

Grid Deployment and Operations for EGEE, LCG and GridPP

Jeremy Coles

GridPP – CCLRC Rutherford Appleton Laboratory (RAL)

Abstract

Over the last year the Enabling Grids for E-science (EGEE) project has seen its deployed resources grow to over 10,000 CPUs in more than 30 countries. During this period the UK and Ireland Regional Operations Centre (UKI ROC) has played a significant role in the development and refinement of the presently used deployment and operations model. This paper reviews the background to the EGEE project and describes these models. As a large part of the EGEE resources are provided by sites of the Large Hadron computing Grid (LCG), and within the UK by the Grid for Particle Physics project (GridPP), this paper describes the deployment planning and operational activities surrounding these organisations and provides some insight into future plans and challenges.

1. Participation in EGEE

The EU funded Enabling Grids for E-science¹ (EGEE) project aims to develop a production service grid infrastructure which is available to scientists 24 hours-a-day.

1.1 The structure of EGEE

The project has activities that fall into eleven “activities” in three main areas: Networking Activities (NA), Specific Service Activities (SA) and Research Activities (JRA). Grid deployment, operation and management come under the area SA1. The management structure consists of a centralised Operations Management Centre (OMC) at CERN, Geneva, and Regional Operations Centres² (ROCs) representing the main participating countries. The ROC managers and OMC representatives ensure that resources are made available, upgraded, monitored and maintained according to policies and service level agreements put forward by them. The project currently has over 10,000 CPUs deployed and over 900 registered users.




Figure 1: Resources in EGEE over time

1.2 The UK and Ireland ROC

The UK and Ireland (UKI) ROC³ has contributed widely to the development of the models for deployment and operations described later. The ROC consists of three partners: Grid Ireland⁴, the UK National Grid

Service⁵ (NGS) and the leading partner, the UK Grid for Particle




Figure 2: UK regional deployment structure

Physics⁶ (GridPP). At present most of the GridPP resources are within EGEE largely because EGEE was founded on the back of advances in grid technology made through the European DataGrid Project⁷ (EDG) and its successor the Large Hadron collider Grid⁸ (LCG) –the latter has a primary goal of producing a grid infrastructure to cope with the vast data output from CERN LHC experiments from 2007 onwards. NGS and Grid Ireland resources in EGEE are expected to increase over the coming year.

2. Deployment

The UKI ROC team consists of a manager, deployment coordinators for each of the geographical regions (see Figure 2), a security officer, helpdesk (provided by the Grid Operations Support Centre at RAL) and a Core Infrastructure Centre (CIC) group. All activities

are overseen by a GridPP Tier-2 board and a GridPP led Deployment Board.

2.1 Middleware installations

The EGEE deployment team release a main middleware tag approximately every 3 months. Between these main releases are minor bug fix and security updates. Site administrators install the latest release (held in a central repository) with help from their coordinators and the UK support team. Support for administrators and users are also provided by a central EGEE group based in Karlsruhe, Germany. The Global Grid User Support⁹ (GGUS) facility is being built to interact directly with regional support desk applications and teams, and to provide a core into which each Virtual Organisations (VOs) can link with user specific support. In addition GGUS provides links to useful tutorials (in conjunction with the National E-Science Centre in Edinburgh) and other support information such as the LCG Wiki¹⁰. With a rapidly growing repository of information and knowledge and a frequently changing environment (both the middleware and deployment tools) many lessons have been learnt and captured in this Wiki and shared through a number of mailing lists and areas like the GridPP deployment web pages.

2.2 Operations

The integrity of all the grid resources is constantly monitored by one of the (currently four) Core Infrastructure Centres. Each CIC takes responsibility for monitoring the daily-run job submission test results and grid statistics (gstat) pages¹¹. Any problems which are found are recorded in a problem tracking tool and followed up using additional information provided by the Grid Operations Centre database (e.g. contact information, node and scheduled downtime details for the site) developed and supported by RAL. Changes in operational responsibility take place at weekly handover meetings where current issues and problems are discussed.

3. GridPP and LCG planning

Due to the operating performance required for the LHC by 2007, current activities are constantly pushing the capabilities of the deployed infrastructure. They also require good interoperability and management with NorduGrid¹² and the US Open Science Grid¹³.

3.1 Service challenges

LCG success is being improved by a decision of

its strategic body, the Grid Deployment Board, to run a series of challenges between now and 2007. GridPP is responsible for the UK contribution to these activities which have become known as Service Challenges.

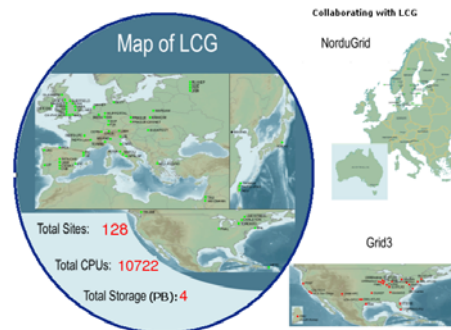


Figure 3: LCG uses 3 major grids

The challenges include targets for reading and writing data between sites for both disk and tape. They have required the provisioning of dedicated light paths and deployment of new Storage Resource Management technologies.

3.2 The future

GridPP has devised a deployment plan which is updated quarterly. As well as details of internal challenges to ensure site readiness, it contains milestones and associated metrics for monitoring the overall progress of deployment. Nevertheless EGEE, LCG and GridPP face a number of related challenges. Broadly these are in the areas of security, site configuration and maintenance, network bandwidth provision and an expected transition from LCG2 middleware to gLite middleware (a Web Services architecture which addresses shortcomings found with LCG2 functionality and stability). The coming year should see many advances in performance and functionality in both the middleware and the surrounding deployment and operations activities.

¹ <http://egee-intranet.web.cern.ch/>

² <http://egee-sa1.web.cern.ch/egee%20sa1/>

³ <http://goc.grid-support.ac.uk/ROC/>

⁴ <http://cagraidsvr06.cs.tcd.ie/>

⁵ <http://www.ngs.ac.uk/>

⁶ <http://www.gridpp.ac.uk/>

⁷ <http://eu-datagrid.web.cern.ch/eu-datagrid/>

⁸ <http://lcg.web.cern.ch/LCG/>

⁹ <https://www.ggus.org>

¹⁰ <http://goc.grid.sinica.edu.tw/gocwiki/>

¹¹ <http://goc.grid-support.ac.uk/gridsite/monitoring/>

¹² <http://www.nordugrid.org/>

¹³ <http://www.opensciencegrid.org/>