# Australian Experimental High Energy Physics Community Activities

**Geoff Taylor**

## Background:

The Australian Experimental High Energy Physics (HEP) community is an active participant in two of the world's largest international scientific collaborations. The ATLAS collaboration is currently constructing the ATLAS detector on the Large Hadron Collider (LHC) at CERN, Geneva, and will be completed by 2007. The primary motivation for this experiment is the discovery of the Higgs particle which is thought to be responsible for mass in ordinary matter. In addition to the search for the Higgs, the energies created at the LHC will also allow ATLAS to probe further into the structure of matter than before. ATLAS is expected to produce in excess of 10 PB of data a year which will be distributed internationally and accessible to the 2000 scientists from 150 institutions in 34 countries. The Belle collaboration has been studying data since 1999 from a running detector situated at the KEK B-Factory in Japan. It is investigating a fundamental violation of symmetry in nature which may help explain the universal matter-antimatter imbalance. The Belle collaboration has produced over 1000 TB (1 PB) to date and is collecting data at an ever increasing rate.

The data processing and access patterns for experimental high energy physics is somewhat different to that of other research areas. Many data intensive eResearch efforts focus on data discovery: selecting individual data sets of interest from many distributed data sets indexed in digital registries or meta-data catalogues (eg. International Virtual Observatory and many Open Archives Initiatives projects). HEP applications generally process and scan the entire communities distributed data sets for features of interest extracting statistical summaries for further processing. This type of processing requires not only access to the data sets but the ability to deploy code designed by the researcher on high performance computing attached to the data storage. The breadth of research within each community excludes the possibility of providing common filtering or transformation services attached to the data. “The Grid” and in particular Data Grid techniques are of particular interest to the community as a solution to our data access requirements. The international HEP community is pouring much effort into the development of Data Grid middleware and is driving the deployment of internationally interoperable infrastructure.

## Past Efforts and Current Status:

To investigate the existing Grid and e-Science software and infrastructure, in early 2003 we constructed an Australian distributed high performance test-bed. It's initial construction occurred over a period of 9 days. The test-bed consists mainly of hardware donated by IBM Australia in collaboration with IBM Asia-Pacific, with version 2 of the Globus Toolkit used as standard middleware. There are 5 nodes making up the test-bed and these are all used for processing and data storage. The nodes are located at the University of Melbourne, the APAC National Facility at the ANU in Canberra, SAPAC at the University of Adelaide, and at the University of Sydney. The APAC National Facility mass storage system was also used for storage of data. There are a number of central services which help to maintain and utilise the test-bed, such as a Certificate Authority (CA), a Grid Index Information Service (GIIS), and a Globus Replica Catalog.

Belle simulation and analysis using the test-bed was successfully demonstrated at the 4th PRAGMA workshop and ICCS2003 conferences, and at the international Super Computing conference (SC2003) and associated Grid 2003. The test-bed was also utilised in the high performance computing challenge at SC2003. More recently our experience of the Grid through this test-bed has been presented to the ATLAS collaboration in CERN, at a Global Grid Forum in 2004, and to the Belle collaboration and KEK facility in Japan. This "High Energy Physics Test-bed" continues to be maintained and used for developing and testing new technologies. Access to such a distributed test-bed has been essential for our research into Data Grids and has allowed existing production resources to be unaffected and used in parallel.

In 2004 the Belle Experimental Collaboration reached a critical stage in their computing requirements. Due to an increased rate of data collection an extremely large amount of simulated (Monte Carlo) data was required to correctly analyse and understand the experimental data. The resulting simulation effort consumed more CPU power than was readily available to the experiment at the host institution, KEK, Japan. In order to meet requirements the simulated data production was distributed to remote institutions who were able to contribute CPU power. The Australian Belle collaborators participated in this production successfully utilising resources at number of Australian facilities, including APAC (Australian Partnership for Advanced Computing), AC3 (Australian Centre for Advanced Computing and Communication), University of Melbourne ARC (Advanced Research Computing centre), and VPAC (Victorian Partnership for Advanced Computing). This production involved the use of a globally accessible data catalogue and resource management system, SRB (Storage Resource Broker), and tools developed in-house for the central dispatch and monitoring of jobs.

In 2005 the production has continued and is now utilising the newly implemented APAC National Grid infrastructure. The Australian high energy physics community has played an active role in the deployment of Grid services within APAC, the associated partners, and at the University of Melbourne. The software used for this Grid production effort is SRB for data management and LHC Computing Grid (LCG) middleware from the Enabling Grids for E-sciencE (EGEE) project.

## Requirements for the Near Future:

The LHC Computing Grid (LCG) will provide the computational, storage, and network resource necessary for the management and processing of data from the ATLAS collaboration. The LCG is the leading infrastructure deployment effort within the Enabling Grids for E-sciencE (EGEE) project funded by the European Commission. The projected cost of the ATLAS requirements for the LCG is 35 million CHF in hardware alone. A large portion of this cost will be in-kind contributions from regional facilities and funding bodies. We plan to establish Australian EGEE/LCG Data Processing centres to be integrated with the international EGEE/LCG infrastructure.

A major function of the EGEE/LCG infrastructure is to disseminate and provide access to distributed collaborative data within 48 hours of production. The LHC data is collected from experiment then distributed internationally for further processed and filtering before researcher consumption. Simulation of experimental data is computationally intensive but necessary for understanding efficiencies and uncertainties in analysis. In total, the ATLAS collaboration expects to generate more than 10 PB of data. The Replica Location Service (RLS) data management software will be used throughout the LCG to distribute and access this data. RLS will provide provide a globally accessible logical file catalogue overarching many distributed file catalogues and storage resources.

To better define levels of hardware and support contributions the LCG have categorised facilities into "Tiers". Tier 1 denotes international facilities storing the majority of the ATLAS data set and providing a high level of support to a number of Tier 2 facilities. Tier 2 denotes large national level facilities storing mainly summarised data sets, but will be used for the bulk of common processing tasks. Tier 3 usually denotes institution level resources for the sole use of local collaborators tuned to their specific research needs. These facilities must expand over time to match the annual increase in ATLAS data. An Australian Tier 2/3 facility would be essential for rapid access to the ATLAS data and continued research outcomes. By 2007 an estimated 30 Tier 2 facilities will be contributing world wide. A typical start up Tier 2 facility in 2007 will consist of 66 kSI2k (thousand SpecINT 2000) and have a storage capacity of 46 TB. This equates to roughly 70 modern Pentium4 processors. The Tier 3 facility requirements are expected to be similar but can be tailored to individual research groups. By 2008 this facility must increase in processing power and storage size 6 times then nearly double each subsequent year.

To support the Australian ATLAS effort and the wider collaboration we are in the process of establishing an Australian Tier 2/3 facility to peer with the Grid Operations Centre and Tier 1 at Academia Sinica, Taiwan. An Australian Tier 2/3 facility is essential for rapid access to the ATLAS data and continued research outcomes. A Tier 2 facility must have direct access to advanced research networks that peer with international networks, Tier 1 facilities and the experiment at CERN. The speed and reliability of these networks directly affects data replication and the ability of researcher to collaborate and contribute

internationally. It is also sensible to co-locate both Tier 2 and Tier 3 resources providing direct access to Tier 2 storage for researcher analysis. Following the recent APAC National Grid Program, in which we aided in the deployment of an EGEE/LCG infrastructure layer, we intend to integrate this facility with Australian Grid infrastructure utilising existing services where possible.

## Future of Information Management:

One of the larger problems being addressed by eResearch initiatives is information management. All international collaborative science projects which involve sharing large data sets will encounter difficulties with information dissemination. The larger and more distributed the collaboration the more difficult it is to define common semantics/terminology for describing data sets and the tools to analyse them. Historically, in experimental HEP, we have stored a great deal of "meta-data" in community databases but some information remained embedded in log books, file names, directory structures, and "common" knowledge. Information is often disseminated by presentation at meetings and private communication. For large international collaborations this form of communication is essential but not the most efficient way of rapid dissemination or preserving community knowledge for future collaborators.

As a member of the Belle collaboration we are currently investigating meta-data standards and schema for the description of the Belle data sets. Belle has accumulated 1000 TB of data both experimental and simulated. To reduce the burden on the researcher meta-data must be automated where possible. It is our intention to provide a system to capture:

- Meta-data descriptions at several levels (Item, Collection, Task and Research)
- Provenance/History meta-data (capture workflow and how the data was derived)
- Structural meta-data (format information helpful for determining processing tools)
- Discipline based meta-data (describing the research undertaken that lead to the data produced)
- Researcher annotation meta-data (allowing individual and community evaluation of data quality and usability)

It is expected that a lot can be gained by capturing the workflow and task descriptions of researchers as they produce data. Common tasks could be more easily identified; typical or even incorrect use of software tools could be determined or rectified; community expertise can be expressed and searchable. The nature of some research tasks within the collaboration may mean that data sets become less valuable over time, however, the long term preservation of associated meta-data is essential. Meta-data can allow us to regenerate the data if lost or deleted and preserves the history and progression of the research, analysis tools, and community expertise.

As a part of this effort we are investigating the use of several standards to improve future access of this data set and leverage international infrastructure. These include the universal Dublin Core meta-data standard (DC) and the CCLRC Scientific Metadata Model (CSMDM, defined by the Council for the Central Laboratory of the Research Councils from the UK e-Science program). We hope to build on these base meta-data frameworks to construct a Belle community specific meta-data schema. It is our intention to support the Open Archives Initiative's Protocol for Metadata Harvesting (OAI-PMH) to provide national and international access to the data description. This protocol may be used to federate distributed meta-data repositories within the Belle consortium. In the long term future this could promote future research projects that span experiment based and collaboration datasets, reducing the boundaries to cross collaboration.