THE QUANTUM CHROMODYNAMICS GRID

James Perry, Andrew Jackson, Matthew Egbert, Stephen Booth, Lorna Smith.
EPCC, The University Of Edinburgh (The QCDgrid project, on behalf of GridPP).

Key words to describe the work: DataGrid, Quantum Chromodynamics, Globus, European DataGrid.

Key Objectives: To implement a 'QCDgrid' to become a production environment for UKQCD, a collaboration of UK Scientists carrying out Quantum Chromodynamics (QCD) simulations.

Motivation for the work (problems addressed): The advent of more and more powerful supercomputing resources has led to many HPC applications generating huge amounts of data, often in the terabyte to petabyte range. This in turn has led to significant challenges in data storage. One such application area is QCD, which is constantly demanding larger and larger supercomputing resources and generating larger and larger amounts of raw physical data. UKQCD currently stores around five terabytes of data on the EPCC J90 mass store. In addition, this figure is expected to grow dramatically as the collaboration's purpose built HPC system, QCDOC, comes on line. Hence it is essential that UKQCD has access to a secure, reliable and expandable distributed storage systen to handle data effectively. The QCDGrid project, which is part of the GridPP project, aims to address this issue, by developing a mulit-terabyte storage system, which supports distributed data management across different UK sites.

Please complete the text of your abstract here. Please note that your completed document, including answers to compulsory questions, diagrams and references, must not exceed two pages

The first phase of this project involved developing a prototype of the datagrid, which will run across a series of Red Hat Linux PCs with large RAID arrays of hard disks located at four sites across the country: Edinburgh, Glasgow, Liverpool and Swansea.

The data to be stored on the grid is of great scientific value and therefore the most important requirement of the software is that the files it manages are protected from accidental loss or damage. This is achieved through redundancy - on each individual storage element the mass storage system is a RAID array which means the data is duplicated on a second disk automatically. Thus no single disk failure can cause data loss.

However, the QCDgrid project has even more stringent requirements – the grid must be able to cope with the loss of an entire site without losing any data. This means that each file must be stored on at least two storage elements, and these elements must be physically located at different sites.

Despite this fault tolerance requirement, the software cannot compromise on usability or efficiency. Ease of use is ensured by hiding most of the underlying complexity of the datagrid from the user – each file on the grid is referred to by its logical filename, and the user need not know or care where the file is actually stored.

Efficiency is always a concern when files are distributed across the whole country - obviously it generally takes longer to transfer a file from Swansea to Edinburgh than it would to transfer it from another machine at Edinburgh. This problem can be overcome in the software by allowing users to register interest in having a particular file stored on a storage element located physically close to them. The grid software will then take this request into account when deciding where to store the file.

A prototype version of this grid software has been developed, running across systems located at three different sites. Although the project is still at a relatively early stage, most of the functionality described above has been implemented in the test grid.

Technically, the software uses the Globus grid toolkit for basic grid operations such as data transfer, security and remote job execution. It also uses the Globus replica catalogue to maintain a directory of the whole grid, listing where each file is currently stored. The machine hosting the replica catalogue also hosts a central control thread which constantly scans the grid, making sure the storage elements are all working and the files are all stored in at least two suitable locations.

In addition to storing the raw physical data, the project aims to provide an efficient and simple mechanism for accessing and retrieving this data. This is achieved by generating metadata, structured data which describes the characteristics of the raw data. The metadata is in the form of XML documents and a part of this project is focussed on developing an XML schema, which defines the structure and content of this metadata in an extensible and scientifically meaningful manner. The metadata is stored in an XML Database server (XDS) and queried using command line tools or the QCD metadata browser, a prototype of which has been already been developed.

QCDGrid is part of the GridPP project, a collaboration of Particle Physicists and Computing Scientists from the UK and CERN, who are building a Grid for Particle Physics.

For further details see:

http://www.epcc.ed.ac.uk/computing/research_activities/grid/qcdgrid/

http://www.gridpp.ac.uk