

THE GRIDPP TIER1 CENTRE

Martin Bly, John Gordon, George Prassas, Derek Ross, Andrew Sansum, Steve Traylen, Jonathan Wheeler, Nick White

The Tier1 centre operated by CCLRC staff at RAL provides GRIDPP with a large scale computing resource, allowing the UK Particle Physics community to meet large scale international commitments such as: data challenges, production data processing and physics analysis work. The Tier1 has taken a leading role in the UK for the deployment of GRID testbeds from the European Datagrid project (EDG) and the LHC Computing GRID (LCG) as well as supporting the rollout of GRID middleware to other GRIDPP sites. Many other GRIDPP project deliverables have received infrastructure or technical support from the Tier1 centre.

The centre has usually been one of the earliest deployment sites (internationally) for releases of both EDG and LCG software and has built up considerable experience operating this technology on the largest scales yet deployed within the UK. Access to the Tier1 LCG-CORE service (over 40% of total Tier1 CPU capacity) is now solely via the GRID interfaces and significant amounts of real scientific work are beginning to be carried out on this service.

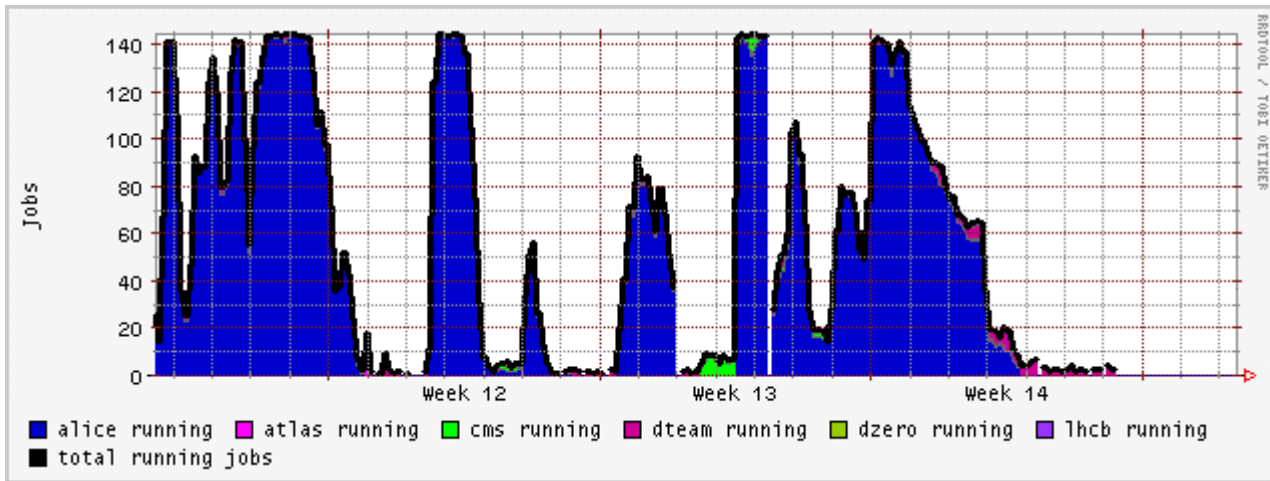


Figure 1: GRIDPP Tier1 LCG-CORE utilisation (March/April 2004)

The UK Tier1 centre has contributed to GRIDPP in many areas, providing computing resources, core infrastructure support and expertise. It has:

- Provided both processing and storage to the worldwide production data challenges of Atlas, CMS, and LHCb
- Been the lead site in the development by GridPP of a persistent grid infrastructure for Particle Physics in the UK.
- Been a key site in both the production and development EDG testbeds.
- Contributed to the development of UK grid infrastructure by developing and running the initial Certificate Authorities (CA) for GridPP and providing this technology to the UK Core eScience Programme.
- Been an early participation in all the LCG testbed releases including the most recent LCG-CORE.

Closely associated with the Tier1 Service is the LCG grid operations centre (GOC) which monitors the status of middleware components deployed throughout the LCG

grid. A secure (https) X509 certificate based front-end (GridSite), allows remote site administrators to maintain resource centre configuration details in a SQL database used by the automated monitoring tools. PBS/LSF filters, provided by the RGMA software infrastructure, allow for the aggregation of accounting information per resource centre and per Virtual Organisation (VO). The GOC also uses the GridIce monitoring tool which is based on Nagios.



Figure 2: LCG Grid Monitoring Map

The Tier1 Centre operates an integrated service both as an LHC Tier-1 Centre for the LHC experiments and as a regional computing centre (Tier A) as part of the BaBar experiment's international distributed computing model. One of the greatest challenges has been to maintain the traditional access mechanisms into legacy part of the service while providing complimentary GRID gateways for a diverse range of projects.

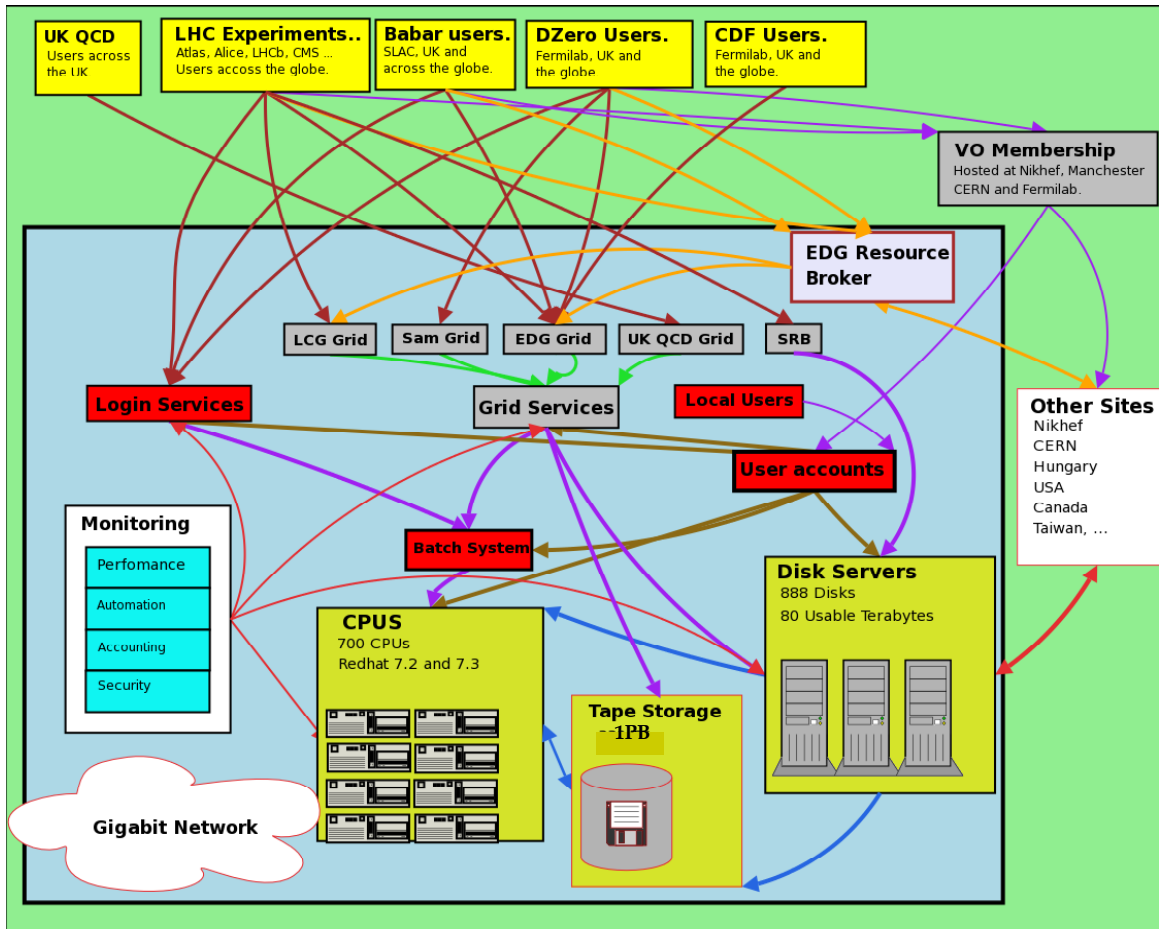


Figure 2: Logical Layout of Tier1 Service infrastructure

This service consists of

- A 700 processor cluster providing roughly 430K SpecInt2000 controlled by the PBS scheduler.
- 80 TeraBytes(TB) of usable disk spac, hosted on commodity PC server and IDE RAID hardware.
- The Atlas Datastore (ADS). A network storage management system backed by an STK Powderhorn robot and IBM 3590, STK 9940A and 9940B tape drives. An SRB interface is provided into the ADS and an SRM interface is currently being tested.
- About 50 other systems which provides various services from Network Monitoring and an AFS Cell through to mini-clusters of Testbeds for the European DataGrid (EDG), LCG and shortly EGEE JRA1.
- There are 50 intelligent 16amp APC power distribution units. There is also a considerable amount of non-intelligent power distribution.

Fabric management has become increasingly complex as the service has grown. By the summer of 2004 this service will have grown to

- A 1000 CPU cluster providing 950 KSpecInt2000.
- 200 TeraByte(TB) of usable disk space.

The staff support over 1000 users directly registered on the service and many more users who have indirect GRID access by virtue of their membership of a GRID virtual organisation.

Over the coming year, access to the service will increasingly be only by the GRID and the service becomes closely integrated into LCG in preparation for the start of LHC data taking in 2007.

