Practical Grid Interoperability: GridPP and the National Grid Service

Jens JENSEN, STFC
Graeme A STEWART, Univ. of Glasgow
Matthew VILJOEN, STFC
David WALLOM, Univ. of Oxford
Steven YOUNG, Univ. of Oxford

April, 2007

Abstract

In this paper, we discuss experiences with making the two large national Grids interoperable, the National Grid Service, NGS, and the UK Grid for particle physics (GridPP). We describe the current state, ongoing convergence efforts, and suggest future directions. It is our hope that the practical experience presented in this paper may be of use to other Grids embarking on interoperability exercises.

1 Introduction

The UK has two large national Grids, the National Grid Service, NGS, and GridPP. Both provide services for scientific and e-Science communities in the UK. The NGS consists of four funded sites and a growing number of partners and affiliates. It is a heterogeneous grid with different CPU architectures and operating systems (including prospective Windows sites) and partners which offer services not based around computation and storage (e.g., GridSAM and service hosting). Some of the UK supercomputing resources like EPCC and HPCx in Edinburgh and at Daresbury, and previously CSAR in Manchester, are affiliated to the NGS.

GridPP is the UK Grid for particle physics, including the UK contributions to the Large Hadron Collider Computing Grid, LCG. It is a more homogeneous Grid with the "standard" LCG operating system and a coordinated middleware upgrade cycle.

Despite the apparent differences, it is worthwhile to get the Grids to interoperate. This unfunded activity parallels the OpenGridForum "Grid Interoperability Now!" (GIN) activities, but is targeted specifically at the UK national Grids. Both GridPP and NGS participate directly in GIN activities, and interoperability experiences are of course shared with GIN when relevant.

2 Comparing GridPP and the NGS

2.1 Authentication

Both the NGS and GridPP use X.509 certificates and Globus security (GSI) which is an extension of X.509, so technically, NGS and GridPP are already interoperable here, albeit with some caveats (§2.1.2). Furthermore, the NGS has made significant advances in this field which we cover very briefly in this paper (§ 2.1.1, 2.1.3, and 2.1.4).

2.1.1 Login Users in GridPP log into the *User Interfaces* (UIs, machines running user client software) using `ssh`. Thus, GridPP users must have both X.509 certificates and `ssh` key pairs. This is considered acceptable, since most GridPP users are familiar with both `ssh` and digital certificates.

The NGS, in contrast, uses `gsissh`, that is, `ssh` amended to accept X.509 and GSI, thus ensuring that users only need X.509. They use X.509 certificates also to log into the Grid cluster head nodes, and run jobs. The process is facilitated by a MyProxy server which is used both to get a proxy to log into the head node, and to get a proxy to run jobs (and potentially to renew the proxy if the job runs longer than the proxy). This process is further facilitated by GSI-SSHTerm, a Java `gsissh` terminal enhanced by NGS to talk to MyProxy server [11].

2.1.2 X.509 The UK e-Science Certification Authority (CA) provides "medium assurance" digital certificates to Grid and e-Science users in the UK. As the IGTF-accredited [8] CA covering the UK, it covers both user communities.

Although the decision is left to the site admin, most Grids make recommendations regarding which CAs sites should trust. As a rule both Grids trust all IGTF-accredited CAs, but GridPP also needs to trust the CA from Fermilab because LCG does, despite the fact that this CA is not yet IGTF-accredited. Likewise, the NGS trusts certain CAs accepted by the TeraGrid, such as the Pittsburgh SuperComputing Centre CA, to promote interoperability with TeraGrid.

We expect these will converge, as the US CAs get IGTF accredited. This is an on-going and time consuming process, and many of these CAs have for historical reasons different practices from the CAs in the rest of the Grid.

Another practical difference is that GridPP has a "catch-all" CA via LCG (namely, the CERN CA). This means that a user who does not have a "local" CA can get a certificate via this CA. NGS does not have such a facility, although could perhaps use EGEE's catch-all (CNRS France). Another possibility is to trust certain educational PKIs (Public Key Infrastructures), or commercial ones. Most NGS customers are of course covered by the UK e-Science CA, but UK scientists also collaborate with scientists from other countries who need access to UK resources.

Finally, it should be mentioned that many novice users find digital certificates "hard to learn." Usability studies show this, also the ones performed by the NGS [12], and the NGS has done work to ease this with applets and other relatively simple Java applications to ease certificate management.

2.1.3 Shibboleth - NGS Shibboleth is being deployed by JISC to replace Athens for the higher and further education institutions in the UK. If we can make use of the UK Shibboleth infrastructure for authenticating to the Grid, then the potential user base becomes much larger.

This was the aim of the ShibGrid project [13] (the sister project SHEBANGS [16] looked at embedding VOMS attributes). Briefly, users log in to a portal and generate a full certificate or GSI proxy transparently from a MyProxy server, and other tools enable credentials to be uploaded and downloaded from MyProxy.

2.1.4 Credential Conversion - NGS Instead of "converting" the internal renewable token to a Shib login (§2.1.3, we can generate certificates or GSI proxies transparently [14] using the site authentication token. Thus, the user does not even need to go through the slightly clumsy Shibboleth process of selecting their home institution and typing their username and passwords, and it works directly with X509 instead of requiring portals (this will also be addressed in Shib version 2 and in associated further middleware developments). The site needs to run a conversion CA, also based on MyProxy, so the home token never leaves the home organisation. Such a site can potentially make Grid access completely transparent for its users.

Figure 1: EGEE sites in April 2007

2.1.5 EGEE The European *Enabling Grids for eScience*[5] project, which is now in its second phase, grew out of the European Data Grid Project (EDG), which developed the first set of testbed middleware for the European grid.

EGEE has now grown to encompass more than 32 countries, many outside of Europe, and the EGEE production grid, of which GridPP forms a part, consists of well over 200 sites (see Figure 1). The total available resources consist now of more than 37,000 job slots and more than 20PB of storage resources.

With such a large production infrastructure much emphasis in EGEE is now placed on providing a production quality grid with well supported middleware. Operations therefore play a significant role in the project, with the various monitoring systems (§2.6) being actively used by teams of *Grid Operators on Duty* (GOoDs), grid experts provided by the *Regional Operations Centres* (ROCs) provided by EGEE partners. GOoD teams will raise tickets against failing sites to resolve problems.

The tickets are channeled back through the ROCs, which play the role of co-ordinating EGEE activities and operations each region. In the UK the ROC is based at RAL and runs grid services for both GridPP and NGS, e.g., the UK/I certification authority (§2.1.2).

The NGS operates a central helpdesk to which users can send tickets appertaining to problems in both the core and partner nodes. These are then passed to the individual site concerned and tracked in both locations. Therefore the user will individually interact with the member of NGS staff nearest the location of their problem.

Work has been done to make the NGS and GridPP helpdesks interoperate; however, this work is complex and well beyond the space limitations of this paper.

2.1.6 Training EGEE runs the GILDA project ([17]), which issues short lived certificates from a dedicated non-IGTF CA that can be used on a small testbed grid, run by some EGEE partners. The EGEE NA3 group at NeSC in Edinburgh run training courses for NGS as well, with certificates from another dedicated training CA, run by the NGS.

2.2 Information System

An integral component in many Grids is the Information System. It makes this information available to clients for service selection and accounting. Information servers typically publish via LDAP in the GLUE schema, which is now an OGF standard. Both GridPP and NGS use version 1.2 [4], and although 1.3 has been finalised, there are currently no plans to switch to 1.3. We will need at least a clear migration path from LCG or EGEE.

Resources provide information via "Grid Resource Information Services" (GRIS) typically running locally (on the same host), and this information is aggregated in one or more higher level servers known as Grid Index Information Services (GIIS). In NGS, there is a single top-level GIIS for the whole Grid, whereas in GridPP, larger sites often have GIIS hierarchies.

Prior to the launch of the NGS, various Information Systems were tested in the UK e-Science Level 2 Grid pre-production environment. Initially the Globus Toolkit Monitoring and Discover Service (MDS) [1] was used. However, the MDS at the time relied upon a client push model where the GIIS waited for information to be "pushed" by the individual GRISes and this proved to be unstable as the individual GRIS daemons were prone to locking, resulting in the catastrophic effect of freezing the top level GIIS. As an alternative, the Berkeley Database Information Index (BDII) [3] from LCG was tested and found to be a more stable model. Unlike the Globus MDS, BDII uses a server pull model, where the GIIS regularly polls individual GRISes for information. BDII proved to be far more reliable than the Globus MDS and more scalable.

Thanks to the common LDAP format and the GLUE schema, interoperation between Globus-based and gLite based Grids is straightforward.

2.3 Storage Element

Storage services in Grid environments are dominated by two, essentially separate, systems. These are *Storage Resource Manager* (SRM)[24] and *Storage Resource Broker* (SRB)[25]. SRM is in fact simply a control protocol which is then layered by sites on top of existing storage, catalog and metadata systems[23], whereas SRB is a more all-encompassing single solution for data management.

There is little current interoperability between SRM and SRB; this means that storage services are essentially divided into two islands. However, both GridPP and the NGS are active in the GIN forums and track possible developments. Previously it has been suggested to build SRB interfaces to SRM, and vice versa, but we have proposed to transfer files between them instead. This could be done by a higher level data transfer service, such as EGEE's FTS ("File Transfer Service") which can access SRB via its GridFTP interface. Investigations are ongoing, in collaboration with SDSC, SRB's parent organisation, and EGEE.

2.3.1 Common SRM Implementation: DPM The gLite Storage Element (SE) is the *Disk Pool Manager*[9] (DPM), a lightweight disk-only implementation of SRM.

The Grid Engineering Task Force [22] has evaluated DPM, and recommended that the NGS start deploying DPM as an additional storage solution. The main obstacle is that DPM (or its dependencies) are currently not very portable, and run only on Scientific Linux (§2.5).

DPM is widely deployed within GridPP, being the production SRM at 12 of the 19 GridPP Tier-2 sites. Initial deployments were on a small scale, however, as GridPP Tier-2 centres ramp up their storage in advance of LHC data, DPM is now deployed in sites managing storage of 50TB, and this will shortly rise to well over 100TB.

GridPP have also contributed to the development of DPM, and currently maintain the "generic information provider" plugin which allows DPM sites to correctly publish used and available storage into the BDII.

GridPP has also been very active in helping to test DPM to ensure that its performance matches the expectations, particularly those of the LHC VOs, who will need to analyse terabytes of data at their Tier-2 centres ([20]). This testing has shown that the read rates achieved with DPM over the LAN are good (Figure 2),

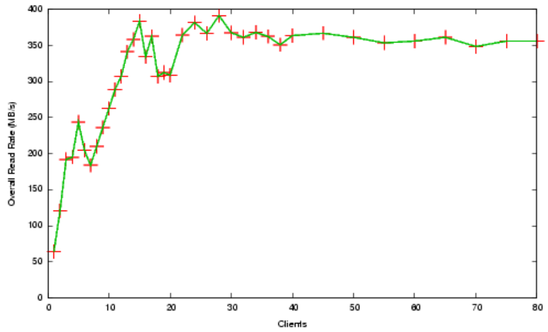

Figure 2: Total read rate with DPM

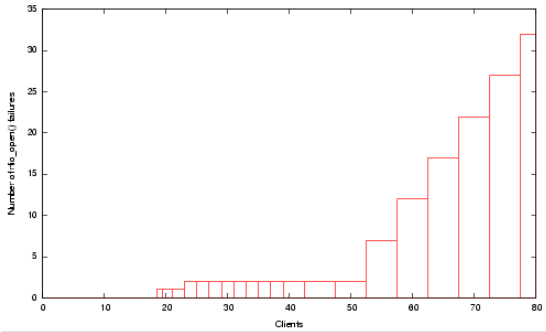

Figure 3: Open failures in rfio_open() call in DPM

but that some problems exist with trying to open large numbers of files simultaneously (Figure 3). These problems are being addressed now by the developers and tests will be repeated when a new release of DPM is available.

Very loosely speaking, GridPP tends to focus on performance, interoperability, and scalability, whereas the NGS tends to focus on maintainability, usability, and features. Nevertheless, the experience gained with DPM[9] within GridPP is used directly to help with deployments within the NGS, just as it helps other Tier 2 sites across the world. Conversely, the NGS attempt to compile DPM for Solaris can help improve the portability of DPM.

2.4 Computing Element

2.4.1 Resource Broker and Job Submission A *resource broker* (RB) presents the user with a common interface to all resources that are available to that user at any one time. The user can submit jobs to a resource broker specifying job requirements, and the resource broker makes use of the information provided by the Information System to choose the best suited resource to send the job.

The NGS is currently trialing two resource brokers, the gLite Workload Management System and Logging and Bookkeeping Service (WMS-LB) developed by EGEE [5] and the GridWay Metascheduler [6]. The latter is in its very early stages of testing, but the NGS gLite WMS-LB (hereafter known as the NGS RB) was deployed in pre-production in March 2007, and is able to send jobs to sites published in the NGS BDII server. Users can make use of the RB using the gLite client tools from UIs (§2.1.1).

Currently all core NGS sites can accept jobs from users making use of the NGS RB. Work is under way to configure the non-core NGS sites to interoperate with the NGS RB.

The fact that the NGS RB is a component of gLite means that it can easily interoperate with Grids running EGEE middleware: in principle users can submit jobs to both the NGS and GridPP from the same UI (as long as the versions are compatible), although redirecting a UI to another RB non-trivial. Only with a common RB and BDII does it make sense to use the same UI. In fact, more work was required to configure the NGS RB to work with NGS itself whose computing nodes are not based on EGEE middleware. If the RB does not submit jobs to a LCG or gLite CE, it does not fall back to a standard Globus job submission.

Transparent interoperability can be achieved on both GridPP and the NGS if, among other requirements, the following conditions are fulfilled:

- An common Information System is set up, publishing all the resources from both Grids. Presently the development NGS BDII server aggregates data from both core NGS sites and a sample subset of GridPP consisting of all GridPP resources at RAL.

- An RB is provided which uses the common information system. The NGS currently has an RB which is configured to query the development NGS BDII server described above. Both the NGS and GridPP have deployed and are using VOMS (Virtual Organization Management Service) [10] as a means of managing user authorization in the context of VOs. Using the VOMS hosted by Manchester, users from both the NGS VO and the "deployment team" DTEAM VO are currently authorised to the NGS RB.

- There are resources across both Grids that users from either Grid are authorized to use. All NGS core sites are configured to authorise members of the DTEAM and NGS

VOs. Other GridPP VOs can apply to NGS for resources. Conversely, GridPP have also agreed to authorise members of the NGS VOs to use their resources. However, the NGS VO cannot yet access GridPP resources via the NGS RB since NGS uses the new DNS style VO names which are not recognised yet by the GridPP middleware.

2.4.2 Computing Element GridPP

2.4.2 Computing Element GridPP: The most widely deployed production Computing Element (CE) in gLite is the LCG-CE. This is a slightly modified version of the Globus Toolkit 2 gatekeeper, incorporating performance and security fixes.

The default job wrapper used, `lcgpbs`, has been modified to limit resource consumption on the gatekeeper. This is done by having a single gatekeeper process monitor all of the jobs submitted through the same resource broker for a user. The `lcgpbs` job manager is also able to manage jobs on clusters where the pool accounts do not share home directory spaces (this is helpful in larger clusters where, e.g., NFS scaling problems can be seen).

However, the standard GT2 `pbs` job manager is still available. Although it is not enabled by default, a number of GridPP sites which support the NGS have enabled it to have better compatibility with NGS job submission.

This is important because the way that certain options to the `pbs` and `lcgpbs` job managers are handled is different. e.g., the `-np` option (number of processes) will fork `N` processes within a single job slot with the `lcgpbs` job manager, which is often not the desired or expected behaviour for users used to the standard GT2 job manager.

However, in the longer term, the intention is to migrate to the new gLite CE. The gLite CE is based on Condor-C, which can submit jobs to all common batch systems (Torque, LSF, SGE). X509 authentication has been added to integrate with the gLite WMS-LB system.

The gLite CE also supports BLAH (Batch Local ASCII Helper) protocol, which makes it considerably easier to pass job parameters from the resource broker into the local batch system.

Despite being released in production as part of gLite the gLite CE is not yet considered stable enough for the majority of sites to use. However, as it is intended that the next major release of the gLite middleware (compiled for Scientific Linux 4) will not include support for the LCG-CE an intensive program of robustness and scalability testing is now going on.

It would seem likely that both CEs will coexist for sometime, even if the LCG-CE is running on the older SL3 architecture.

NGS: In the NGS a standard Globus Toolkit 2 gatekeeper is deployed. The job managers vary from site to site. The list of batch systems used within the NGS include PBSPro, LSF, SGE, Condor.

Within the partners of the NGS it has been necessary to make several changes, both to how the grid middleware has been installed, either on the head node of the system or a connectivity box but also the actual software installed itself.

The connectivity boxes have been necessary particularly where we have been connecting resources which do not have good middleware support, for example Windows, SGI or Solaris systems. These connectivity systems are normally installed with a Linux variant that is Scientific Linux compatible. These are designed to be simple submission nodes for the scheduling system which is installed onto the main system as well as all the services of the middleware such as jobmanager, information server etc. This has also allowed the fears of some resource owners to be calmed. If they are unhappy with the resource utilisation of the Globus services themselves or their behaviour, for example if there was a security threat, then switching them off and removing connectivity from the cluster is merely a matter of removing a network cable.

Due to the heavy usage it has also been necessary to alter the Globus Jobmanagers. Globus 2 jobmanagers operate in a constant polling cycle which when there has been a very large number of tasks submitted can introduce a significant load onto the head node of the system, in some cases with enough submitted tasks making the system unusable. As a work around for this IO it is necessary introduce a wait period within the polling routine so that the status of the task that is connected to that instance of the running job manager is only checked once per minute. Since most tasks run within the grid operate for longer than three or four minutes then a 30 second wait was deemed acceptable. This has reduced load in some cases from approximately 75% to approximately 10% on the relevant system head node. This is an alteration that has been made to all jobmanagers for resource scheduling systems.

Other changes to jobmanagers have been made to individual systems, in particular Condor resources which have a specially altered jobmanager for situations where the pool that is being supported does not have a shared filesystem. This results in needing to alter not only the scheduler specific file 'condor.pm' but also the jobmanager.pm file to ensure that input and output files are correctly transferred between the master and workers within the pool.

2.4.3 CondorG Within the Condor ecosystem Condor-G has been designed to connect a grid resource to a standard condor Condor system [2]. This includes support for native submission with the appropriate options added to the job submission script. Condor includes a different 'universe' for submission through different mechanisms. To use grid resources the user must include the 'universe = grid' in the submission script as well type of grid software that the resource has installed. For the NGS the minimum software stack the submission script looks as below for submission to the Oxford NGS-1 node (last line broken):

```
executable = /bin/date
universe = grid
grid_type = gt2
globusscheduler = grid-compute.
 oesc.ox.ac.uk/jobmanager-pbs
```

Thus, users specify the exact resource that they wish the job to be run on. Within the Condor-G system itself though there is no automated way of discovering resources available to the user. It is though also possible to add options that can be passed through to the remote resources such as queue name etc. As shown above, submission to the NGS uses the basic options where queue name etc. are not necessary. However, when submitting to GridPP resources it is necessary to define the queue on the remote resource since Condor-G does not support VOMS additions to the user credentials.

2.5 Installed software

The method by which the nodes of the NGS (both core and partner sites) deal with installed software is very different. Within the NGS-1 core sites there was a conscious decision to use the system provided by Clustervision to load preinstalled software as modules which the user specifies before they then run remote jobs. This includes commercial software that requires specific licensing agreements. It should be noted that not all of the core sites have the same software available as independent deals have been agreed by each of the four hosting institutions.

An example of the way that licensed software has been dealt with is the Gaussian computational chemistry package[7] which has been installed on the STFC node. In this case when each new user wants to use the package they must first locally register with the STFC e-Science group so that they can add the pool account name into a physical Unix group that has permissions to access the package. This is only a local solution, though, independent for each node since there is currently no method of propagating both the software itself and the permissions/license agreements between nodes. For open source licensed software there has been an attempt to install a common set of applications across all nodes for which the particular application is supported.

Unlike the NGS as a whole, GridPP is much more homogeneous, both in its hardware infrastructure, operating systems, middleware, and applications. Most sites are running a flavour of Scientific Linux 3 (SL3, based on Red Hat Enterprise Linux) and are currently upgrading to SL4. The gLite middleware was developed for SL3, and as of release 3.1, gLite developers are working on migrating to SL4. The original LCG goals of wide portability have not been met. This affects the NGS with its more heterogeneous infrastructure.

Unlike the NGS, each GridPP site supports typically only one major experiment, but provides some resources to most GridPP VOs. Various mechanisms have been invented to allow the experiments to propagate non-standard (not included in the default release) software to the sites, such as the "VO box" or "pilot jobs" which pull in applications for the real jobs. Software is always open source so there is no licencing issues, except for the "3D" (distributed database deployment) which uses Oracle.

The GridPP deployment model is based on official releases, which are then usually deployed as a "pre-production service." A deployment team is responsible for coordinating the deployment in the UK and testing the software on GridPP resources. The NGS deployment is based on the user's requirements, but also on the recommendations of the Engineering Task Force (ETF). The role of the ETF is mainly to make technical evaluations of Grid middleware, to provide guidelines for the diverse NGS sites.

The NGS deployment model for the core sites could thus be brought closer to the GridPP model—or vice versa. Where GridPP has a deployment team, the NGS has the "operators", system administrators from the NGS sites, with a similar role. The NGS as a whole aims to cater for diverse sites, and GridPP could fit into this diversity. At the same time, plans should be made to test interoperability between the homogeneous core sites and the non-standard sites, insofar as they run the same services.

2.5.1 Globus The NGS makes use of the Globus Pre-WS software, Globus Toolkit 2. Since this has been unsupported by the Globus Alliance for some time, the software is currently supplied by the Virtual Data Toolkit. This includes several patches to give extra functionality. Firstly the standard version of the Globus toolkit requires a local user account and within a large multi-institutional grid this is not sustainable. The VDT system includes a patch which allows for the remote user's Distinguished Name (the name in the certificate) to be matched with a transient local pool account. This patch was originally designed by the GridPP collaboration since they encountered exactly the same problem.

There are plans to update this installation to Globus Toolkit 4 Pre-WS so that the upgrade path to WS based services is easier. This will be done through the use of the latest stable version of the VDT system. This will also allow for running a stable production service alongside a development environment using the WSRF based services. This can also be allowed to decide the full upgrade path dependant on real user demand.

2.6 Monitoring

NGS monitors sites using the GITS (Grid Integration and Test Scripts), developed by the NGS's Engineering Task Force. This is an active monitoring system, where jobs and services on each of the NGS sites are probed directly.

The INCA system from SDSC, also used by TeraGrid, has now been adopted by the NGS as part of its monitoring framework with the GITS tests integrated within INCA, and the framework implements a "some-to-all" testing model, that is, testing is performed from some (not all, but more than one) sites and tests are performed against all sites. Further work on INCA is planned to allow it to implement monitoring as specified in the NGS Monitoring Specification document [15].

A number of monitoring schemes are in place within GridPP. The first of these is the gstat monitor. This system polls the information as published by each GridPP site's BDII 2.2 and carries out basic sanity checks. Although it checks only the BDII, problems at this level can severely impact the usability of the site. For example, if the SRM storage system end points are not properly published then jobs may not be able to store their output, leading to wasted CPU cycles. Worse, in certain circumstances a poorly configured site may affect the performance of the whole grid, e.g., if empty queues are advertised on a site which is "blackholing" jobs, then large numbers of jobs on the grid may disappear.

The system is also used for storage accounting. GridPP developed a storage monitor ([21]) for EGEE.

However, gstat is only a passive system and thus rather limited in what it can probe at a site. For this reason a much more active monitoring system, the Service Availability Monitoring (SAM) system is also deployed.

SAM tests a site actively, either by directly testing advertised services at a site, e.g., the SRM storage interface, or by submitting jobs to the site which perform basic functionality. The core tests consist of ensuring that:

- A job can be submitted successfully to the site's CE through a resource broker.
- That the version of the gLite software found on the batch worker is valid.
- That the CA certificates found on the site are valid and up to date.
- The hostname of the job's broker is defined and is consistent.
- That the `csh` environment is sane.
- That basic replica management tasks can be performed:
  - A valid host for the information system provider is found.
  - A file can be copied from the worker node to the site's storage system and registered in the file catalog.
  - The file can be copied off the storage system
  - This file can be replicated to an off site storage system.
  - The file can be deleted from both storage systems.

If any of these tests fail then the site is considered to be unavailable. If tests fail for any length of time a ticket will be raised against the site.

Experience suggests that testing sites in this way does improve sites' performance, as shown in Figure 4.

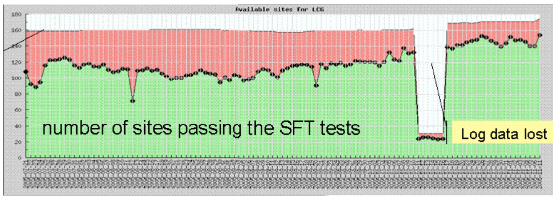

Figure 4: EGEE Sites Passing Basic SAM Tests

3 Conclusion

In this paper, we have presented a whirlwind tour of getting the two national Grids, NGS and GridPP, to interoperate. We have focused the areas that are either most mature, or most urgent to resolve. We have covered areas which are fully interoperable (X509 authentication, information systems), which are mostly or nearly interoperable (RBs, CEs, authorisation, job submission, monitoring), and where more work is required (SEs, other authentication such as Shibboleth, operating systems, WSRF, and middleware portability). Future developments will see further interoperability, to the benefit of UK science.

Acknowledgment

This work was funded by GridPP (PPARC) and NGS (JISC, EPSRC, and CCLRC). G.Stewart was supported by the EGEE project, co-funded by the European Commission (under contract number INFSO-RI-031688) through the Sixth Framework Programme. Document typeset with LaTeX.

References

[1] http://www.globus.org/toolkit/mds/

[2] http://www.cs.wisc.edu/condor

[3] https://twiki.cern.ch/twiki/bin/view/EGEE/BDII

[4] S Andreozzi *et al*: *GLUE Schema Specification version 1.2*, OGF GLUE-WG.

[5] http://public.eu-egee.org/

[6] http://www.gridway.org/

[7] http://www.gaussian.com/g_brochures/g03_intro.htm

[8] http://www.gridpma.org/

[9] http://www.gridpp.ac.uk/wiki/Disk_Pool_Manager

[10] A Forti, M Jones, S Dolgobrodov: *VOMS deployment in GridPP and NGS*, Proc. UK All Hands (2006)

[11] http://www.grid-support.ac.uk/content/view/81/61/

[12] J Jensen and M Viljoen: *Usability of the UK e-Science Certification Authority*, Proc. UK AHM 2005.

[13] D Wallom, D Spence, K Tang, M Viljoen, J Jensen, A Trefethen: *Shib-Grid, a Shibboleth based access method to the National grid service*, submitted to AHM 2007.

[14] J Jensen, D Spence, M Viljoen: *A scalable PKI for a national grid service*, Proc. 6th Annual PKI workshop, NIST.

[15] http://www.ngs.ac.uk/ops/mwg/NGS_Monitoring_ Specification_0_5.pdf]

[16] http://www.sve.man.ac.uk/Research/AtoZ/SHEBANGS

[17] https://gilda.ct.infn.it/

[18] http://www.camtology.com/

[19] http://www.total.com/

[20] http://indico.cern.ch/contributionDisplay.py?contribId=101&sessionId=13&confId=3738

[21] http://goc02.grid-support.ac.uk/accountingDisplay/index.php

[22] D McBride, S Young, T Harmer, T Parkinson, D Wallom: *UK Engineering Task Force: Preliminary Report on Lightweight SRM Evaluation, v0.2*

[23] G A Stewart, D Cameron, G A Cowan, G McCance: *Storage and Data Management in EGEE*. Proc. Conf. in Research and Practice in Information Technology, Volume 68. pp69-77 2007.

[24] http://sdm.lbl.gov/srm-wg/.

[25] http://www.sdsc.edu/srb/index.php/Main_Page