

A Southern Tier-2 Computing Centre for HEP

DRAFT2 – 16th August 2002

D. Bailey, N. Brook, D. Newbold[†]
University of Bristol

F. Harris, T. Huffman[†], P. Jeffreys, I. McArthur
University of Oxford

D. Kelsey[†], R. Middleton, G. Patrick
Particle Physics Department, RAL, CLRC

Over the past two years, the particle physics groups at the Universities of Bristol and Oxford, and at the Rutherford Appleton Laboratory, have each established a strong local programme of Grid-based computing in direct support of their experimental activities. In each case, the development of a high-performance local computing system has been accompanied by a leading involvement in UK, European and Worldwide Grid prototype projects. We now feel that there is a valuable opportunity to combine our expertise and resources to provide a prototype Southern Tier-2 computing centre for particle physics. The goals of the project are: to investigate the concept of a Tier-2 centre and define what role the centre should play in the UK; to increase the efficiency with which existing resources may be used; and to strengthen links with our core e-Science centres and other potential collaborators.

The establishment of a Tier-2 centre will offer important benefits to the experimental collaborations participating within GridPP, and to the project as a whole. Almost the entire range of current and future UK HEP experiments is represented within one or more of our groups, all of which will gain access to the Tier-2 resources. Highly experienced computing personnel are in post at each group, whose areas of expertise are in many cases complementary. The overall system management and operations manpower requirement will be reduced through close cooperation between sites, and through collaboration with local e-Science centres. The active collaborations that already exist between the three groups, for example through the LHCb Grid computing programme, will be enhanced by the provision of a dedicated regional test platform for distributed computing, and there will be new opportunities for collaboration between personnel associated with other experiments.

The proposed Tier-2 centre will be distributed in nature, consisting of hardware and manpower at our three sites. Computing resources will include both CPU power and high-performance disk storage. Redundancy amongst sites will ensure a high availability of resources. The operation of such a ‘virtual’ computer centre relies upon excellent wide-area network connections between sites, and it will therefore be vital to ensure that network performance continues to keep step with the evolution of the local resources. The centre will initially consist of existing or planned facilities at the three sites. In the run up to the LHC era, the scale of the centre will expand to match the increased requirements of the experiments; it is likely that the final scale will correspond to roughly a quarter of the processing and storage capacity of the UK Tier-1/A centre.

We estimate that roughly four FTE of effort will eventually be required to operate the centre, one FTE being a dedicated Tier-2 manager, with the remaining effort divided between system management and operations at the three sites. Some of this effort is already in place, with the remainder to be sought through GridPP and through collaboration with our local e-Science Centres. Wherever possible, we will develop system management tools and methods in cooperation with the Tier-1/A centre, building on the close collaboration between personnel at our institutes and the CLRC e-Science team.

[†] Contact persons

At the University of Bristol, the particle physics group has recently benefited from JREI and JIF awards, and from industrial sponsorship. Along with PPARC and Departmental funding, this has enabled the deployment of some 150 modern CPUs and around 5TB of high-performance disk storage. This facility has contributed at over 80% efficiency for the last year towards the simulation and analysis programmes of the BaBar, CMS and LHCb experiments. The group is shortly to benefit from an upgrade to the University network, which will provide a dedicated 622Mbit/s connection to JANET; this will prove invaluable in the operation of a distributed Tier-2 centre. It is planned to further expand these computing facilities, both through additional University and PPARC funding, and through more effective coupling to other large computing facilities within the University. Bristol personnel are heavily involved in both management and operational aspects of the UK HEP Grid. The group has been awarded two GridPP posts and two e-science studentships, and currently expends an additional 0.5FTE of system management effort on Grid-related activities. Alongside computing experts from the LHCb, CMS and BaBar collaborations, group personnel have already made essential contributions towards experimental data challenges and in the roll-out of the first EDG grid software, and their expertise will be essential in setting up the distributed Tier-2 centre.

Oxford University has been awarded two GridPP posts and an e-science studentship to grid-enable systems based on the GAUDI/GANGA framework for LHCb and ATLAS and a further GridPP post and studentship to make the CDF analysis system grid-aware. In addition, the University Division of Physical Sciences has authorized a new ATLAS Grid Lecturer for the group to be filled in the autumn of 2002. Oxford was awarded an e-Science Centre, which has successfully attracted industrial funds and partnership. Oxford has also benefited from a £500K JIF award in computing and an SRIF award to improve the University network. The Particle Physics group anticipates working very closely with the e-Science Centre and will look to them to provide the operations and hardware support required by Tier-2 facilities. The Grid project has also promoted closer ties with the Oxford computer science groups who now play an important role in supervising and training our grid personnel. By 2003 the Oxford group anticipates having 10TB of RAID disk, and between 20 and 30 CPU's available for compute intensive work. The group expects to also receive a share of the centrally supported e-Science supercomputer, which has 300 CPUs and the associated support staff. Due to the presence of the e-Science centre in Oxford, a strong case can be made to the SRIF funds for a substantial increase in data processing and storage infrastructure. These proposals will be coordinated with the other members of the distributed Tier-2 centre.

Members of the Particle Physics Department at CLRC/RAL are active participants in GridPP across all areas of Grid computing. The department has been awarded 5 GridPP posts. The physicists in the department, just like those in the universities, will require access to a Tier-2 centre to enable them to contribute fully to physics analysis. The department has purchased a rack of 7 dual-processor systems plus a 1TB disk server, as an initial contribution to this project. This hardware will be upgraded as required by the ongoing PPD physics / Grid programme, and as funds allow; by the end of 2003, at least 50CPUs and 5TB of storage will be available. The RAL connection to the Thames Valley Network MAN has recently been upgraded to 2.5Gbit/s and there is a Gigabit backbone across site. A more important contribution than hardware, however, is the fact that experienced staff are willing and able to work on this project. This includes both system management effort from the PPD Computing Group, and experiment / application effort from, for example, LHCb who have already started exploring how potential Tier 2 centres integrate into the Grid. Other experiments in PPD, ATLAS, CMS and BaBar, will also participate in this project.

In summary, we feel that there is an excellent opportunity to combine the expertise and resources of our three groups, to provide a facility of both immediate and long-term benefit to UK particle physics. The proposed Tier-2 centre will directly assist a wide range of UK experimental collaborations in the achievement of their science goals. Moreover, it will provide a focus for future collaboration in the area of high-throughput Grid computing, and provide a valuable opportunity to share expertise, infrastructure and resources with the local e-Science centres, and potentially with other interested partners.

