridPP operations team participation. The LHC experiments also have members within the WLCG operations coordination team where GridPP operations team members are also present.

GridPP and the experiments both have operations representation at the GridPP PMB and the WLCG Grid Deployment Board. Issues are also discussed in a joint context within the WLCG Management Board.

4.2 LHC Experiment support posts

As GridPP operates a computing infrastructure our ‘users’ are, to first order, the experiments rather than individuals; that is, GridPP supports a range of workflows and services for each experiment. GridPP has three experiment support posts (one for each of ATLAS, CMS and LHCb) based at the Tier-1 that provide operational support and maintain a formal channel of direct communication with the experiments. The post holders have an expert understanding of the individual experiment computing models and workflows from an implementation perspective. These have become vital for obtaining the most out of the infrastructure for the experiments. The posts have ensured excellent engagement with the central LHC experiment computing coordination activities and have led UK participation in pilot projects to ensure effective strategies are pursued. They been core to developing Tier-1 strategies in light of frequent changes to the operational models and have assisted with the UK wide needs of the experiments. With their good knowledge of user issues and close association with GridPP operations and deployment activities, these roles have become valuable in communicating issues back to the experiments and seeking their resolution.

In addition to outward facing responsibilities, the support posts have acted as RAL Tier-1 liaison for ATLAS, CMS and LHCb – this has ensured an effective communication chain between core experiment operations and the RAL Tier-1 team. In addition they have assisted in Tier-1 activities that have directly benefitted the experiments – such as with future batch system considerations.

As we move towards Run-2, and as the computing models put increasing emphasis on the role of Tier-2 sites, these experiment support roles must evolve to increasingly drive engagement across the Tier-2s. Close working with the Tier-1 deployment and operations team has endowed the post holders with a wealth of experience that can directly benefit the TIer-2 sites as those sites absorb more responsibilities.

4.3 Other specific support

The project part funds two **GANGA support** posts. These posts have provided ATLAS and LHCb with the necessary effort to ensure Ganga bugs are fixed and interfaces remain working. As well as taking responsibility for Ganga Core they have also consolidated and run the necessary support infrastructure for ongoing Ganga development – which in includes making it more usable by non-LHC communities.

5. Support for non-LHC VOs

GridPP has enabled support for over 20 VOs on the VO Management Server run at Manchester (and now also hosted on redundant 'core' servers at Oxford and Imperial College). This includes VOs that we agreed to support following the demise of the NGS and GridIreland. The VOs[1] cover non-LHC HEP, physics and non-physics areas. These are supported at sites according to demand. In addition a number of international VOs hosted on VOMS overseas are enabled[2] at GridPP sites.

In GridPP there has been a **User Technical Support** post to aid new VOs in setting up their job submission frameworks and develop a processing strategy. Responsibility has been taken for setting up a DIRAC instance for the use of non-LHC VOs – so far NA62, SNO+ and the landslides VOs have used it. The more complicated case for T2K is currently being worked upon. There is an increasing need for such support if the project aims are wider than supporting LHC activities. The project must assume that any community wishing to make use of Grid resources has manpower available to port applications to the Grid and to invest in developing and supporting approaches for job submission and data management.

6. National context

GridPP makes up a significant part of the UK NGI and runs the services on which it depends. This includes maintaining a VOMS service and hosting WMSes and the UK top-level BDIIs.

Some other areas are joint funded international (EGI) obligations and include:

i) **GOCDB maintenance** Developing and maintaining the Grid Operations Centre DataBase (used by all of WLCG and EGI sites). GOCDB provides an essential repository of relatively 'static' information that helps to identify sites, their contacts and services. It is used for publishing outage information.

ii) **APEL maintenance**: the Account Processing Events Log (APEL) accounting system is an international service that collects job accounting data. This data is used in multiple ways from checking how resources are being used and

[1] See https://www.gridpp.ac.uk/wiki/GridPP_approved_VOs for more details.
[2] For a matrix of VOs supported across sites see http://pprc.qmul.ac.uk/~walker/votable.html.

shared to comparing site contributions. With millions of jobs run daily on the infrastructure processing the associated data is a full-time task.

For completeness mention should be made of areas where GridPP has some dependencies but which are funded outside the project:

i) The UK **Certificate Authority** (CA) is operated by STFC and maintains a signing service that acts as part of a global trust federation to authenticate UK users.

ii) The **UK helpdesk** is a service is part of RAL/Daresbury Scientific Computing Department run to provide help to UK users in areas under the CA and e-Science supported applications/communities.

7. International dimensions

The GridPP infrastructure forms part of the WLCG production grid. As such the hardware and middleware services must be maintained to levels agreed in the WLCG MoU.

The UK NGI also forms part of the EGI infrastructure and contributes essential components to EGI such as the GOCDB and APEL services mentioned in the last section. EGI provides some project funding for these.

In addition, GridPP takes a lead in Security Policy, Trust and general coordination work in both EGI and WLCG. Maintaining trust and agreed methods across infrastructures is essential. Recent work has included extending policies to encompass new resource types such as cloud. There is increasing recognition of the need for federated identity management and single sign on services, which needs further development.

In addition to these project leading contributions, GridPP contributes at a number of other levels to ongoing work in both WLCG and EGI. One of the most obvious is allocating effort to the cyclical task of middleware /product testing as part of a well defined ‘staged rollout process’. This is evolving into engagement with a WLCG initiative on middleware readiness that now forms one task force within the WLCG Operations Coordination Team work.

Other activities falling under this WLCG team and in which GridPP staff are involved include: glexec deployment; perfSONAR deployment; tracking tools evolution; HTTP proxy discovery; FTS3 integration and deployment; XrootD deployment; SHA-2 migration; machine/job features; IPv6 validation and deployment; multi-core deployment. GridPP had strong involvement and leadership in several other task forces that have now completed including SL6 migration; CVMFS deployment and Squid monitoring.

8. Distribution of effort

The roles described in this paper have been allocated across the GridPP institutes according to needs and skills needed at and existing at those institutes. Table-2 shows the distribution and FTEs in each area as at January 2014 (note these figures have evolved during GridPP4).

	Institute	ATLAS group analysis priority	CMS Group analysis priority	ALL VOs unscheduled analysis + simulation	EGI	Experiment support	*General support & managem ent*	*Storage and DM*	*Security*	*GOC (GCDOB & APEL*
London	**Brunel**		1.4							
	Imperial		2			0.755	*1*			
	QMUL	2								
	RHUL			1						
	UCL									
NorthGrid	**Lancaster**	2								
	Liverpool			1						
	Manchester	2								
	Sheffield			0.5						
ScotGrid	**Durham**			0.25						
	Edinburgh			0.5				1		
	Glasgow	2					*1*	1.5		
SouthGrid	**Birmingham**			0.5		0.5				
	Bristol									
	Cambridge			0.5			1			
	Oxford	1.5								
	RAL-PPD		1.5						1.5	
RAL	**RAL SCD**									1
	RAL general					3.5		1		
TOTAL										

Table 2: The allocation of deployment, operations and support effort across the GridPP sites.

9. Appendix A – Roles, tasks and responsibilities

To provide an overview of the important work done in different categories the following table provides a summary of the areas, roles and tasks covering all areas of deployment, operations and support work within GridPP.

Area	Role	Funded FTE	Role description
Management	Production manager	1	
	Represent GridPP deployment and operations on the WLCG Grid Deployment Board, EGI Operations Management Board and within the UK NGI. Ensure site support staff remains aware of (and have input to) project wide developments: news; decisions; milestones; targets; strategy. Take up and resolve site issues with wider project(s). Coordinate the operations team – which means overseeing each of the core operations areas (making sure there is an on-duty person, middleware is tested, updates are performed in a timely fashion, performance issues are understood and acted upon). Follow up on escalated user and site issues. Manage UK authorisations. Administer various mailing lists and CERN eGroups. Contribute to WLCG ops coordination and decisions. Participate in conferences and workshops to exchange ideas, issues, solutions with peer projects both inside and outside HEP. Keep the GridPP PMB informed about deployment, operations and support issues. Attend daily and weekly operations meeting with international partners and follow-up on issues. Compile quarterly reports.		
Site support	Site expert	11	
	Manage site level operations. Procure and test hardware. Deploy and maintain hardware. Deploy and maintain cluster management tools. Deploy and update/patch operating systems and Grid middleware. Support local users. Respond to experiment/VO tickets and requests.		
Core operations	Group analysis expert	8.5	
	Undertake Site Expert duties at sites plus: take responsibility for core operations tasks, liaise with experiment computing and operations teams and users to develop and maintain the service. Remain aware of technical developments. Participate in testing work and investigate opportunities for improvements (e.g. with monitoring tools).		
	Tier-2 support	2	
	Undertake the Core Ops tasks without Site Expert responsibilities plus: Coordinate activities on a more local level. Offer second level support to sites. Participate in forward planning and testing activities (e.g. adoption of IPv6, network optimisation and monitoring).		

Data Storage & Management	Data Management	1	
	Contribute to the DPM collaboration (support, maintenance, testing, feature enhancement). Provide site and user support for DPM. Responsible for the DPM Admin Toolkit. Support sites with storage configuration and management of storage elements – to ensure efficient local file access. Evaluate new storage technologies. 'Small' VO Data Management support.		
	Data storage - DPM	1	
	Contribute to the DPM collaboration (support, maintenance, testing, feature enhancement). Provide site and user support for DPM. ATLAS T2 data roles. Liaison and expert for ATLAS data (e.g. moving forwards with data redirection). Advise on storage configuration. VO Data Analysis support. Testing of new protocols and approaches. Tracking of upgrades. Liaising with WLCG's storage teams, including reporting to the GDB (Grid deployment board, WLCG's steering group)		
	Data storage - other	1	
	Review file catalogue integrity for VOs. Run data checks (e.g for 'dark data'). Provide Tier-1 liaison for data management. FTS configuration and management to optimise experiment data flows (to Tier-1s and Tier-2s).		
	Data processing	0.5	
	Provide support on data processing configurations and efficiency options. Investigate approaches and opportunities to improve performance (e.g. with file data structures).		
Networking	Oversight	0	
	Advise on issues arising. Liaise with JANET. Develop strategy. Support GridPP forward look and international peering approaches.		
Security	Operational	1	
	Provide expert support and advice to the GridPP Operations Team and the UK Tier-1/Tier-2 sites in all aspects of security, with the aim of achieving secure and reliable operation of the GridPP computing Grid. Participate in the EGI CSIRT including regular weekly duties on a rota. Run security courses. Coordinate GridPP response to security incidents. Develop and maintain incident response procedures. Monitor security/patching dashboards. Provide expert input to technology and software evaluations. Contribute to EGI security work.		
	Policy	0.5	
	Provide expert guidance to and coordinate activities in GridPP, WLCG and EGI in developing and extending security policies and procedures. Contribute to and lead activities moving towards an integrated federated identity management system. Ensure policies ease interoperation between infrastructures and are easy to understand and adopt.		

	Vulnerabilities	0	
	Lead security risk assessment activities for the benefit of WLCG and GridPP to determine how software vulnerabilities should be handled. Develop clear and concise advisories for sites. Identify security issues and provide mitigation strategies. Coordinate the security assessment of software components.		
VO support	ATLAS	1	
	Assign ATLAS T1 disk deployments and ensure deadlines met. Manage ATLAS batch farm needs and oversee ATLAS forward testing for products such as Condor and hyper-threading. Coordinate ATLAS UK Cloud activities (including LOCALGROUPDISK management). Deal with ATLAS software support at Tier-1. Support RAL FAX and run RAL UK pilot factory for ATLAS. CVMFS support. Oversee RAL ATLAS Frontier service and coordinate Frontier within ATLAS. ATLAS IPv6 contact. Sub-convene ATLAS B-physics group to remain aware of user issues. Contribute to the ATLAS software harmonisation team. Run user work to gain insights into issues. As ATLAS computing expert contribute to shift work and operational discussions. Distributed Analysis Coordinator for ATLAS.		
	CMS	1	
	Oversee CMS utilised services at Tier-1 – FTS, myProxy, ARC CE and batch. Provide support to users. Maintain the services. Contribute to on-call cover. Maintain accurate capacity and usage information including monthly GridPP and WLCG resource accounting. Act as CMS RAL Tier-1 contact (i.e. main point of liaison). Ensure reliable and efficient handling of CMS data storage, transfer and processing activities. Maintain CMS specific configurations, services and systems. Contribute to CMS HLT farm opportunistic cloud resource activities. Act as CMS level 3 computer manager so oversee smooth processing of CMS data and MC samples across all CMS Tier-1s so that users have timely access to data for analyses.		
	LHCb	1	
	Focus on getting storage and CPU resources (Tier-1 and Tier-2) available for LHCb and work with LHCb experts to resolve issues. Attend various meetings to co-ordinate actions and stay abreast of LHCb developments. Undertake LHCb 'Grid Expert on Call' shifts. Ensuring hardware meets LHCb needs (via role within Hardware Advisory Group). Take on LHCb-UK computing coordination role. Review, develop forward look for and quarterly reporting on LHCb activities. General LHCb user support. Use full analysis chain software framework to actively spot problems.		
	Other VOs	0.225	
	Embed within small VOs to understand their computing needs and support the adaptation of their code. Provide guidance to VOs on and help them implement efficient job submission platforms. Support VO activities (e.g. MICE reconstruction and NA62 production systems)		
	ATLAS Ganga	0.5	
	Support for ATLAS use of Ganga. Fixing of bugs within Ganga Core. Assessing/meeting wider Ganga needs. Consolidation of the infrastructure (test machines, bug tracker, monitoring services) to		

	ensure that these can be maintained by the Ganga team.		
	LHCb Ganga	0.5	
	Support for LHCb use of Ganga including implementation of bulk submission to the DIRAC backend. Support of User Interface Layer on top of summary files provided by Gaudi jobs. Improve the interactive experience of Ganga.		
	CMS Castor	0.5	
	Support for CMS use of CASTOR		
GOC	GOCDB	0.5	
	Support and development of the Grid Operations Centre DataBase (GOCDB) upon which all of WLCG and EGI site information relies. Support users of the system. Respond to feature requests. Maintain production systems.		
	APEL	0.5	
	Develop, maintain and support the APEL middleware packages. Run the APEL databases for consumption of all (CPU and storage) accounting data from publishing sites. Develop new interfaces for emerging technologies.		
	CA	0	
	Run the UK e-Science Certificate Authority. Plan for hardware and software evolution. Liaise with RAs around the UK. Provide tools to ease certificate renewal.		

10. Appendix B – GridPP5 Data Management

Data management in High Energy Physics in the coming years faces a number of challenges. LHC storage requirements are expected to grow to almost 400 PB by 2017 (figure 1): an amount that is greater than the resources that could be delivered with a constant budget. Therefore this will not only be an important part of operations activity in the coming years, but also the community must find ways to use all of its storage and data resources more efficiently. A programme to achieve this within GridPP5 is outlined below.

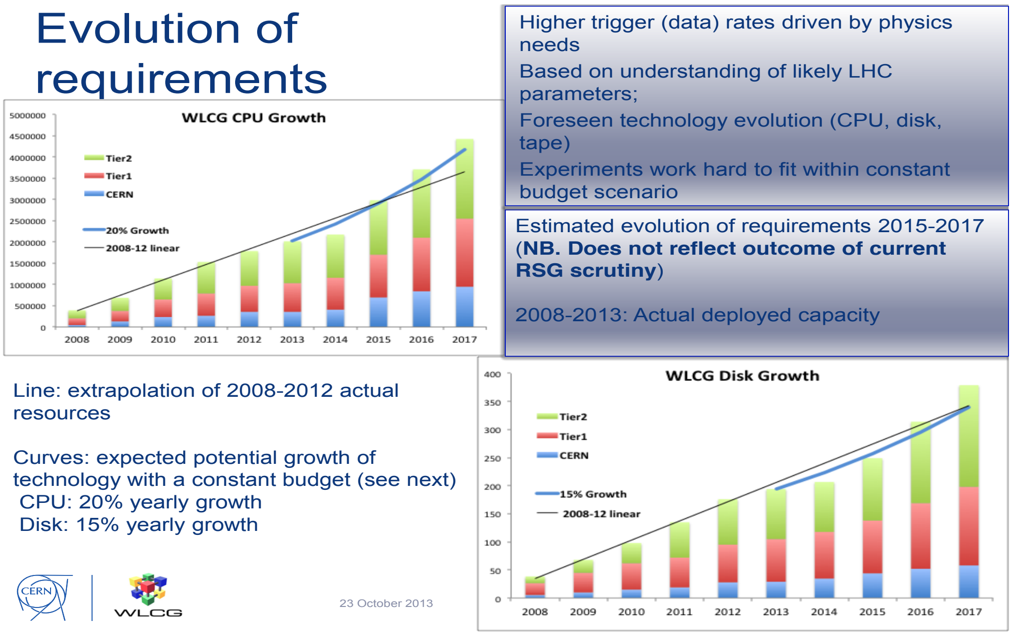

Figure 1. Projected storage requirements for the LHC (PB) [1]

10.1 Data Operations

Storage and data resources are perhaps the most challenging area of LHC computing to deliver. While manageability has improved and will continue to do so, in particular through the middleware developments described in the next section, it is expected that responsive operation of grid storage will still be an important part of GridPP5. Both site expertise and some level of “central” activity will be needed to support installing and configuring storage in terms of both the hardware and middleware (est. 0.25 FTE effort) and additionally to respond to experiment demands such as the cleaning of ‘dark’ data and the management of space allocations (0.5 FTE). The UK has gained an excellent reputation for responsiveness to the demands of LHC experiments and it will be important to maintain this to meet the physics goals of LHC Run-2.

10.2 Data Middleware

GridPP, and WLCG in general, currently relies on storage "middleware" to make its resources available. A variety of products are in use at UK Tier2s (see figure 2) with the majority of sites using the DPM storage solution, though considerable resources have also been very successfully made available with dCache and StoRM.

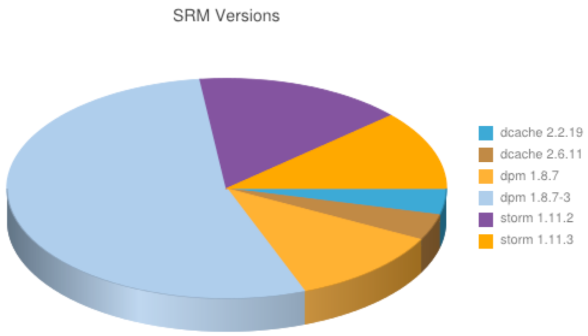

Figure 2: Storage Middleware in use at UK Tier 2s

A storage "middleware" layer of some sort will continue to be needed in GridPP5. Various activities are underway to simplify the requirements of this layer (such as deprecation of the HEP specific protocol SRM), and the UK is leading on those efforts through chairing the relevant working groups in WLCG. However some requirements will remain, including the provision of detailed accounting, the community's strong security model, as well as its challenging i/o requirements. However, underlying data storage technologies have seen considerable recent development, for example for cluster filesystems and cloud interfaces. Furthermore, to be able to scale up to larger data stores, the community must consider use of simpler storage models (such as achieved by object stores rather than posix filesystems). Therefore HEP will need to work towards being able to operate on top of such technologies with thin-layer "future-looking", middleware and access interfaces. As well as being essential for delivery, such development will reduce the future support effort needed for storage.

In particular, the DPM storage solution has undergone a recent transition to a more flexible plugin architecture ("DMLite") (see figure 3.) that allows it to interface several components including "cloud" object storage, via the S3 interface, as well as cluster filesystems such as GPFS and Hadoop/HDFS with WAN transfer protocols such as http and storage federation technology (discussed in the next section). GridPP was heavily involved in the establishment of a "DPM collaboration" in 2013 that ensures long-term support for this product. This collaboration is strongly supported by CERN, but also with commitments from France, Italy, the Czech republic and Taiwan. So far it has proved very successful and through participation in the collaboration, the UK has increased awareness and influence of a product that is key to its effective storage delivery. This effort should be continued in GridPP5 (0.75 FTE) in order to continue our current collaboration commitments that overlap closely with UK operational needs (ie. testing and monitoring, support and documentation and administrator tools) but also to enable exploration of the modern storage backends or access protocols discussed above.

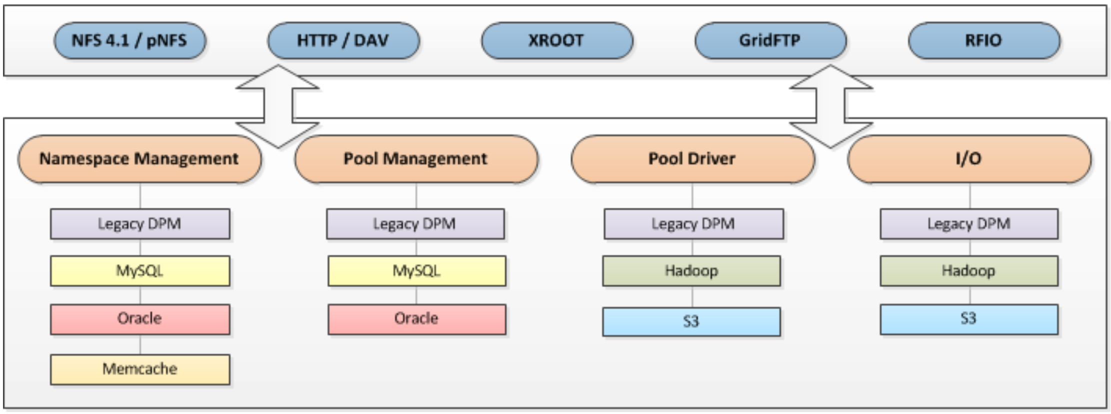

Figure 3. DPM (DMLite) plugin architecture [2]

In addition, it is expected that in GridPP5 those sites currently very successfully using dCache and StoRM are likely to continue to do so, and resources behind them will grow. This remains important because of existing expertise at those sites and in terms of diversity and avoiding technology lock-in. With the increasing demands on Tier-2s and growing deployed resources using these alternative storage solutions, some effort (0.25 FTE) will be needed in GridPP5 to support them, particularly in testing and pushing their development to suit UK requirements.

10.3 Data technology transfer

Large-scale data management and processing, often now termed "Big Data", is of increasing interest to many other academic communities which are now facing challenges that HEP has already had to deal with for some time. Additionally, "Big Data" has become business-as-usual in parts of the private sector (evidenced by the popularity of Apache Hadoop) including, but not restricted to, big names such as Yahoo, Twitter and Google. While the LHC is the oft-quoted archetypal example of big data science, the HEP community uses little of the same technology as these other communities and has not engaged with it at the level it could. In many cases, HEP has developed tools (such as those connected with ROOT) that could be a great asset, but are unknown, to industry while in other cases HEP is incapable of using complimentary industry tools. This should be addressed in GridPP5 with a programme of work to bridge the technology divide. Doing so will improve interoperation of the tools, reveal strengths and weakness in each, and enable efficiencies within all communities.

This effort needs to incorporate both technical projects, which engage with industry and other academic communities to address their challenges with existing tools, but also organizational, bringing together those involved. This activity has begun at the Big Data conference organized by GridPP in London in July 2013 but there is much more than could be achieved with focused effort (0.5 FTE).

10.4 Data interfaces and processing

While the highly distributed (grid) model employed by LHC computing has brought many benefits, it brings challenges in resource utilisation. Recently the wide-area network infrastructure has improved considerably, opening up new possibilities and interest in methods for more dynamic data and remote access: such as in so-called "Data federations". These are a mechanism to unify access to diverse storage systems using protocols such as http and xrootd [3]. They provide a uniform namespace reading data from multiple sources with transparent fallback in the event of error. We played a leading role in this area in GridPP4, successfully incorporating most UK resources into a "federated" infrastructure for both ATLAS and CMS. However this technology is still at relatively early stages and further effort will be needed in GridPP5 (~0.5 FTE) to support and evolve the required infrastructure. Doing this offers considerable potential for more resilient and flexible resource delivery.

[1] Bird, I. (2013) HEP vision and long term planning ,WLCG Workshop Copenhagen http://indico.cern.ch/getFile.py/access?contribId=9&sessionId=2&resId=1&materialId=slides&confId=251191
[2] DPM https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm
[3] Vukotic, I. *et. al* (2013) [ATLAS Collaboration] Data Federation Strategies for ATLAS using XRootD http://indico.cern.ch/getFile.py/access?contribId=265&sessionId=6&resId=1&materialId=slides&confId=214784