

## Appendix V APAC, the National Facility, the National Grid<sup>6</sup>

APAC was established to strengthen Australia's advanced computing capabilities by installing a peak computing system at the APAC National Facility and by increasing the national expertise to support users of advanced computing systems.

Over the last three years the partnership has grown in expenditure from \$33M pa to \$45M pa, by research users from around 1200 to nearly 2000 and in staff either employed or supported from around 90 to 168.

The context for the second stage of APAC (2004-06) was influenced by the increasing demands for advanced computing services in the Australian research community and by the major initiatives related to 'cyber-infrastructure' in many other countries.

These demands were being addressed by major government initiatives in many other countries such as the USA, nations of the European Union (especially the UK), Japan, Korea, Taiwan and Singapore. The initiatives were focussed on the provision of advanced computing facilities, grid services and broadband communications to support the emerging eResearch 'revolution'.

The need for Australia to participate in this global revolution guided APAC to develop a mission around seamless access to its facilities as part of the global research infrastructure, which was expressed as:

*'Providing the national advanced computing and grid infrastructure for eResearch'*

APAC therefore supports researchers by providing a peak computing service based on the National Facility and grid services that enable seamless access to distributed computing and data management facilities through the National Grid.

The 'open' architecture for the National Grid enables access to institutional resources, on-line instruments, sensors and other data sources and enables interoperability with grids in institutions and other countries. The impact of this architecture is being seen in recent discussions with organisations on incorporating their facilities into the National Grid.

The major sources of funding into these initiatives over the last three years is indicated to the right, noting the difference in the Total Cash Income and Expenditure represents a cash carryover of around \$2m at the end of 2006.

<b>TOTAL for 2004-2006</b>	<b>CASH</b>	<b>IN-KIND</b>	<b>TOTAL</b>
	(\$'000)	(\$'000)	(\$'000)
<b>INCOME</b>			
<b>SII</b>	29,000		29,000
<b>Partners</b>	9,000	24,730	33,730
<b>Other</b>	4,200		4,200
<b>TOTAL</b>	42,200	24,730	66,930
<b>EXPENDITURE</b>			
<b>National Facility Program</b>	19,810	6,975	26,785
<b>CT&amp;T Program</b>	3,320	3,150	6,490
<b>Grid Program</b>	11,800	12,400	24,200
<b>EOT Program</b>	3,000	2,205	5,205
<b>Management</b>	2,170		2,170
<b>TOTAL</b>	40,120	24,730	64,850

While the remainder of this summary focuses on the National Facility and the National Grid, APAC has also operated a Computational Tools and Techniques Program (the CT&T line in the above table) and an Education, Outreach and Training program (the EOT line in the above table).

### National Facility

The National Facility operates on a resource share model whereby APAC, its partners and other organisations are able to obtain resources on the National Facility in proportion to their cash and in-kind contributions. Under the resource share model, APAC's funding to the National Facility purchases high end compute resources available to all Australian researchers through a Merit Allocation Scheme.

The National Facility is also supported by 15.5 effective full time staff. The roles of the staff are to manage the facility, to help users optimise codes, organise software and the user environment, perform system development and administration, develop training courses, and to service the granting committees.

<sup>6</sup> These notes are extracted from the extensive documentation provided by APAC to the APAC review panel.

The current major system provided by the APAC National Facility is an SGI Altix comprising:

- 1928 1.6Ghz Itanium2 processors (1920 dedicated to batch jobs) with 5.6Tbytes main memory
- SGI's NUMalink4 interconnect providing 3.2 Gbytes/s bandwidth per link and <2us MPI latency
- 30 Tbytes of global file system with an aggregate bandwidth of > 2Gbytes/s
- A total of over 70 Tbytes of scratch disk local to the compute nodes

In addition, the National Facility operates a small 156 processor cluster fitted with 1.4Tbytes of shared file space and 15 Tbytes of scratch space, intended to take small jobs off the main system. The load shift has been highly successful and an upgrade to a larger cluster is foreshadowed.

The National Facility also provides a large Data Storage System for researchers which is also in the process of being upgraded to include:

- Large SAM-QFS cluster managed HSM system for data management
- 2 Sun Opteron V40z providing metadata servers to the cluster file system and management of the HSM
- Large, expandable online disk storage for projects requiring more on-line data
- Migration to tape storage for large or deep storage using existing data robots and updated tape capacity
- A data intensive computational cluster with mounted SAM-QFS file system
- A cluster of systems dedicated to database access for SQL services and fast index searching
- Scalable front end to enable portals for data management
- Scalable high speed data transfer for grid and streaming services

APAC supports access to the data store for Data Intensive projects on a merit basis, with the following current aggregate allocations:

Projects	Requests (TBytes)	Maximum Usage (TBytes)
APAC DI Projects	72.52	39.68
APAC + Partner Projects	78.93	39.93
Total	151.45	79.61

## The National Grid

The main objective of the APAC Grid is to develop, install and operate a grid infrastructure across the APAC and partner facilities to:

- allow users easier access to the facilities through secure mechanisms,
- allow users to access the most appropriate computing and data storage system in the National Facility,
- provide users with grid services to support dynamic virtual organisations, and
- support users participating in national and international research programs.

The National Grid has been designed as an integrated virtual system with interfaces to external systems and grids; such as institutional computing systems, on-line instruments, sensors and data sources. It is also designed for inter-operation with other grids at the institutional, national and international levels.

Researchers are able to see the resources in the APAC National Grid as an extension of their personal work environments; accessible through a common security system and via portal interfaces. In this sense, the National Grid is intended to become an infrastructure that is transparent to most of its users.

The architecture of the National Grid involves a 'gateway' system at each partner site configured to support a range of grid services and to receive and process grid service requests. The aim of the architecture is to:

- limit the number of systems that need grid components installed and managed thereby reducing overall grid management overheads

- enhance security as many grid protocols and associated ports only need to be open between the gateways and as only the local gateway needs to interact with site systems
- support roll-out and control of production grid configuration through the implementation of standardised grid support across all grid sites
- support production and development grids and local experimentation without significant hardware investment through a Virtual Machine implementation where different services and different quality of services are provided on separate grid installations

The APAC Grid program has engineered, installed and operates grid infrastructure at all APAC partner sites. This infrastructure supports:

- internationally accredited PKI certificates, shared myproxy services, user controlled DN to local account mapping and VOMS/VOMRS based project grouping
- remote data staging and job submission to all significant compute resources within the partnership using gram2, gram4 and LCG protocols
- direct and third party transfers between all sites using GridFTP
- multi-site shared data using federated SRB servers
- portals based on Gridsphere and related tools

The APAC grid middleware stack is based on the Virtual Data Toolkit used in the US Open Science Grid (which incorporates the Globus Toolkit). This deployment has provided the basis for conforming to the emerging international standards for grid services and for interoperation with large-scale grids being developed for researchers in other countries.

It has also facilitated participation in international forums in order to contribute to the development of grid middleware and to gain experience with the early implementation of standards. The participation is fostering cooperation on grid technologies, standards and applications with groups in the US, UK and Asia-Pacific.

Examples of application services developed include:

- workflow scheduling for genome annotation at Murdoch University (not an APAC site) computing using resources at iVEC and SAPAC
- workflow scheduling for geomechanics from iVEC to resources at iVEC and CSIRO HPSC, staging data using federated SRB
- LCG-managed job submission from Melbourne University (not an APAC site) to resources at VPAC, SAPAC, ac3 and ANU
- Belle experiment management using federated SRB to import data internationally, the National Grid to manage jobs and SRB to export data internationally
- digital repositories and portal access solutions for marine and astronomical data sets
- Web-based access to chemistry packages for molecular modelling.

The following application and infrastructure projects are in progress using the APAC Grid.

#### Astronomy

- Astronomy Data Warehouse: Katherine Manson, University of Melbourne. Providing efficient, standardised access to key Australian and international astronomical data collections, providing the ability to discover and query archives and to extract data from those archives.
- Astrophysics Portals: Sarah Maddison, Swinburne University. Providing user-level services for job configuration, submission and monitoring using Gridsphere portlets for theoretical astrophysics codes.
- Gravitational Wave Research: Susan Scott, ANU. Providing the Australian Consortium for Interferometric Gravitational-wave Analysis (ACIGA) to establish a National Gravitational Wave Research Grid which to enhance the collection and analysis of data within Australia, the sharing of this data with international collaborators, and the enhancement of data access and retrieval from these collaborators.

### Bioinformatics

- Bioinformatics Grid Services, Kimberly Begley, Griffith University. Deploying grid-based services for local implementations of major overseas databases and their associated tools.
- Genome Annotation: Matthew Bellgard, Murdoch University. Providing a grid-enabled Blast system; installing a local Ensembl database and the relevant compute/data intensive tools and developing a distributed Genome Annotation System using the rice genome as a model.
- Molecular Docking: Rajkumar Buyya, University of Melbourne. Providing Web-based services to molecular docking applications and a wide variety of chemical databases, as well as tools to assist in the analysis and archival of the screening results.
- Parasitology Grid: Tim Langtry, University of Technology Sydney. Providing services for the sharing and remote processing of biological data for the ARC Network for Parasitology, by adapting tools developed in the European myGrid program (in particular Taverna).
- pBlast: Albert Zomaya, University of Sydney. Providing grid-based services for Web-based BLAST searches, query and annotation commands and bioinformatics portals.

### Chemistry

- Chemistry Portal: Andrey Bliznyuk, ANU. Providing access to advanced computing systems for the computational molecular science community through portal development.

### Earth Systems

- Digital Libraries and Grid Computing: Glenn Hyland, University of Tasmania. Enabling Australian data sets related to oceans, atmospheres, Antarctica and the climate to be made transparently available to the Australian and international research community using grid protocols and standards.

### Geosciences

- Earthbyte: Dietmar Muller, University of Sydney. Providing the foundation for grid-based data access and earth process modelling by linking geological and geophysical observations to palaeogeographic models for constraining mantle convection and lithospheric deformation.
- Geosciences Workflow: Robert Woodcock, CSIRO. Providing workflow services for the geoscience community to access geoscience data sets and modeling software.

### High-Energy Physics

- Belle and ATLAS Experiments: Glenn Moloney, University of Melbourne. Deploying a common data grid infrastructure allowing participation in two international frontier experiments in High Energy Physics: the Belle experiment at the KEK laboratory in Japan; and the ATLAS experiment at the CERN laboratory.
- International Lattice Data Grid: Paul Coddington, University of Adelaide. Providing services for the International Lattice Data Grid (ILDG), that aims to allow highenergy physicists to publish, locate and access hundreds of Terabytes of data from computationally intensive lattice Quantum ChromoDynamics (QCD) simulations.

### Infrastructure Projects

- Computing Infrastructure: David Bannon, VPAC. Deploying the computing systems infrastructure required for an operational APAC Grid, including services for authentication, authorization and job management.
- Information Infrastructure: Ben Evans, ANU. Providing services for an integrated national data grid including data transfer, data federations, information services, and access control.
- User Interface and Visualisation Infrastructure: Raj Chhabra, Queensland University of Technology. Enabling easy access to the National Grid through application portals and visualisation tools for multiple simulations and large-scale data sets.
- Collaborative Support Infrastructure: Chris Willing, University of Queensland. Enhancing Access Grid functionality by adding visualisation and usability enhancements, in order to support collaborative working.