A Grid for Particle Physics – from testbed to production

P. Clarke,
on behalf of the GridPP Collaboration
Department of Physics and Astronomy
University College London
London WC1E 6BT

Abstract

The GridPP project, in close association with the European DataGrid project, will reach a key milestone this year. The projects will have completed their first phase leading to a very successful prototype production Grid testbed. We describe the pros and cons of this testbed, the value added middleware necessary to make it function over diverse users across the globe, the use which applications have made of it, and lessons learned. We then describe briefly the strategy for the future – whereby we intend to move from a testbed oriented project to a production service able to meet the needs of the LHC and other particle physics programmes.

Introduction

The GridPP project, in close association with the flagship European DataGRID project (EDG), has achieved remarkable successes in the last three years, leading to a pervasive Grid testbed existing throughout Europe and including the UK as a major component. It is perhaps not always appreciated that this is a working production Grid prototype, to which real user production jobs can be submitted in a properly manageable and scalable way across national boundaries. Put simply, a user can submit a job to a broker in Italy, and find it running in Glasgow or Amsterdam a few minutes later. There are few such Grids in operation in the world today. This success was perhaps best underlined at the recent EU Review of EDG, where all participants were commended for the level of achievement attained by the project. Figure 1 shows a snapshot of the testbed just before the final review.

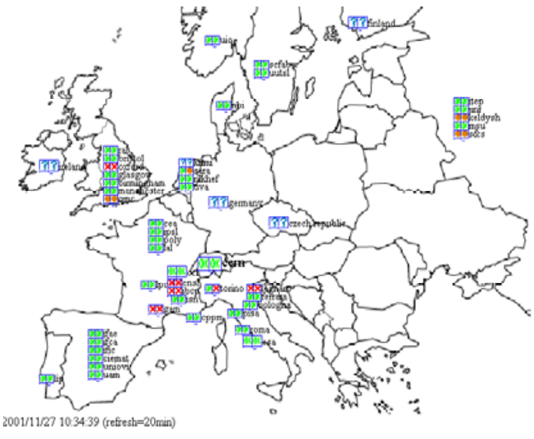

Figure 1: European DataGrid Testbed sites throughout Europe

The infrastructure

What has GridPP achieved ? Deploying low level tools, such as Globus, does not in itself provide a Grid. To enable a diverse set of users scattered throughout the world to access the Grid requires several "value added" components listed in the table below. GridPP and EDG have developed good working prototypes of all of these, and in this presentation we will briefly describe the current status and plans for evolution of each.

Resource Information Services	*An enhanced information services system was developed. This is based upon the R-GMA publish and subscribe architecture, and was required to provide the much richer information set representing "computing", "storage" and "network" elements needed by the workload management system.*
Workload Management and logging services	*A high capacity workload management framework was developed to accept user requests and match these to available resources around the globe. Several of these are now in full time operation, including one at ICSTM in the UK.*
Virtual Organisation Management	*For a Grid to function a scalable and comprehensive identity and credential management system was required to allow resource providers to verify identity and capability of people belonging to large distributed virtual organisations. This was a key development without which a scalable grid, serving many and varied users cannot be constructed*
Installation and Configuration tools	*An automated system for installation and configuration of code based was developed and used to simplify prototype deployment.*
Storage and Data Management	*A Grid interface to heterogeneous storage systems, exposing access via different transport bindings was required to access and manage storage of files across resources. The data itself was managed by a comprehensive replica location system, and a web services based metadata management facility layered on to of this.*
Application Software Support	*The particle physics community consists of many and varied experiments at all stages of production. These were supported to interface their simulation and analysis code to the Grid infrastructure.*

Usage Examples

The prototype production grid shown in the Figure 1 was used for several real life particle physics applications, as well as exemplar demonstrations from the fields of Earth Observation and the Biomedical sector.

The plot below depicts the time evolution of a production run within the CMS experiment (an experiment to run at the Large Hadron Collider in CERN) , in which 250,000 simulated events were created. Simulation jobs were submitted to the resource broker, and then farmed out to centres in the UK, Italy and France over a period of nearly three weeks.

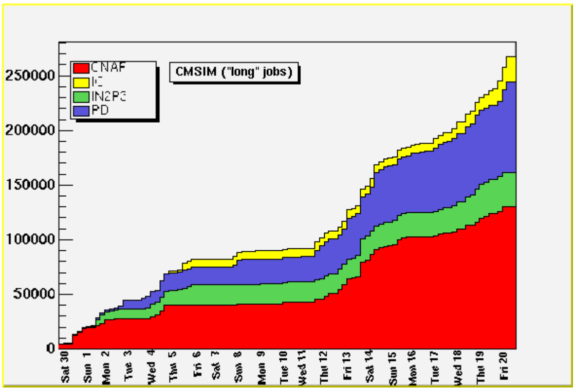

Figure 2: Number of events produced as a function of time (in days) in a CMS-Experiment Data Challenge. The different colours represent different centres across Europe.

In addition to the LHC experiments, the technology has been successfully used to benefit several other Particle Physics interests, including the distributed data handling and job processing facility used by the CDF and D0 experiments at Fermilab, the UK QCDGRID and the BaBar experiment.

In the context of a completely different discipline, the EDG grid testbed was used in the field of radiotherapy to treat brain tumors. Current analytic dose calculation techniques can make errors of to 20%. Monte-carlo simulations are potentially much more accurate, but without the grid these are unacceptably time consuming from the point of view of the consultant. Instead, the grid testbed was used to demonstrate the ability to make a rapid dose distribution calculation in near real time.

In the presentations further example of the usability of the grid testbed will be described.

The future – toward a production Grid

The success of GridPP needs to feed to the future. The PPARC funded GridPP-2 project has just been approved, and in GridPP-2 one of the principle goals is to move from an era where middleware development and testbed construction has been the focus, to an era where we focus upon building a production quality Grid running production quality services to meet the LHC challenge.

The very limitations which meant that the prototype grid could be successfully deployed on a relatively short timescale, now become the principle issues for the next few years. These are (i) removing the dependence on the limited Linux deployment platform arising due to the complex and nature of the middleware prototypes and particle physics analysis and simulation jobs (ii) the need for true robust production quality code (iii) the need for scalable and businesslike operational procedures.

All of these issues will be addressed by GridPP-2 through its collaboration with the UK eScience Grid operations centres, the LHC Computing Grid (LCG), the flagship European Production Grid infrastructure project known as EGEE, and the OMII. The next few years will see Grids for Science becoming the norm for on-demand computing. In GridPP-2 we foresee a much more "production" oriented organisation than in GridPP-1. The project will contain two boards. A "user board", will have the task of ensuring user requirements are prioritised for the "production and deployment board" whose job will be to ensure that a stable and robust service is deployed. This latter will of course work in close cooperation with the UK Grid operations centre, and indeed many of the resources will be shared. GridPP-2 will be particularly closely linked to the CERN based LCG project, with the mission to provide the Grid needed for the LHC experiments for turn on 2007/8.

The planning of GridPP-2 and its interaction with other projects, bodies and initiatives will be described.