GridPP: Running a Production Grid

Olivier van der Aa[1], **Stephen Burke**[2], Jeremy Coles[2], David Colling[1], Yves Coppens[3], Greig Cowan[4], Alessandra Forti[5], Peter Gronbech[6], Jens Jensen[2], David Kant[2], Gidon Moont[1], Fraser Speirs[7], Graeme Stewart[7], Philippa Strange[2], Owen Synge[2], Matt Thorpe[2], Steve Traylen[2]

[1]Imperial College London [2]CCLRC [3]University of Birmingham [4]University of Edinburgh [5]University of Manchester [6]University of Oxford [7]University of Glasgow

Abstract

GridPP operates a Grid for the UK particle physics community, which is tightly integrated with the Enabling Grids for E-sciencE (EGEE) and LHC Computing Grid (LCG) projects. These are now reaching maturity: EGEE recently entered its second phase, and LCG is ramping up to the start of full-scale production in 2007, so the Grid must now be operated as a full production service. GridPP provides CPU and storage resources at 20 sites across the UK, runs the UK-Ireland Regional Operations Centre for EGEE, provides Grid-wide configuration, monitoring and accounting information via the Grid Operations Centre, and takes part in shift rotas for Grid operations and user support. It also provides support directly for its own system managers and users. This paper describes the structure and operations model for EGEE and LCG and the role of GridPP within that, focusing on the problems encountered in operating a worldwide production Grid with 24/7 availability, and the solutions which are being developed to enable this to succeed.

1. EGEE and LCG

Enabling Grids for E-sciencE (EGEE) [1] is an EU project to develop a large e-science Grid in Europe, which also has extensive links with projects in other countries. The first phase started in April 2004 and completed successfully in March this year. The second phase runs from April 2006 to March 2008, and discussions are currently underway concerning the future of the infrastructure beyond that date. EGEE is intended to serve the entire scientific community, although particle physics is currently the largest user.

The LHC Computing Grid (LCG) project [2] will provide the computing infrastructure for the Large Hadron Collider (LHC), which is due to start operation in 2007. Much of its infrastructure is provided within the EGEE framework, but it also needs to interoperate with other Grids, notably the US Open Science Grid (OSG) [3].

The combined EGEE/LCG Grid now consists of over 200 sites in around 40 countries (see Figure 1), provided by a large number of affiliated national and regional Grid projects, with in excess of 25,000 CPUs and 12 Pb of storage. There are well over 2000 registered users, organized into nearly 200 Virtual Organisations (VOs). Managing such a large, distributed infrastructure is a non-trivial problem, and the projects have been working to develop tools and procedures to provide a production-quality service.

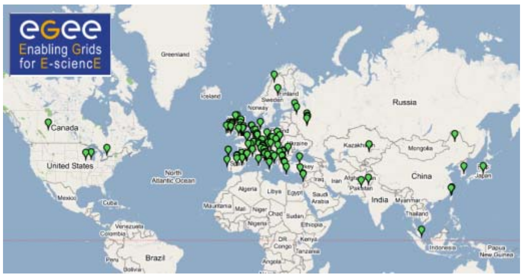

Figure 1: Sites which are part of the EGEE/LCG Grid, displayed on a Google map

1.1 EGEE structure

The EGEE project is split into a number of Activities in three broad areas: operations (SA), middleware development (JRA), and interactions with users and external bodies (NA). The SA1 Activity is responsible for Grid deployment and operations, and GridPP [4] is a major contributor to it. Overall co-ordination of deployment and operations is provided by the Operations Co-ordination Centre (OCC) at the CERN laboratory in Geneva.

EGEE is organised into regions, each of which has a Regional Operations Centre (ROC). The function of the ROC is to provide deployment and operational support for the sites in their region, to contribute to user support, and to represent the region to EGEE as a whole. CERN acts as a ROC for those sites without a regional affiliation.

GridPP is part of the UK-Ireland (UK-I) region, and runs the UK-I ROC. UK-I resources currently come largely from GridPP and GridIreland [5], but we also anticipate some degree of convergence with the National Grid Service (NGS) [6], and we are currently seeing some interest from other e-science sites in the UK.

GridPP also co-ordinates the Grid Operations Centre (GOC) [7]. This maintains a database of information about all the sites in the Grid and provides a variety of monitoring and accounting tools.

In EGEE-1 a subset of the ROCs were classed as Core Infrastructure Centres (CICs). These were involved in running critical services and taking on a rotating responsibility to monitor the state of the Grid and take action to solve problems. In EGEE-II the CIC functionality is being absorbed into the ROCs as more regions gain the necessary expertise. GridPP has been involved with the CIC functions from the start.

1.2 LCG structure

LCG is dedicated to providing for the computing needs of the experiments at the LHC particle accelerator. These are expected to produce several Pb of primary data per year, together with similar volumes of derived and simulated data, so data movement and storage is therefore a major consideration.

LCG organises its sites into a hierarchical, tiered structure. Tier 0 is at CERN, the laboratory which hosts the LHC accelerator. Files are shipped from there to a number of Tier 1 sites, which have substantial CPU, storage and manpower resources, high-capacity network connections, can guarantee high availability, and are able to make a long-term commitment to supporting the experiments. Tier 1 sites will also need to transfer data between themselves.

Each Tier 1 site supports a number of Tier 2 sites, which are allowed to have lower levels of resources and availability. Files are expected to be transferred in both directions between Tier 2 sites and their associated Tier 1, but generally not on a wider scale, although the detailed data movement patterns are still under discussion.

LCG is putting in place Service Level Agreements (SLAs) to define the expected performance of the Tier centres, and is also developing monitoring tools to measure actual performance against the SLAs. A set of Service Challenges are being run, principally to test the routine transfer of large data volumes but also to test the full computing models of the LHC experiments.

In the UK, GridPP has a Tier 1 site at the Rutherford Appleton Laboratory [8] and 19 Tier 2 sites at Universities around the country [9]. For administrative convenience these are grouped into four “virtual Tier 2s”. From a technical point of view each site operates independently, but the Tier 2s each have an overall co-ordinator and are moving towards some sharing of operational cover.

1.3 Virtual Organisations

Users are grouped into Virtual Organisations (VOs), and sites can choose which VOs to support. EGEE is developing formal policies under which VOs can be approved, but at present there are effectively several different categories of VO:

- LCG-approved VOs support the four LHC experiments, and also some associated activities.
- EGEE-approved VOs have gone through a formal induction process in the EGEE NA4 (application support) Activity, and receive support from EGEE to help them migrate to the Grid. However, this is a relatively heavyweight process and so far rather few VOs have come in via this route.
- Global VOs have a worldwide scope, but have not been formally approved by EGEE and do not receive direct support from it. At present these are predominantly related to particle physics experiments.
- Regional VOs are sponsored, and largely supported, by one of the participating Grid projects. The UK-I ROC currently has a small number of VOs in this category, but the number is expected to increase. Regional VOs may also be supported at sites outside the region if they have some relationship with the VO.

Although the status of the VOs varies, as far as possible they are treated equivalently from a technical point of view in both the middleware and the infrastructure. Each VO has a standard set of configuration parameters, which makes it relatively easy for a site to add support for a new VO (unless there are special requirements).

As a particle physics project GridPP is mainly interested in supporting VOs related to the experiments in which the UK participates. However, it has a policy of providing resources at a low priority for all EGEE-approved VOs.

2. Middleware

The EGEE/LCG middleware is in a state of constant evolution as the technology develops, and managing this evolution while maintaining a production service is a major component of Grid deployment activity. The current middleware is based on software from Globus [10] (but still using the Globus toolkit 2 release dating from 2002), Condor [11] and other projects, as collected into the Virtual Data Toolkit (VDT) [12]. A substantial amount of the higher-level middleware came from the European DataGrid (EDG) project [13], the predecessor of EGEE. Many other tools have been developed within LCG. The middleware Activity in EGEE (JRA1) has developed new software under the gLite brand name, and an integrated software distribution has recently been deployed on the production system which has adopted the gLite name [14].

The middleware can broadly be split into three categories: site services which are deployed at every site, core services which provide more centralised functions, and Grid operation services which relate to the operation of the Grid as a whole.

2.1 Site services

Every site needs to deploy services related to processing, data storage and information transport. There has also recently been some discussion about providing facilities for users to deploy their own site-level services. The individual services are as follows:

- Berkeley Database Information Index (BDII) – this is an LDAP server which allows sites to publish information according to the GLUE information schema [15].
- Relational Grid Monitoring Architecture (R-GMA) – another information system presenting information using a relational model. This is also used to publish the GLUE schema information, together with a variety of monitoring and accounting information, and is also available for users to publish their own data.
- Computing Element (CE) – this provides an interface to computing resources, typically via submission to a local batch system. This has so far used the Globus gatekeeper, but the system is now in transition to a new CE interface based on Condor-C. A tool called APEL (Accounting Processor using Event Logs) is also run on the CE to collect accounting information, and transmit it via R-GMA to a repository at the GOC.
- Worker Nodes (WN) – these do not run any Grid services (beyond standard batch schedulers), but need to have access to the client tools and libraries for Grid services which may be needed by running jobs, and often to VO-specific client software.
- Storage Element (SE) – this provides an interface to bulk data storage. Historically the so-called “classic SE” was essentially just a Globus GridFTP server, but the system is now moving to a standard interface known as the Storage Resource Manager (SRM) [16].
- VOBOX – there have recently been requests to have a mechanism to deploy persistent services on behalf of some VOs. This is currently under discussion due to concerns about scalability and security, but some such services are currently deployed on an experimental basis.

2.2 Core services

Some services are not needed at every site, but provide some general Grid-wide functionality. Ideally it should be possible to have multiple instances of these to avoid single points of failure, although this has not yet always been achieved in practice. The services are:

- Resource Broker (RB) – this accepts job submissions from users, matches them to suitable sites for execution, and manages the jobs throughout their lifetime. A separate Logging and Bookkeeping (LB) service is usually deployed alongside the broker.
- MyProxy [17] – a repository for long term user credentials, which can be contacted by other Grid services to renew short-lived proxies before they expire.
- BDII – the information from the site-level BDII servers is collected into Grid-wide BDIIs which provide a view of the entire Grid. BDIIs are usually associated with RBs to provide information about the available resources to which jobs can be submitted.
- R-GMA Registry and Schema – these are the central components of the R-GMA system, providing information about producers and consumers of data and defining the schema for the relational tables. At present there is a single instance for the whole Grid.
- File Catalogue – this allows the location of file replicas stored on SEs using a Logical FileName (LFN). The system is currently in transition from the EDG-derived Replica Location Service (RLS) to the LCG File Catalogue (LFC) which has improved

performance. At present this is generally deployed as a single instance per VO, but the system is designed to allow distributed deployment.

- File Transfer Service (FTS) – this provides a reliable, managed service to transfer files between two SEs.
- User Interface (UI) – this has the various client tools to allow jobs to be submitted and monitored, files to be managed, R-GMA to be queried etc. This can be a shared service or installed on a user's own machine.
- VO server – this stores the membership information for a VO. The system is currently in transition from a solution based on LDAP servers to a more powerful system called the VO Membership Service (VOMS). At present these servers are normally deployed as a single instance per VO, although there is some scope for redundancy.

2.3 Grid operation services

It has proved to be necessary to develop some infrastructure services/tools to support the operation of the Grid as a whole. These generally relate to testing and monitoring the system:

- The GOC Database (GOCDB) has information about every site in the system, including its status, contact details, and a full list of nodes. This is used by various other tools, e.g. to generate lists of sites/nodes to monitor. The information is also republished into R-GMA.
- The operations portal [18] is a web site which collects a variety of operational information and tools for users, VO managers, site and ROC managers, and operations staff.
- The Global Grid User Support (GGUS) portal [19] is a web site which enables support tickets to be entered and tracked, and also has links to a wide variety of user support information and documentation.
- The Site Functional Tests (SFTs) [20] are test jobs which run regularly at each site and test the major Grid functionality. The results are reported back to a central database, and can be displayed in various formats (see Figure 2). They can be used both to flag problems at the sites, and to measure the availability and reliability of sites and services. This is currently being upgraded to a new system known as SAM (Site Availability Monitoring).
- The Grid status monitor (GStat) [21] reads information published by each site according to the GLUE schema, and uses this to display the current status and history, and flag potential problems.
- The GridView tool collects file transfer information, published from GridFTP into R-GMA, and displays it.
- Freedom of Choice for Resources (FCR) is a tool which allows the selection of sites seen by an RB through a BDII to be customised dynamically. In particular, sites which are failing either standard or VO-specific tests can be removed, and individual sites can also be white- or black-listed.
- The Real Time Monitor (RTM) [22] collects information about jobs from LBs. This can be used to display the state of the system in real time, and also allows historical data to be processed to derive statistics on such things as job volumes, duration and failure rates.
- GridICE [23] displays monitoring information published both in the GLUE schema and in GridICE-specific extensions.

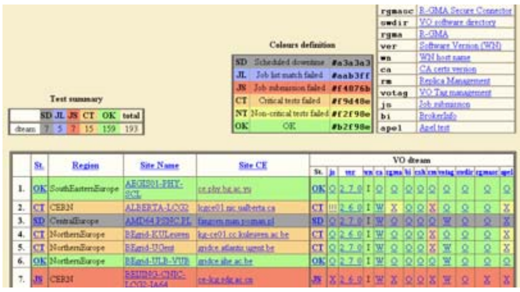

Figure 2: Part of the Site Functional Test display

2.4 Deployment issues

Middleware components are often developed by relatively small groups, with a single development platform and with installation and testing on a small system which is under their complete control. They are also typically developed as standalone services, rather than being designed from the start to work with other components in an integrated way. This can lead to the emergence of a number of problems when the middleware is deployed on a large-scale, heterogeneous system.

A basic issue is scalability; many components encounter problems coping with the large number of sites/jobs/VOs/users/files in the production system, and with continuous operation over long periods. Unfortunately this is effectively confronted only on the production system, as test systems are too small to detect the problems.

It is often very difficult to port middleware to new operating systems (or even new versions of a supported operating system), compilers and batch systems. EGEE has an objective of supporting a wide range of platforms, but so far this has not been achieved in practice, and porting remains a painful exercise.

The middleware typically depends on a number of open-source tools and libraries, and the consequent package dependencies can be difficult to manage. This is particularly true when the dependency is on a specific version – which can lead to situations where different middleware components depend on different tool versions. This is partly a result of poor backward-compatibility in the underlying tools. These problems can be solved, but absorb a significant amount of time and effort.

In a large production Grid it is not practical to expect all sites to upgrade rapidly to a new release, so the middleware needs to be backward-compatible at least with the previous release and ideally the one before that. Any new services need to be introduced in parallel with the existing ones. This has generally been achieved, but is sometimes compromised by relatively trivial non-compatible changes.

The most important issue is configuration. Much of the middleware is highly flexible and supports a large number of configuration options. However, for production use most of this flexibility needs to be hidden from the system managers who install it, and Grid deployment therefore involves design decisions about how the middleware needs to be configured. There is also a scalability aspect, for example services should be discovered dynamically wherever possible rather than relying on static configuration. The middleware should also be robust in the sense that small changes to the configuration should not prevent it working, which is not always the case. In practice, finding a suitable configuration and producing a well-integrated release is a major task, which can take several months.

2.5 Operational issues

A production Grid requires continuous, reliable operation on a 24/7 basis. In a large system there are always faults somewhere, which may include hardware failures, overloaded machines, misconfigurations, network problems and host certificate expiry. Middleware is often written under the assumption that such errors are exceptional and must be rectified for the software to function. In a large Grid such "exceptions" are in fact the norm, and the software needs to be able to deal with this. In addition, error reporting is often regarded as having a low priority, and consequently it can be very hard to understand the reason for a failure. Ideally:

- Middleware should be fault tolerant; wherever possible it should retry or otherwise attempt a recovery if some operation fails.
- Services should be redundant, with automatic failover.
- Logging and error messages should be sufficient to trace the cause of any failure.
- It should be possible to diagnose problems remotely.

In practice these conditions are usually not fulfilled in the current generation of middleware, and it is therefore necessary to devise operational procedures to mitigate the effects of failures.

3. Grid deployment

Deployment relates to the packaging, installation and configuration of the Grid middleware.

3.1 Middleware installation and configuration

Various installation/configuration tools have been developed in a number of Grid projects which allow much of the middleware flexibility to be exposed in the configuration. However, as mentioned above, in practice system managers prefer to have the vast majority of configuration flexibility frozen out. LCG has therefore developed a tool called YAIM (Yet Another Installation Method) which uses shell scripts driven from a simple configuration file with a minimal number of parameters.

3.2 Middleware releases

Release management is a difficult area in a situation where the middleware continues to evolve rapidly. So far there have been major production releases about three times a year, with all the software being upgraded together. (This excludes minor patches, which are released on an ad-hoc basis.) However, this has caused some difficulties. On one side it takes a long time for new functionality, which may be urgently needed by some users, to get into the system. On the other hand, this is still seen as too rapid by many sites. Practical experience is that it takes many weeks for most sites to upgrade to a new release, with some sites taking several months.

Deployment is now moving towards incremental upgrades of individual services to try to alleviate this situation. However, this may bring problems of its own, as the backward-compatibility requirements may become much greater with a more complex mix of service versions at different sites. The introduction of a Pre-Production System (PPS) to try out pre-release software should however enable more problems to be found before new software goes into full production.

Support for release installation is provided principally on a regional basis; GridPP has a deployment team which provides advice and support to sites in the UK.

3.3 Documentation and training

Documentation is needed for both system managers and users. Unfortunately there is little or no dedicated effort to produce it, and the rapid rate of change means that documentation which does exist is often out of date. Documentation is also located on a wide range of web sites with little overall organisation. However, more effort is now being made to improve this area.

For system managers the central deployment team provide installation instructions and release notes. There is also a fairly extensive wiki [24] which allows system managers to share their experiences and knowledge.

The biggest problem with user documentation is a lack of co-ordination. The main User Guide is well-maintained, but in other areas it relies on the efforts of individual middleware developers. Many user-oriented web sites exist, but again with little co-ordination. GridPP has its own web site with a fairly extensive user area [25], which attempts to improve the situation rather than add to the confusion by providing structured pointers to documentation held elsewhere. EGEE has recently set up a User Information Group (UIG) with the aim of co-ordinating user documentation, based around a set of specific use-cases.

Training within EGEE is provided by the NA3 Activity, which is co-ordinated by NESC [26]. The training material is also available on the web. However, again there is a problem with keeping up with the rate of change of the system, and also to some extent with reaching the highly diverse community of people who might require training.

One especially difficult area is the question of the induction of new VOs, to allow them to understand how to adapt their software and methods to make best use of the Grid. EGEE is intending to run periodic User Forum events at which users can share their experiences to improve the transfer of such information, with the first having been held in March this year [27]. In GridPP, so far most users have been from the large LCG VOs who have a fairly well-developed engagement with the Grid, but we will need to develop induction and training for internal VOs which may only have a small number of members.

3.4 VO support at sites

EGEE is intended to support a large number of VOs, and at present there are getting on for 200 VOs enabled somewhere on the infrastructure. This implies that it should be easy for a site to configure support for a new VO. However, there are currently some problems with this:

- The key configuration information is not always available. The operations portal has templates to allow VOs to describe themselves, but in most cases this is not yet in use.
- Supporting a new VO implies changes in various parts of the configuration, so this needs to be done in a parameterised way to make it as simple as possible - in the ideal case a site should simply have to add the VO name to a configuration file. YAIM now incorporates a web-based tool to generate the required configuration for those sites (the majority) which use it, which goes a long way towards this goal. Unfortunately, at present this also suffers from the fact that the configuration information for many VOs is not yet present.
- Some VOs may require non-standard configurations or extra services, and this decreases the chances that they will be supported by sites which do not have a strong tie to the VO. This is partly a question of user education; VOs need to realise that working within a Grid environment can accommodate less site-specific customisation than when working with a small number of sites. It may also imply the development of a standardised way to deploy non-standard services, the “VO box” concept mentioned above.

- Sites, and the bodies which fund them, may also need to adapt their attitude somewhat to the Grid model. Sites are often funded, and see themselves, as providing services for a small number of specific VOs, rather than a general service to the Grid community as a whole.

3.5 VO software installation

VOs often need to pre-install custom client tools and libraries which will be available to jobs running on WNs. The current method, which has been found to work reasonably well, is to provide an NFS-shared area identified by a standard environment variable, which is writeable by a restricted class of VO members. These people can also use GridFTP to write version tags to a special file on the CE, which are read and published into the information system, allowing jobs to be targeted at sites which advertise particular software versions. This tagging method can also be used to validate any other special features of the execution environment which are required by the VO.

4. Grid operations

As discussed above, failures at individual sites are a fact of life in such a large system, and the middleware currently does not deal with errors very effectively in most cases. The "raw" failure rate for jobs can therefore be rather high. In the early days it was also not unusual for sites to appear as "black holes": some failure modes can result in many jobs being sent to a site at which they either fail immediately, or are queued for a long period.

To get to a reasonable level of efficiency (generally considered to be a job failure rate below 5-10%) it has been necessary to develop operational procedures to mask the effects of the underlying failures. The key tool for this is the SFT system, which probes the critical functionality at each site on a regular basis. Failing sites can be removed from being matched for job submission, using criteria defined separately for each VO, using the FCR tool. In parallel with this an operations team submits tickets to the sites to ensure that problems are fixed as quickly as possible, as described below. All these systems and tools continue to be developed as we gain more experience. Performance is measured by a variety of tools, e.g. the RTM, which allow the current situation and trends over time to be monitored.

As mentioned earlier, LCG uses resources from other Grids, notably OSG in the US, and efforts are currently underway to integrate these into the EGEE/LCG operations framework as far as possible.

4.1 Site status

New sites are initially entered into the GOC DB as uncertified. The ROC to which they belong is expected to perform basic tests before marking the site as certified, at which point it becomes a candidate for regular job submission, and starts to be covered by the routine monitoring described below. Sites can also be marked as being in scheduled downtime, which temporarily removes them from the system.

4.2 Regular monitoring

An operations rota has been established to deal with site-level problems, with ROCs (formerly just CICs) being on duty for a week at a time. While on duty the various monitoring tools are scanned regularly (information is aggregated for this purpose on the operations portal). If problems are found tickets are opened against the relevant sites, and followed up with a specified escalation procedure, which in extreme cases can lead to sites being decertified. At present, with around 200 sites there are about 50 such tickets per week, which illustrates both that active monitoring is vital to maintaining the Grid in an operational state, and that there is never a time when the system is free of faults.

Tickets are initially opened in the GGUS system, but each ROC has its own internal system which it uses to track problems at the sites it serves, using a variety of technologies. An XML-based ticket exchange protocol has therefore been developed to allow tickets to be exchanged seamlessly between GGUS and the various regional systems.

Summaries of the operations activity are kept on the CIC portal, and issues are discussed at a weekly operations meeting.

4.3 Accounting

Accounting information is collected by the APEL tool and published to a database held at the GOC [28]. At present this only relates to the use of CPU time, but storage accounting is in development. Accounting information is currently only available for entire VOs (see Figure 3) due to concerns about privacy, but user-level information is likely to be required in the future once the legal issues are clarified. Accounting information is also available per site and with various degrees of aggregation above that. So far accounting data is only used for informational purposes, but it is likely that it will be used in a more formal way in future, e.g. to monitor agreed levels of resource provision.

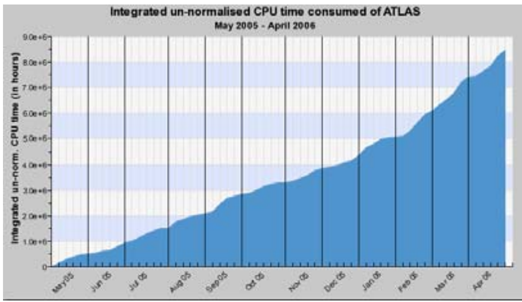

Figure 3: Accounting - integrated CPU time used by the atlas VO

4.4 User support

The question of user support initially received relatively little attention or manpower, but it has become clear that this is a very important area, and in the last year or so a substantial amount of effort has been put into developing a workable system. This is again based on the GGUS portal, where tickets can be entered directly or via email. People known as Ticket Processing Managers (TPMs) then classify the tickets, attempt to solve them in simple cases, or else assign them to the appropriate expert unit, and also monitor them to ensure that they progress in a timely way. TPMs are provided on a shift rota by the ROCs, and GridPP contributes to this, as well as providing some of the expert support units.

At present the system is working reasonably well, with typically a few tens of tickets per week. However, as usage of the system grows the number of tickets is also likely to grow, and it is not clear that the existing system can scale to meet this. The biggest problem is not technical, but the lack of manpower to provide user support, especially given that supporters need to be experts who generally have many other tasks. In the longer term we will probably need to employ dedicated support staff.

5. Summary

GridPP runs the UK component of the worldwide EGEE/LCG Grid. This is a very large system which is now expected to function as a stable production service. However, the underlying Grid middleware is still changing rapidly, and is not yet sufficiently robust. It has therefore been vital to develop Grid operations procedures which allow users to see only those parts of the system which are working well at any given time, and to manage software upgrades in a way which does not disrupt the service. This is still being developed but is now working well, and usage of the Grid has been rising steadily. Many challenges nevertheless remain, especially as the LHC starts to take data.

References

[1] http://www.eu-egee.org/
[2] http://lcg.web.cern.ch/LCG/
[3] http://www.opensciencegrid.org/
[4] http://www.gridpp.ac.uk/
[5] http://cagraidsvr06.cs.tcd.ie/index.html
[6] http://www.ngs.ac.uk/
[7] http://goc.grid-support.ac.uk/gridsite/gocmain/
[8] http://www.gridpp.ac.uk/tier1a/
[9] http://www.gridpp.ac.uk/tier2/
[10] http://www.globus.org/
[11] http://www.cs.wisc.edu/condor/
[12] http://vdt.cs.wisc.edu/
[13] http://www.edg.org/
[14] http://www.glite.org/
[15] http://infnforge.cnaf.infn.it/glueinfomodel/
[16] http://sdm.lbl.gov/srm-wg/
[17] http://grid.ncsa.uiuc.edu/myproxy/
[18] http://cic.in2p3.fr/
[19] http://www.ggus.org/
[20] https://lcg-sft.cern.ch/sft/lastreport.cgi
[21] http://goc.grid.sinica.edu.tw/gstat/
[22] http://gridportal.hep.ph.ic.ac.uk/rtm/
[23] http://gridice2.cnaf.infn.it:50080/gridice/site/site.php
[24] http://goc.grid.sinica.edu.tw/gocwiki/FrontPage
[25] http://www.gridpp.ac.uk/deployment/users/
[26] http://www.egee.nesc.ac.uk/
[27] http://indico.cern.ch/conferenceDisplay.py?confId=286
[28] http://goc.grid-support.ac.uk/gridsite/accounting/