

NorthGrid - A Tier-2 Computing Facility for HEP

CLRC Daresbury
R Allan[?], R Blake, R Tasker

IA Bertram, RWL Jones^{??}, PN Ratoff
University of Lancaster

T Bowcock, PP Allport, MA Houlden[?]
University of Liverpool

R Barlow[?], I Duerdoth, A McNab
University of Manchester

D Tovey[?], P Hodgson, L Thompson
University of Sheffield

The Daresbury Laboratory and the university HEP groups in the North of England are already commissioning large-scale computing hardware for the Grid, with leading involvements in Grid prototype projects. These successful developments have led us to the belief that the North would be well served if our individual facilities were combined into a distributed Tier-2 facility for particle physics, linking existing machines and personnel. We propose to call this NorthGrid. It will serve the HEP groups in the region by efficient exploitation of existing and future resources, and could be the basis for a PetaGrid that will serve the wider community in the North.

The institutions in NorthGrid are involved in all of the major Grid-client experiments except CMS. In particular, all of the groups are in ATLAS, and will be participating in the Grid-based Data Challenges. Lancaster, Liverpool and Manchester are involved in Tevatron experiments and the Gridification of SAM. Manchester and Liverpool are in BaBar and are responsible for much of the BaBar Grid activity. So there is already close interworking between the groups, and well-developed lines of communication between the experts from the different sites. The region includes a regional eScience centre, the eBusiness centre, and one arm of the CLRC eScience centre; the establishment of a Tier 2 organisation will enable existing ties with these agencies to be developed and strengthened.

Our existing resources are already well on the way to Tier 2 scale, and these will be expanded using SRIF2, NWDA and White Rose funds. We anticipate a CPU capacity of the same order as the Tier-1. This will require resources in terms of support effort and some hardware contributions, and a high-performance network between all of the sites: at present some parts of the region are well served but others are not. We request GridPP support in improving this situation, both directly and through lobbying with other agencies. We hope that this will produce a system at the Peta scale, benefiting many academic disciplines in the North, allowing our experience and expertise to be of benefit to the wider community, and will work towards this goal.

Facilities at the individual sites are as follows:

Daresbury is a centre for high performance computing (both COTS and purpose-built) and networking. It recently took delivery of the HPCX facility, an AIX system with 1028 IBM Regatta boxes, 20 Tbytes disk and 30 Tbytes tape. This is the 9th most powerful computing facility in the world, and will be expanded over the next four years. Although intended particularly for close-coupled problems, it is available for particle physics farm-like activities. Daresbury also provides eScience support through the CLRC eScience centre (employing six people), and is extremely active in network monitoring and in middleware development. Surprisingly, it is only served by a 155Mbit/s connection to JANET.

Lancaster is a centre for the Grid-enabling of the D0 Monte Carlo production system and of the SAM system for D0 and CDF. This has already achieved success, with worldwide SAMGrid production demonstrated at SC02. It provides management effort for D0 and for ATLAS computing and Grid activities, and for the ATLAS UK Data Challenges. It has a 200 CPU Computing Farm, with considerable disk space, initially intended for D0 work.

[?] Local contact person

[?] Overall contact

This will be upgraded in two stages to allow full exploitation of the facility with minimal interruption for the D0 work, and to allow testing of SAM/JIM software and EDG tools simultaneously. The tape store can currently hold 26 Tbytes, and could be expanded beyond a Pbyte; this needs support effort and money now the JREI funding is at an end. Lancaster has 1Gbit/s connections cross-campus and up to the SJ4 bar, but is currently limited to 155Mbit/s to the Janet hub; several bids are underway to upgrade this to at least 10Gbit/s.

Liverpool pioneered the COTS computing farm concept with the 300 CPU Monte Carlo Array Processor (MAP). MAP2 is now being constructed with SRIF funds as part of ULGRID. This will be a ~80kSI95 facility, with 120Gbyte local disks for each worker; ready in April 2003. It also has a farm of 80 CPUs for BaBar and an 8 CPU IBM Xeon machine for CDF, and EDG test cluster of 6 machines. The campus is well networked, with 4x155Mbit/s connections to the JANET hub. Liverpool is active on BaBar, ATLAS, CDF and LHCb, and the facilities are used by all experiments as part of their Monte Carlo productions and Data Challenges. It is a centre of expertise in large fabrics and distributed computing.

Manchester is on ATLAS, D0 and BaBar, and is using Grid computing in all three experiments. It also acts as a distribution point for Grid software, and provides vital help on installation and deployment to the region and the country. It has a farm of 80 CPUs for BaBar, a small Grid test bed of six machines, and a farm of 60 CPUs for D0/ATLAS use. These are linked so that it is possible to submit to the two experiment farms using the test bed as a Grid interface. Manchester hosts the VO servers for BaBar and GridPP, and also the GridPP website. All facilities have Gigabit to 100Mbit switches and there is a dedicated Gbit connection to the campus network and to the JANET backbone. Manchester contributes in many ways to the Grid middleware, notably in the areas of networking and security. The HEP group has a close working relationship with the Manchester regional eScience centre.

Sheffield have a HEP Grid cluster of 10 CPUs used by ATLAS, ANTARES and UKDMC and they are expanding this to 50 nodes using rolling grant money. At present, the Grid interface is being tested using a further dedicated three-machine cluster. They also have a 4Tbyte tape robot and a 1 Tbyte disk array. Sheffield also provides access to the distributed White Rose Computing Grid facility with machines funded by SRIF and RDA funds. This includes a 256-processor Linux cluster at Leeds and high performance shared memory machines at Sheffield, York and Leeds, all potentially available for Grid submission of HEP jobs. Sheffield HEP expertise is being used in the Grid integration of White Rose. There will be a SRIF-funded Gigabit upgrade to the campus network although JANET connectivity remains at only 155Mbit/s.

In summary, we feel that the formation of a NorthGrid Tier-2 'centre' from our existing facilities is the sensible and inevitable path forward, providing a really useful resource for HEP. Beyond this, we hope that it will also provide the basis an interdisciplinary facility for the North, allowing the PPARC developments to be disseminated to and recognised by a wider community. NorthGrid is supported by our existing experimental collaborations, and will help fill in the gaps in resources that presently limit the exploitation of the equipment and expertise already existing in our groups.