# GridPP

## The GridPP testbed

**GridPP has built a prototype Grid of more than 1,000 computers, with equipment at many of the particle physics centres in the UK. Over the next three years this will increase by an order of magnitude, as preparations ramp up for the introduction of CERN's next particle accelerator, the Large Hadron Collider, in 2007.**

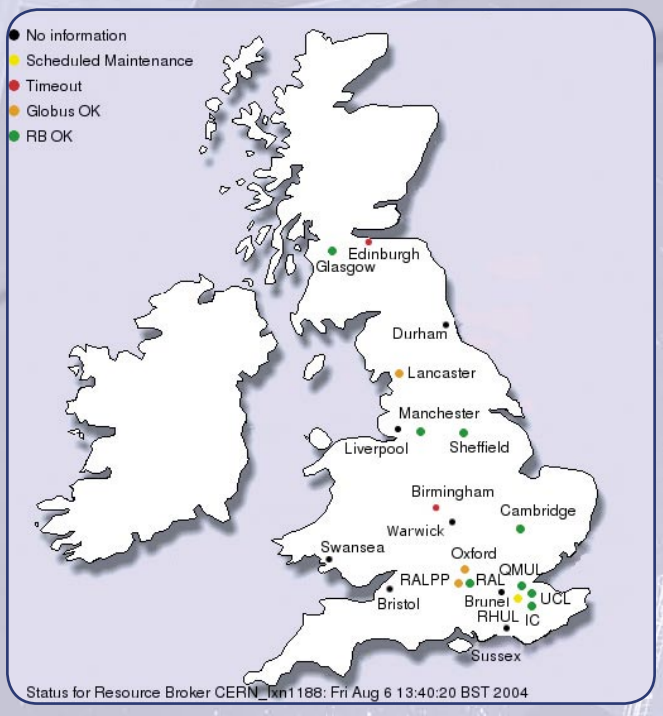


*Monitoring map of the GridPP testbed*

*Monitoring map of the LHC Computing Grid*

### • The LHC Computing Grid

The LHC will be the most powerful accelerator in the world and will produce unprecedented amounts of data, estimated to be around 10 Petabytes per year, for offline analysis by teams of physicists all over the world. To analyse this data and to generate the Monte Carlo simulated data necessary to understand it will require huge amounts of computing and data storage facilities. By 2007, 100,000 computer processors will be distributed across the world, linked together as a Grid. Already, the LHC Computing Grid is the largest functioning Grid in the world, with over 5,000 CPUs and almost 4,000 TB of storage at more than 70 sites around the world. With more than 4,000 jobs being run on the LCG system simultaneously, it is becoming a true production Grid.

### • The GridPP testbed

Currently, 12 sites in the UK are part of the LCG testbed, contributing more than 1,000 CPUs. This GridPP testbed is being developed on a hierarchical model, reflecting the overall structure of the wider LCG testbed. CERN provides the "Tier-0" centre, where the LHC data will be produced. GridPP has contributed £5m to CERN for this, which has been used to support staff and buy hardware. The UK's "Tier-1" centre at Rutherford Appleton Laboratory focuses on data storage and access. In addition there are four smaller, regional, "Tier-2s" in the UK, with a focus on provision of computing power for generating simulated Monte Carlo data and for analysis of data by individual physicists.

PPARC GridPP UK Computing for Particle Physics

## • Tier-1/A –1,000 processors for 1,000 users

The "Tier-1" UK regional centre at RAL also doubles as a "Tier-A" Centre for the BaBar experiment in the USA. It is operated by CCLRC staff and provides GridPP with a large scale computing resource, allowing the UK particle physics community to meet international commitments such as data challenges, production data processing and physics analysis work.

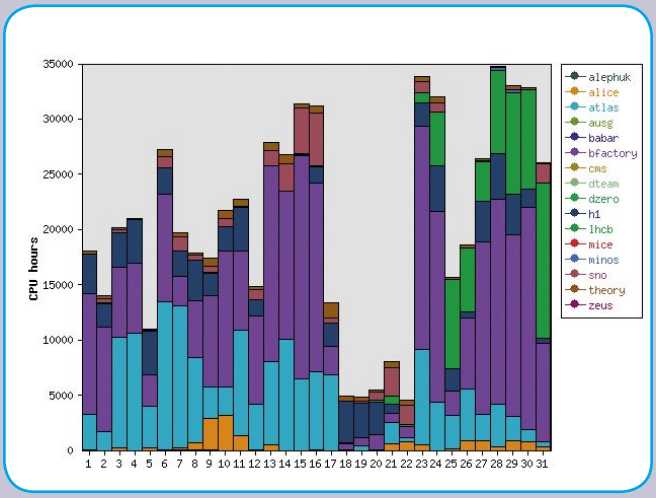


*CPU usage at the Tier-1/A centre by particle physics groups in July 2004*

In June 2004, following a recent upgrade, the service consisted of:

- 1,000 processors, of which around half are for the LCG, and half for BaBar. LCG processors are accessible solely via Grid interfaces. The BaBar service has a number of Grid interfaces, but mainly offers a classical service where users log in to a front end to prepare and submit jobs via a linux batch system.
- A large, highly modular disk service. Its 57 servers give a total available capacity of 200 TB (over 1500 spinning disks). This has been used for a range of services, including UKQCD, CMS, CDF and DZERO.
- The Atlas DataStore (ADS), a networked mass storage system, with a maximum capacity using current tape technology of about 1 Petabyte.
- About 50 other systems which provided various services such as: performance and network monitoring; fabric management and exception monitoring; security logging and scanning; disk space and access to a worldwide distributed filesystem; remote power management; and 40 network switches.

1,000 users are registered on the service and many more have indirect Grid access by virtue of their membership of a Grid virtual organisation.

The centre is usually one of the earliest deployment sites internationally for releases of LCG software and has built up considerable experience operating this technology on the largest scales yet deployed within the UK. Over 40% of the equipment at RAL is accessible solely via the Grid, and significant amounts of real scientific work are being carried out on this service, with over 10 million 1GHz CPU hours used in the last year.

## • Tier-2s

GridPP is also developing four Regional Tier-2s: London, NorthGrid, ScotGrid and SouthGrid. Each of these includes several institutes in reasonable close geographical proximity (see map), and is the major focus for local user support operations, dissemination and outreach activities. Each Tier-2 is associated with at least one Core e-Science Centre and many Centres of Excellence, which encourages continuing close collaboration between GridPP and the Core e-Science programme, with efficient sharing of expertise and resources. Overall, by 2007, the Tier-2s will provide around 6,500 CPUs, with 900TB of storage.

*Racks of computers at Queen Mary, University of London, part of the London Tier 2*

## • The Grid Operations Centre

The Grid Operations Centre (GOC), based at RAL, monitors the operational status of resources deployed internationally through LCG and in the UK through GridPP. Coverage times are improved by sharing this work with Taiwan. Using a secure X.509 certificate based front-end, remote site administrators can maintain configuration details in an SQL database used by the automated monitoring tools. The GOC produces active maps providing real-time information on the state of GridPP and international resources. It also deals with accounting, monitoring the resources consumed by virtual organisations and provided by sites.


*The GridPP Tier 2 structure*